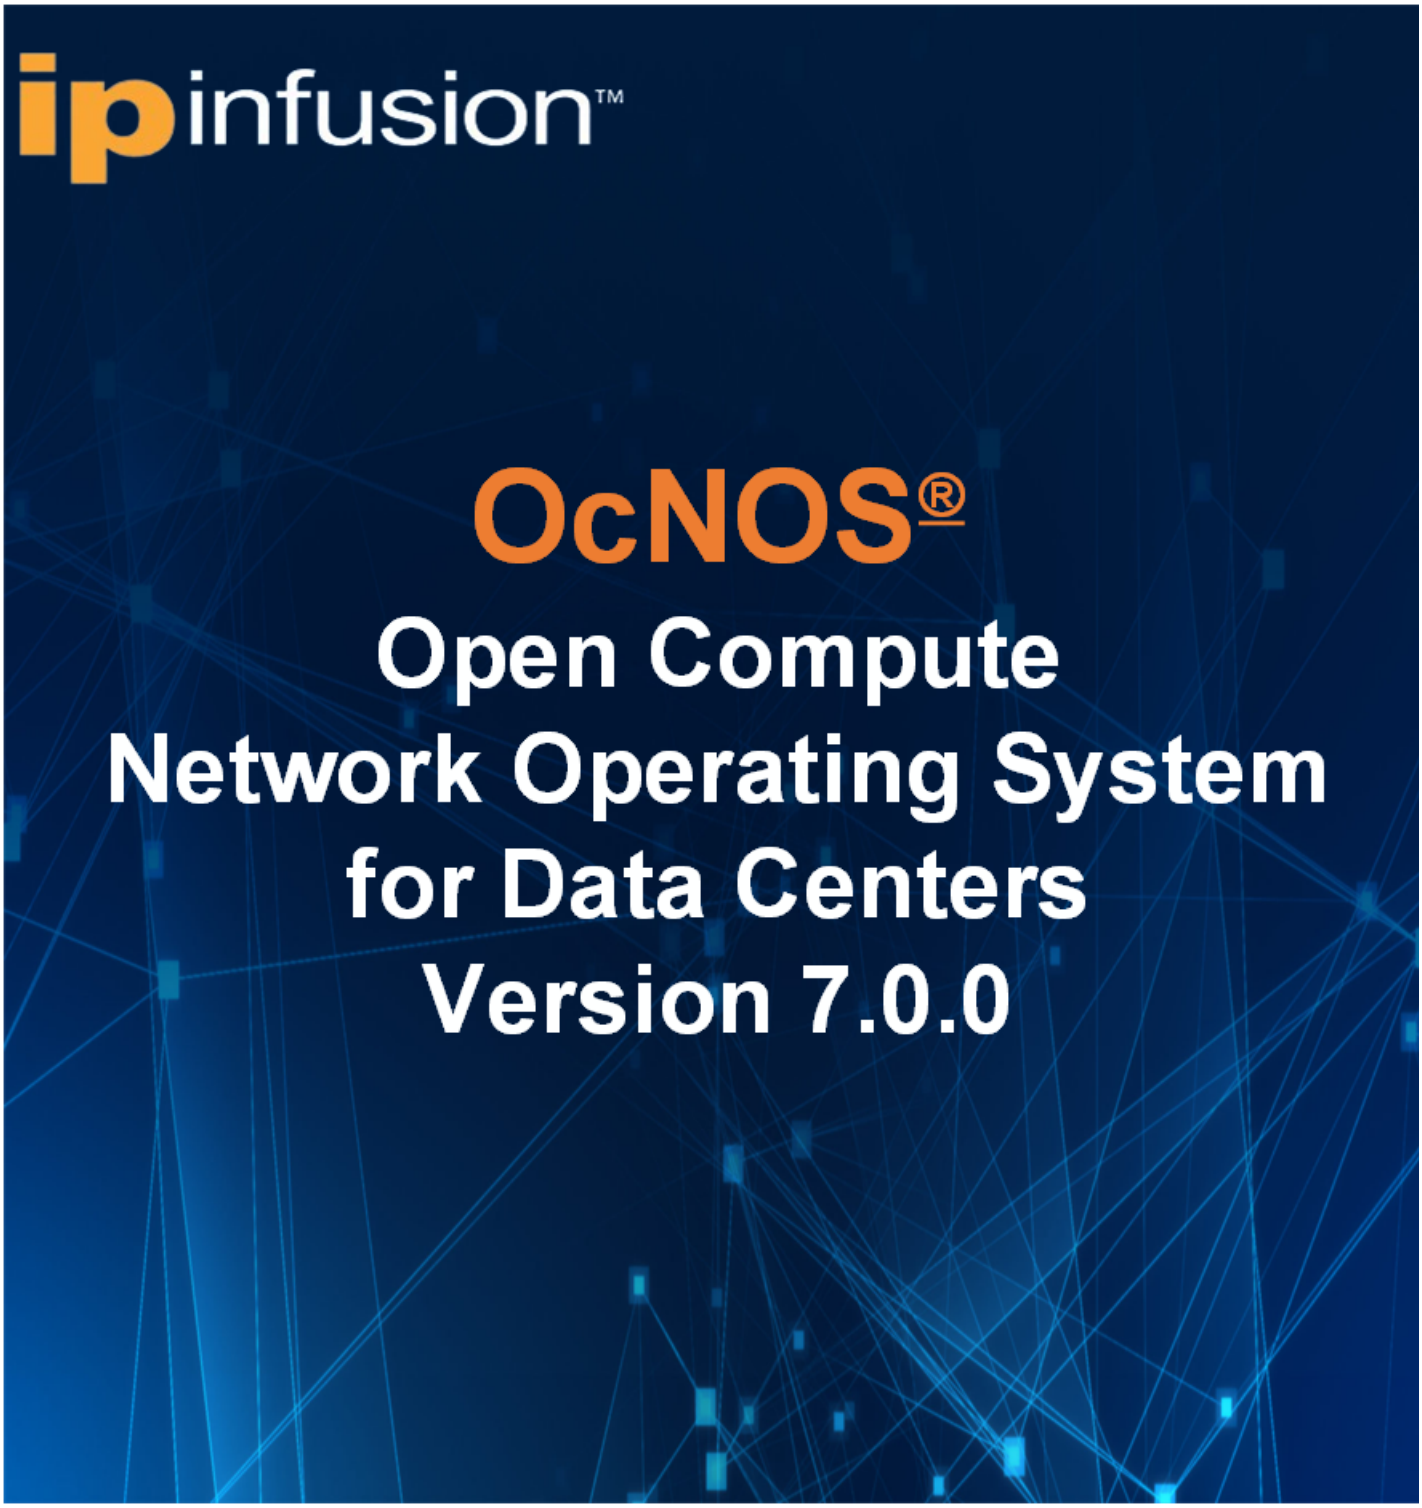

Virtual eXtensible Local Area Network Guide

February 2026

IP Infusion Inc.

3979 Freedom Circle

Suite 900

Santa Clara, California 95054

http://www.ipinfusion.com/

For support, questions, or comments via E-mail, contact:

support@ipinfusion.com

Contents

Preface

This guide describes how to configure OcNOS.

Audience

This guide is intended for network administrators and other engineering professionals who configure OcNOS.

Conventions

Table 1 on page 9 shows the conventions used in this guide.

Table 1: Conventions

Convention	Description
Italics	Emphasized terms or titles of books
Note:	Special instructions, suggestions, or warnings
`monospaced type`	Code elements such as commands, parameters, files, and directories

IP Infusion Product Release Version

Each integer in release number indicates Major, Minor, and Maintenance release versions. Build numbers that follow the release numbers are for internal tracking and verification of the software build process and are visible to customers as part of the software version number.

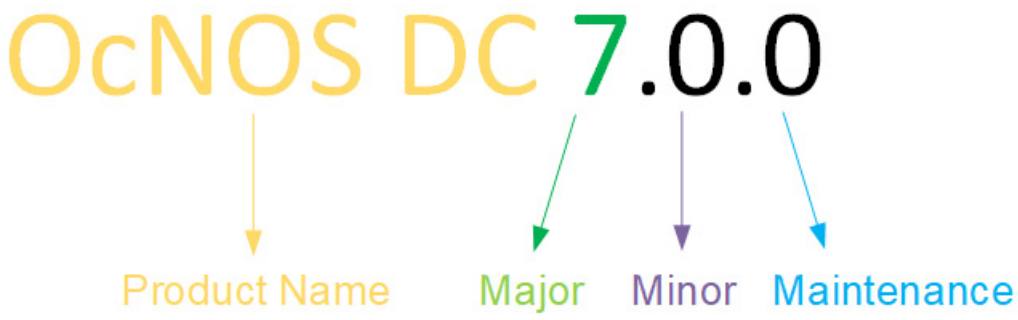

Product Name: IP Infusion Product Family

Major Version: New customer-facing functionality that represents a significant change to the code base; including, a significant marketing change or direction in the product.

Minor Version: Enhancements or extensions to existing features, changes to address external needs, or internal ements might be motivated by improvements to satisfy new sales regions or marketing initiatives.

Maintenance Version: A collection of product bugs or hotfixes usually scheduled every 30 or 60 days, based on the number of hotfixes.

Related Documentation

For information about installing OcNOS, see the *Installation Guide* for your platform.

Feature Availability

Each OcNOS SKU contains a set of supported features. For a list of available features based on the SKU that you purchased. Refer to the *Feature Matrix*.

Migration Guide

Check the *Migration Guide* for necessary configuration changes before migrating from one version of OcNOS to another.

IP Maestro Support

Monitor devices running OcNOS Release 6.3.4-70 and above using IP Maestro software.

Technical Support

IP Infusion maintains an online technical support site that provides a variety of technical support programs for licensed OcNOS customers at the Technical Assistance Center.

Customers and partners enjoy full access to the support website. The site allows customers and partners to open technical support calls, update open calls with new information, and review the status of open or closed calls. The password-protected site includes technical documentation, Release Notes, and descriptions of service offerings.

Technical Sales

Contact the IP Infusion sales representative for more information about the OcNOS solution.

Technical Documentation

For core commands and configuration procedures, visit: Product Documentation.

For training videos, visit: OcNOS Free Training Videos.

For a list of supported platforms and SKUs of OcNOS features, refer to the OcNOS Feature Matrix.

Disclaimer

The global documentation site is evolving to provide an enhanced website user experience for select topics included in this release. Some guides are now available outside the existing documentation library and can be accessed directly from custom documentation landing pages. These guides offer robust in-built search functionality.

For the latest documentation, visit the product-specific documentation landing page and select the relevant guide.

Comments

If you have comments, or need to report a problem with the content, contact techpubs@ipinfusion.com.

Command Line Interface

This chapter introduces the OcNOS Command Line Interface (CLI) and how to use its features.

Overview

You use the CLI to configure, monitor, and maintain OcNOS devices. The CLI is text-based and each command is usually associated with a specific task.

You can give the commands described in this manual locally from the console of a device running OcNOS or remotely from a terminal emulator such as `putty` or `xterm`. You can also use the commands in scripts to automate configuration tasks.

Chapter Organization

The chapters in command references are organized as described in Command Description Format.

The chapters in configuration guides are organized into these major sections:

- An overview that explains a configuration in words
- Topology with a diagram that shows the devices and connections used in the configuration
- Configuration steps in a table for each device where the left-hand side shows the commands you enter and the right-hand side explains the actions that the commands perform
- Validation which shows commands and their output that verify the configuration

Command Line Interface Help

You access the CLI help by entering a full or partial command string and a question mark “?”. The CLI displays the command keywords or parameters along with a short description. For example, at the CLI command prompt, type:

```
> show ?
```

The CLI displays this keyword list with short descriptions for each keyword:

```
show ?
  application-priority          Application Priority
  arp                           Internet Protocol (IP)
  bfd                           Bidirectional Forwarding Detection (BFD)
  bgp                           Border Gateway Protocol (BGP)
  bi-lsp                        Bi-directional lsp status and configuration
  bridge                        Bridge group commands
  ce-vlan                       COS Preservation for Customer Edge VLAN
  class-map                     Class map entry
  cli                           Show CLI tree of current mode
  clns                          Connectionless-Mode Network Service (CLNS)
  control-adjacency             Control Adjacency status and configuration
  control-channel               Control Channel status and configuration
  cspf                          CSPF Information
  customer                      Display Customer spanning-tree
  cvlan                         Display CVLAN information
  debugging                     Debugging functions
```

```
  etherchannel                    LACP etherchannel
  ethernet                        Layer-2
...
```

If you type the ? in the middle of a keyword, the CLI displays help for that keyword only.

```
> show de?
debugging  Debugging functions
```

If you type the ? in the middle of a keyword, but the incomplete keyword matches several other keywords, OcNOS displays help for all matching keywords.

```
> show i? (CLI does not display the question mark).
interface  Interface status and configuration
ip         IP information
isis       ISIS information
```

Command Completion

The CLI can complete the spelling of a command or a parameter. Begin typing the command or parameter and then press the tab key. For example, at the CLI command prompt type `sh`:

```
> sh
```

Press the tab key. The CLI displays:

```
> show
```

If the spelling of a command or parameter is ambiguous, the CLI displays the choices that match the abbreviation. Type `show i` and press the tab key. The CLI displays:

```
> show i
  interface  ip         ipv6       isis
> show i
```

The CLI displays the `interface` and `ip` keywords. Type `n` to select `interface` and press the tab key. The CLI displays:

```
> show in
> show interface
```

Type ? and the CLI displays the list of parameters for the `show interface` command.

```
> show interface
  IFNAME  Interface name
  |       Output modifiers
  >       Output redirection
  <cr>
```

The CLI displays the only parameter associated with this command, the `IFNAME` parameter.

Command Abbreviations

The CLI accepts abbreviations that uniquely identify a keyword in commands. For example:

```
> sh int xe0
```

is an abbreviation for:

```
> show interface xe0
```

Command Line Errors

Any unknown spelling causes the CLI to display the error `Unrecognized command` in response to the ?. The CLI displays the command again as last entered.

```
> show dd?
% Unrecognized command
> show dd
```

When you press the Enter key after typing an invalid command, the CLI displays:

```
(config)#router ospf here
                        ^
% Invalid input detected at '^' marker.
```

where the ^ points to the first character in error in the command.

If a command is incomplete, the CLI displays the following message:

```
> show
% Incomplete command.
```

Some commands are too long for the display line and can wrap mid-parameter or mid-keyword, as shown below. This does *not* cause an error and the command performs as expected:

```
area 10.10.0.18 virtual-link 10.10.0.19 authent
ication-key 57393
```

Command Negation

Many commands have a `no` form that resets a feature to its default value or disables the feature. For example:

- The `ip address` command assigns an IPv4 address to an interface
- The `no ip address` command removes an IPv4 address from an interface

Syntax Conventions

Table 2 on page 14 describes the conventions used to represent command syntax in this reference.

Table 2: Syntax conventions

Convention	Description	Example
`monospaced font`	Command strings entered on a command line	`show ip ospf`
`lowercase`	Keywords that you enter exactly as shown in the command syntax.	`show ip ospf`
`UPPERCASE`	See Variable Placeholders	`IFNAME`
`()`	Optional parameters, from which you must select one. Vertical bars delimit the selections. Do not enter the parentheses or vertical bars as part of the command.	`(A.B.C.D|<0-4294967295>)`

Table 2: Syntax conventions (Continued)

Convention	Description	Example
`()`	Optional parameters, from which you select one or none. Vertical bars delimit the selections. Do not enter the parentheses or vertical bars as part of the command.	`(A.B.C.D|<0-4294967295>|)`
`()`	Optional parameter which you can specify or omit. Do not enter the parentheses or vertical bar as part of the command.	`(IFNAME|)`
`{}`	Optional parameters, from which you must select one or more. Vertical bars delimit the selections. Do not enter the braces or vertical bars as part of the command.	`{intra-area <1-255>|inter-area <1-255>|external <1-255>}`
`[]`	Optional parameters, from which you select zero or more. Vertical bars delimit the selections. Do not enter the brackets or vertical bars as part of the command.	`[<1-65535>|AA:NN|internet|local-AS|no-advertise|no-export]`
`?`	Nonrepeatable parameter. The parameter that follows a question mark can only appear once in a command string. Do not enter the question mark as part of the command.	`?route-map WORD`
`.`	Repeatable parameter. The parameter that follows a period can be repeated more than once. Do not enter the period as part of the command.	`set as-path prepend .<1-65535>`

Variable Placeholders

Table 3 on page 15 shows the tokens used in command syntax use to represent variables for which you supply a value.

Table 3: Variable placeholders

Token	Description
`WORD`	A contiguous text string (excluding spaces)
`LINE`	A text string, including spaces; no other parameters can follow this parameter
`IFNAME`	Interface name whose format varies depending on the platform; examples are: `eth0`, `Ethernet0`, `ethernet0`, `xe0`
`A.B.C.D`	IPv4 address
`A.B.C.D/M`	IPv4 address and mask/prefix
`X:X::X:X`	IPv6 address
`X:X::X:X/M`	IPv6 address and mask/prefix
`HH:MM:SS`	Time format

Table 3: Variable placeholders

Token	Description
`AA:NN`	BGP community value
`XX:XX:XX:XX:XX:XX`	MAC address

`<1-5>`
`<1-65535>`
`<0-2147483647>`
`<0-4294967295>`
	Numeric range

Command Description Format

Table 4 on page 16 explains the sections used to describe each command in this reference.

Table 4: Command descriptions

Section	Description
Command Name	The name of the command, followed by what the command does and when should it be used
Command Syntax	The syntax of the command
Parameters	Parameters and options for the command
Default	The state before the command is executed
Command Mode	The mode in which the command runs; see Command Modes
Applicability	The command introduced in a specific release version and modified or updated in subsequent versions.
Example	An example of the command being executed

Keyboard Operations

Table 5 on page 16 lists the operations you can perform from the keyboard.

Table 5: Keyboard operations

Key combination	Operation
Left arrow or Ctrl+b	Moves one character to the left. When a command extends beyond a single line, you can press left arrow or Ctrl+b repeatedly to scroll toward the beginning of the line, or you can press Ctrl+a to go directly to the beginning of the line.
Right arrow or Ctrl-f	Moves one character to the right. When a command extends beyond a single line, you can press right arrow or Ctrl+f repeatedly to scroll toward the end of the line, or you can press Ctrl+e to go directly to the end of the line.

Table 5: Keyboard operations (Continued)

Key combination	Operation
Esc, b	Moves back one word
Esc, f	Moves forward one word
Ctrl+e	Moves to end of the line
Ctrl+a	Moves to the beginning of the line
Ctrl+u	Deletes the line
Ctrl+w	Deletes from the cursor to the previous whitespace
Alt+d	Deletes the current word
Ctrl+k	Deletes from the cursor to the end of line
Ctrl+y	Pastes text previously deleted with Ctrl+k, Alt+d, Ctrl+w, or Ctrl+u at the cursor
Ctrl+t	Transposes the current character with the previous character
Ctrl+c	Ignores the current line and redisplays the command prompt
Ctrl+z	Ends configuration mode and returns to exec mode
Ctrl+l	Clears the screen
Up Arrow or Ctrl+p	Scroll backward through command history
Down Arrow or Ctrl+n	Scroll forward through command history

Show Command Modifiers

Note: The show command output included in the guides is for illustration purposes only. Based on the combination of features enabled and ongoing enhancements made to the commands, the output for these commands may vary. For instance, the actual command output may differ depending on the software version, configuration, and platform. Field names, values, and formats are subject to change.

You can use two tokens to modify the output of a `show` command. Enter a question mark to display these tokens:

```
# show users ?
  | Output modifiers
  > Output redirection
```

You can type the | (vertical bar character) to use output modifiers. For example:

```
> show rsvp | ?
  begin     Begin with the line that matches
  exclude   Exclude lines that match
  include   Include lines that match
  last      Last few lines
  redirect  Redirect output
```

Begin Modifier

The `begin` modifier displays the output beginning with the first line that contains the input string (everything typed after the `begin` keyword). For example:

```
# show running-config | begin xe1
...skipping
interface xe1
 ipv6 address fe80::204:75ff:fee6:5393/64
!
interface xe2
 ipv6 address fe80::20d:56ff:fe96:725a/64
!
line con 0
 login
!
end
```

You can specify a regular expression after the `begin` keyword, This example begins the output at a line with either “xe2” or “xe4”:

```
# show running-config | begin xe[2-4]

...skipping
interface xe2
 shutdown
!
interface xe4
 shutdown
!
interface svlan0.1
 no shutdown
!
route-map myroute permit 2
!
route-map mymap1 permit 10
!
route-map rmap1 permit 2
!
line con 0
 login
line vty 0 4
 login
!
end
```

Include Modifier

The `include` modifier includes only those lines of output that contain the input string. In the output below, all lines containing the word “input” are included:

```
# show interface xe1 | include input
    input packets 80434552, bytes 2147483647, dropped 0, multicast packets 0
    input errors 0, length 0, overrun 0, CRC 0, frame 0, fifo 1, missed 0
```

You can specify a regular expression after the `include` keyword. This examples includes all lines with “input” or “output”:

```
#show interface xe0 | include (in|out)put
    input packets 597058, bytes 338081476, dropped 0, multicast packets 0
    input errors 0, length 0, overrun 0, CRC 0, frame 0, fifo 0, missed 0
    output packets 613147, bytes 126055987, dropped 0
    output errors 0, aborted 0, carrier 0, fifo 0, heartbeat 0, window 0
```

Exclude Modifier

The `exclude` modifier excludes all lines of output that contain the input string. In the following output example, all lines containing the word “input” are excluded:

```
# show interface xe1 | exclude input
Interface xe1
  Scope: both
  Hardware is Ethernet, address is 0004.75e6.5393
  index 3 metric 1 mtu 1500 <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Administrative Group(s): None
  DSTE Bandwidth Constraint Mode is MAM
  inet6 fe80::204:75ff:fee6:5393/64
    output packets 4438, bytes 394940, dropped 0
    output errors 0, aborted 0, carrier 0, fifo 0, heartbeat 0, window 0
    collisions 0
```

You can specify a regular expression after the `exclude` keyword. This example excludes lines with “output” or “input”:

```
# show interface xe0 | exclude (in|out)put
Interface xe0
  Scope: both
  Hardware is Ethernet  Current HW addr: 001b.2139.6c4a
  Physical:001b.2139.6c4a  Logical:(not set)
  index 2 metric 1 mtu 1500 duplex-full arp ageing timeout 3000
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Bandwidth 100m
  DHCP client is disabled.
  inet 10.1.2.173/24 broadcast 10.1.2.255
  VRRP Master of :  VRRP is not configured on this interface.
  inet6 fe80::21b:21ff:fe39:6c4a/64
    collisions 0
```

Redirect Modifier

The `redirect` modifier writes the output into a file. The output is not displayed.

```
# show cli history | redirect /var/frame.txt
```

The output redirection token (>) does the same thing:

```
# show cli history >/var/frame.txt
```

Last Modifier

The `last` modifier displays the output of last few number of lines (As per the user input). The last number ranges from 1 to 9999.

For example:

```
#show running-config | last 10
```

String Parameters

The restrictions in Table 6 on page 20 apply for all string parameters used in OcNOS commands, unless some other restrictions are noted for a particular command.

Table 6: String parameter restrictions

Restriction	Description
Input length	1965 characters or less
Restricted special characters	“?”, “,”, “>”, “|”, and “=” The “|” character is allowed only for the `description` command in interface mode.

Command Modes

Commands are grouped into modes arranged in a hierarchy. Each mode has its own set of commands. Table P-7 lists the command modes common to all protocols.

Table 7: Common command modes

Name	Description
Executive mode	Also called *view* mode, this is the first mode to appear after you start the CLI. It is a base mode from where you can perform basic commands such as `show`, `exit`, `quit`, `help`, and `enable`.
Privileged executive mode	Also called *enable* mode, in this mode you can run additional basic commands such as `debug`, `write`, and `show`.
Configure mode	Also called *configure terminal* mode, in this mode you can run configuration commands and go into other modes such as interface, router, route map, key chain, and address family. Configure mode is single user. Only one user at a time can be in configure mode.
Interface mode	In this mode you can configure protocol-specific settings for a particular interface. Any setting you configure in this mode overrides a setting configured in router mode.
Router mode	This mode is used to configure router-specific settings for a protocol such as BGP or OSPF.

Command Mode Tree

The diagram below shows the common command mode hierarchy.

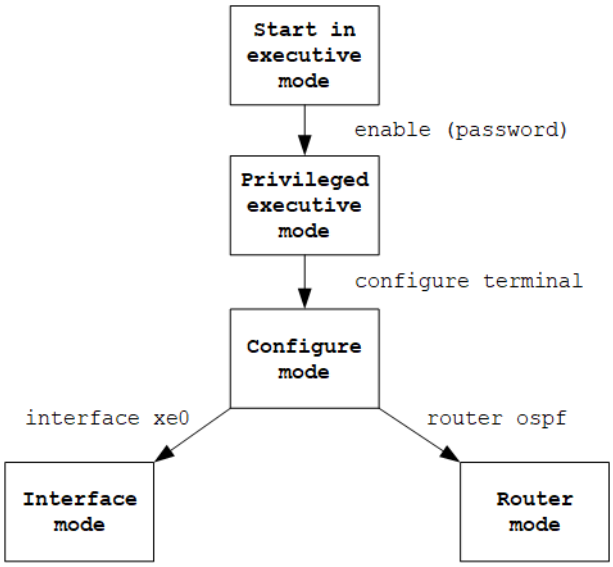

Figure 1: Common command modes

To change modes:

1. Enter privileged executive mode by entering `enable` in Executive mode.
2. Enter configure mode by entering `configure terminal` in Privileged Executive mode.

The example below shows moving from executive mode to privileged executive mode to configure mode and finally to router mode:

```
> enable mypassword
# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
(config)# router ospf
(config-router)#
```

Note: Each protocol can have modes in addition to the common command modes. See the command reference for the respective protocol for details.

Transaction-based Command-line Interface

The OcNOS command line interface is transaction based:

- Any changes done in configure mode are stored in a separate *candidate* configuration that you can view with the `show transaction current` command.
- When a configuration is complete, apply the candidate configuration to the running configuration with the `commit` command.
- If a `commit` fails, no configuration is applied as the entire transaction is considered failed. You can continue to change the candidate configuration and then retry the `commit`.
- Discard the candidate configuration with the `abort transaction` command.
- Check the last aborted transaction with the `show transaction last-aborted` command.
- Multiple configurations cannot be removed with a single `commit`. You must remove each configuration followed by a `commit`.

Note: All commands MUST be executed only in the default CML shell (`cmlsh`). If you log in as root and start `imish`, then the system configurations will go out of sync. The `imish` shell is not supported and should not be started manually.

Virtual eXtensible Local Area Network Configuration

CHAPTER 1 VxLAN Overview

This chapter provides an overview of Virtual Extensible Local Area Network (VxLAN) and its implementation within OcNOS.

VxLAN

Virtual eXtensible Local Area Network (VxLAN) is widely used in Data Centers (DC) networks. VxLAN is an overlay transport virtualization technology commonly used in cloud networks to support the ever-increasing Virtual LAN (VLAN) and multi-tenant networks in data centers. It enables the "stretching" of a Layer 2 network over a physical Layer 3 network.

Overlay Network

VxLAN creates an overlay of virtual L2 LAN segments using a MAC address, and L4 UDP packets in IP encapsulation on top of the physical underlay L3 infrastructure. Technically, it encapsulates L2 Ethernet frames received from a host with L3 IP/UDP packets in a VxLAN header and sends it to the destination in the data center network using IP tunnels. This allows for the extension of L2 networks across data center without changing the underlying physical infrastructure.

EVPN - VxLAN

VxLAN creates LAN segments using MAC-in-IP encapsulation. The encapsulation carries the original L2 frame received from a host to the destination in another host using IP tunnels. The endpoints of the virtualized tunnel formed using VxLAN are called VXLAN Tunnel End Points (VTEPs). The VxLAN segments carry tenant data in L3 tunnels over the network which permits the network to support multiple tenants.The tenant data is not used in routing or switching. This aids in tenant machine movement and allows the tenants to have the same IP or MAC addresses.

Ethernet Virtual Private Network (EVPN) is a protocol based on industry standards used for network virtualization in multi-tenant data center and service provider networks. When used with VxLAN networks, it provides a control plane to create L2 overlays across a L3 network. This enables seamless communication between virtual machines (VMs) or containers across different physical locations as though they are on the same Ethernet segment. EVPN leverages BGP for scalable routing information exchange and distribution of L2 and L3 reachability information across a large network. It also tracks and updates the location of devices based on MAC and IP addresses as they move across the network. It is a critical feature in virtualized environments where VMs or containers may frequently migrate.

Additionally, EVPN supports redundant and active-active multi-homing for robust failover and high availability, by allowing a host to connect to multiple VTEPs. Thus, EVPN ensures, that the traffic is rerouted through an alternate VTEP if a connection to VTEP fails. EVPN is essential for scalable, resilient, and efficient multi-tenant network virtualization when deployed in conjunction with VxLAN in modern data center environments.

Underlay Network

The underlay network consists of a physical L3 infrastructure, which provides the foundation for communication in the network. The underlay network is abstracted in the overlay network, allowing seamless communication of the large virtualized L2 network.

Note: For OcNOS, broadcast, unknown unicast, and multicast traffic is not load-balanced over tunnel L3 next hops or ECMP. Instead, a single next hop is used for forwarding this traffic over the tunnel.

VxLAN Architecture

Typically, VxLAN network operates as an overlay network over an IP underlay network based on a Spine-Leaf CLOS architecture. The underlay network is often referred to as IP fabric or CLOS fabric.

VxLAN allows the network to support several tenants with minimum changes in the network. They carry tenant data in virtual tunnels over the network. The tenant data is not used in routing or switching. This aids in tenant machine movement and allows the tenants to have the same IP or MAC addresses on end devices, hosts, or VMs.

Each overlay tunnel is referred to as a VxLAN segment. VMs can only communicate with each other within the same VxLAN segment, similar to how communication occurs within a traditional VLAN. Each VxLAN segment is identified through a 24-bit segment ID termed the VxLAN Network Identifier (VNI). This allows up to 16 million VxLAN segments to coexist within the same administrative domain. For VMs on different VxLAN segments to communicate, inter-VxLAN routing can be performed on a VxLAN-enabled router or distributed gateway. This is similar to routing between VLANs in a traditional L2 network.

The VNI determines the scope of the inner MAC frame originated from the individual VM. This ensures there can be overlapping MAC addresses across segments, but traffic remains isolated due to the VNI preventing cross-segment interference.

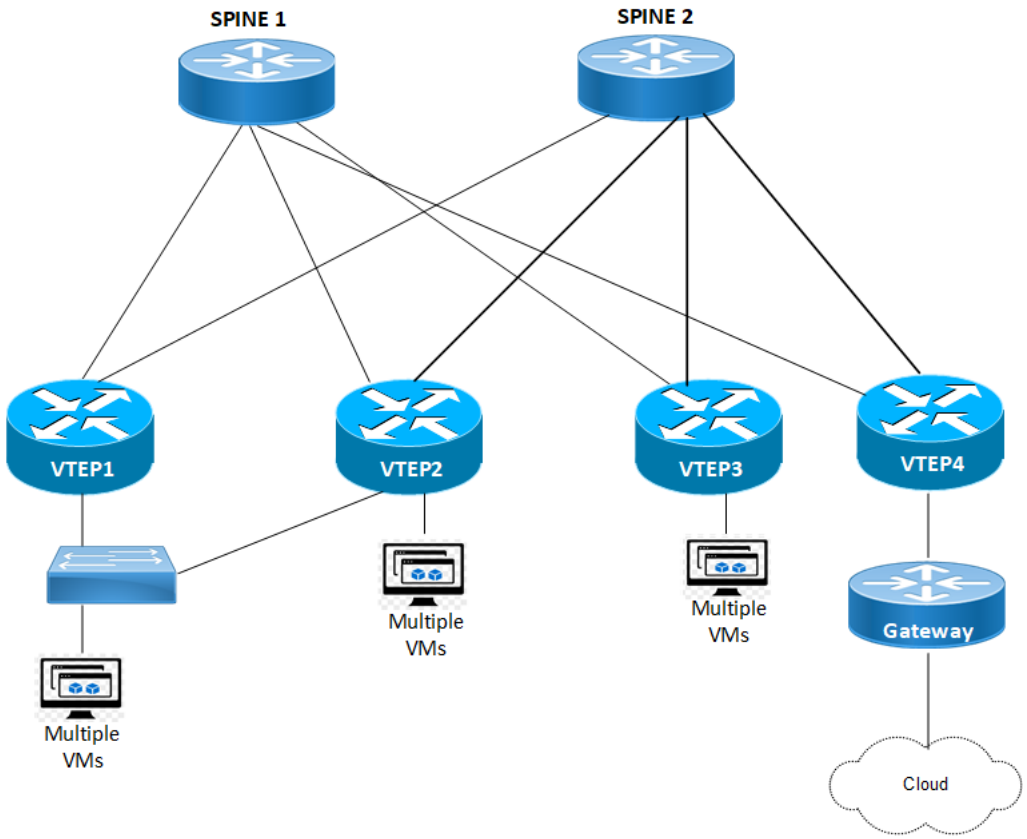

Figure 1-1: VxLAN Deployment - VTEPs across a L3 Network

Features

- Leaf nodes act as VTEP for hosts to connect to the data center and provide VPN services

- With multihoming facility load distribution, link and node level redundancies of the CLOS fabric are extended to hosts
- Hosts are identified using either the port number, port number with VLAN ID or VLAN range, or the port number with a stacked VLAN ID.
- Any packets (including ARP-ND) that are uplifted to the VxLAN CPU queue from any port are rate limited to 500 packets per second. This is done to protect the system and CPU during an ARP storm.
- OcNOS supports VxLAN IPv4 tunnels, but both IPv4 and IPv6 hosts.
- VxLAN works over UDP, with destination port 4789. Source port can be randomized based on L2 frame information carried, hence provides good hashing input for load-sharing on ECMP/LAG paths in the L3 fabric.
- EVPN uses multiprotocol BGP with AFI=25 (L2VPN) and SAFI=70 (EVPN).
- EVPN is used with VxLAN data-plane encapsulations in OcNOS Data Center.
- EVPN helps with discovering VTEPs and learning MAC and IP addresses of the connected hosts in a VxLAN network.
- EVPN is used to implement Integrated routing and bridging (IRB), and E-Tree in OcNOS Data Center. It’s becoming umbrella for multiple service types.
- EVPN provides multihoming service, with redundancy and more bandwidth with multihoming facility load distribution, extends link and node level redundancies of the CLOS fabric to hosts.

Terminology

Terms related to VxLAN configuration are defined in the table below.

VLAN	Virtual Local Area Network
VM	Virtual Machine
VNI	VxLAN Network Identifier (or VxLAN Segment ID)
VTEP	VxLAN Tunnel End Point. An entity that originates and/or terminates VxLAN tunnels
VxLAN	Virtual eXtensible Local Area Network
VxLAN Segment	VxLAN L2 overlay network over which VMs communicate

VxLAN - Data Forwarding Configuration

CHAPTER 1 Static VxLAN Configuration

This chapter explains the static configurations of VxLAN using dynamic OSPF routing protocol configuration.

Configuration

This configuration establishes Layer 2 connectivity over an IP network by setting up two spine switches and VXLAN Tunnel Endpoints (VTEPs) with static VXLAN routing.

In this configuration, each VTEP (VTEP1 and VTEP2) operates as a multilayer switch. Here, `xe34/xe34` serves as the access port for VTEP1/VTEP2 while `ce50/ce49` and `ce12/ce11` acts as the transport network port towards SPINE1 and 2 respectively.

When VM1 sends an untagged IPv4 packet, it arrives at VTEP1 via the access port `xe34`. VTEP1 encapsulates the packet using the configured VXLAN Network Identifier (VNID 3) and forwards it through the tunnel interface over (`ce50/ce49`). The encapsulated packet is then received by VTEP2 on its transport port (`ce12/ce11`). VTEP2 performs decapsulation, extracting the original packet, and forwards it through its access port `xe34` by looking up its local forwarding table using VNID. Finally, the packet reaches the destination VM (VM2), completing the communication.

This process ensures seamless Layer 2 connectivity between the two VMs over the IP network using static VXLAN routing.

Topology

The procedures in this section use the topology shown in Figure 1-2

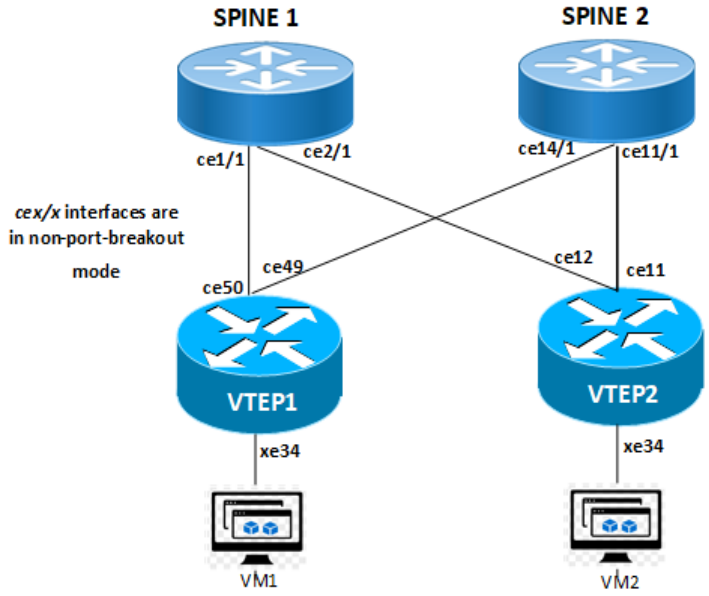

Figure 1-2: Static VxLAN Routing

Procedure to configure Static VxLAN Routing

Configure the VTEP1:

1. Login to VTEP1 and configure VxLAN mode, assign VNID, map network tunnel with VNID, and enter static route for remote VM (VTEP 2) with MAC address and IP address.

```
(config)#hostname VTEP1
(config)#nvo vxlan enable
(config)#nvo vxlan id 3
(config-nvo)#vxlan map-network tunnel Tunnel4
(config-nvo)#vxlan static-entry host-mac 0000.0000.aaaa remote-vtep-ip 11.11.11.1
```

2. Configure VxLAN tunnel with source IP address as VTEP1 IP address and destination IP address as VTEP 2 IP address.

```
(config)#interface Tunnel4
(config-if)#tunnel mode vxlan
(config-if)#tunnel source 10.10.10.1
(config-if)#tunnel destination 11.11.11.1
```

3. Configure VTEP 1 network interface `ce49`.

```
(config)#interface ce49
(config-if)#ip address 4.4.4.4/24
```

4. Configure VTEP 1 network interface `ce50` with IP address.

```
(config)#interface ce50
(config-if)#ip address 2.2.2.4/24
```

5. Configure VTEP1 loop back address.

```
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.1/32 secondary
```

6. Configure VTEP 1 switchport interface `xe34`.

```
(config)#interface xe34
(config-if)#switchport
```

7. Configure a dynamic routing using OSPF router for the network tunnel interfaces `ce50` and `ce49` source IP address.

```
(config)router ospf 1
(config-router)ospf router-id 10.10.10.1
(config-router)network 2.2.2.0/24 area 0.0.0.0
(config-router)network 4.4.4.0/24 area 0.0.0.0
(config-router)network 10.10.10.1/32 area 0.0.0.0
```

8. Map the access port interface `xe34` with network tunnel VNID 3.

```
(config)#nvo vxlan access-if port xe34
(config-nvo-acc-if)#map vnid 3
```

Configure the VTEP2

1. Login to VTEP2 and configure VxLAN mode, assign VNID, map network tunnel with VNID, and enter static route for remote VM (VTEP 2) with MAC address and IP address.

```
(config)#hostname VTEP2
(config)#nvo vxlan enable
(config)#nvo vxlan id 3
(config-nvo)#vxlan map-network tunnel Tunnel4
(config-nvo)#vxlan static-entry host-mac 0000.0000.bbbb remote-vtep-ip 10.10.10.1
```

2. Configure VxLAN tunnel with source IP address as VTEP2 IP address and destination IP as VTEP2 IP address.

```
(config)#interface Tunnel4
(config-if)#tunnel mode vxlan
(config-if)#tunnel source 11.11.11.1
(config-if)#tunnel destination 10.10.10.1
```

3. Configure VTEP2 network interface `ce11`.

```
(config)#interface ce11
(config-if)#ip address 5.5.5.4/24
```

4. Configure VTEP2 network interface `ce12` with IP address.

```
(config)#interface ce12
(config-if)#ip address 3.3.3.4/24
```

5. Configure VTEP2 loop back address.

```
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 11.11.11.1/32 secondary
```

6. Configure VTEP 2 switchport interface *xe34*.

```
(config)#interface xe34
(config-if)#switchport
```

7. Configure a dynamic routing using OSPF router for the network tunnel interfaces `ce50` and `ce49` source IP address.

```
(config)router ospf 1
(config-router)ospf router-id 11.11.11.1
(config-router)network 3.3.3.0/24 area 0.0.0.0
(config-router)network 5.5.5.0/24 area 0.0.0.0
(config-router)network 11.11.11.1/32 area 0.0.0.0
```

8. Map the access port interface `xe34` with network tunnel VNID 3.

```
(config)#nvo vxlan access-if port xe34
(config-nvo-acc-if)#map vnid 3
```

Configure the SPINE1

1. Login to SPINE1 and configure the network interfaces.

```
(config)#hostname SPINE1
(config)#interface ce1/1
(config-if)# ip address 3.3.3.1/24
(config-if)#exit
(config)#interface ce2/1
(config-if)#ip address 2.2.2.1/24
```

2. Configure SPINE1 loop back address.

```
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 12.12.12.1/32 secondary
```

3. Configure a dynamic routing using OSPF router.

```
(config)router ospf 1
(config-router)ospf router-id 12.12.12.1
(config-router)network 2.2.2.0/24 area 0.0.0.0
(config-router)network 3.3.3.0/24 area 0.0.0.0
(config-router)network 12.12.12.1/32 area 0.0.0.0
```

Configure the SPINE2

1. Login to SPINE2 and configure the network interfaces.

```
(config)#hostname SPINE1
(config)#interface ce11/1
(config-if)# ip address 5.5.5.1/24
(config-if)#exit
(config)#interface ce14/1
(config-if)#ip address 4.4.4.1/24
```

2. Configure SPINE2 loop back address.

```
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 13.13.13.1/32 secondary
```

3. Configure a dynamic routing using OSPF router.

```
(config)router ospf 1
(config-router)ospf router-id 13.13.13.1
(config-router)network 4.4.4.0/24 area 0.0.0.0
(config-router)network 5.5.5.0/24 area 0.0.0.0
(config-router)network 13.13.13.1/32 area 0.0.0.0
```

Validation

Execute following show commands to verify the static VxLAN to reach remote virtual machines through the network tunnel interface.

VTEP1

Verify the OSPF neighbors towards SPINEs:

```
VTEP1#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):

Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
12.12.12.1        1   Full/DR         00:00:37    2.2.2.1         ce50                    0
13.13.13.1        1   Full/Backup     00:00:36    4.4.4.1         ce49                    0
VTEP1#
```

Verify the static VxLAN tunnel on VTEP1:

```
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status        Up/Down       Update
=================================================================
10.10.10.1        11.11.11.1       Installed     00:11:58      00:10:13

Total number of entries are 1
VTEP1#
```

Verify the VNID Mapping:

```
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                           VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
3        ----         L2       NW   ----      ----                          ----      ----      10.10.10.1       11.11.11.1
3        ----         --       AC   xe34      --- Single Homed Port ---     0         ----      ----             ----

Total number of entries are 2
VTEP1#
```

Verify the local and remote MAC address configured:

```
VTEP1#show nvo vxlan mac-table
=============================================================================================================================
                                                VXLAN MAC Entries
=============================================================================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI      Type       Status     MAC move AccessPortDesc LeafFlag
_____________________________________________________________________________________________________________________________

3          ----      ----      ----      0000.0000.aaaa 11.11.11.1       Static Remote   -------    0        -------         ----
3          xe34      ----      ----      0000.0000.bbbb 10.10.10.1      Dynamic Local   -------    0        -------         ----

Total number of entries are : 2

VTEP1#
VTEP1#show nvo vxlan mac-table vnid 3
=============================================================================================================================
                                                VXLAN MAC Entries
=============================================================================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI      Type       Status     MAC move AccessPortDesc LeafFlag
_____________________________________________________________________________________________________________________________

3          ----      ----      ----      0000.0000.aaaa 11.11.11.1      Static Remote   -------    0        -------         ----
3          xe34      ----      ----      0000.0000.bbbb 10.10.10.1      Dynamic Local   -------    0        -------         ----

Total number of entries are : 2

VTEP1#
```

VTEP2

Verify the OSPF neighbors towards SPINEs.

```
VTEP2#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
12.12.12.1        1   Full/DR         00:00:35    3.3.3.1         ce12                    0
13.13.13.1        1   Full/Backup     00:00:31    5.5.5.1         ce11                    0
VTEP2#
```

Verify the static tunnel configured on VTEP2:

```
VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status        Up/Down       Update
=====================================================================
11.11.11.1        10.10.10.1       Installed     00:11:32      00:10:28

Total number of entries are 1
VTEP2#
```

Verify the VNID Mapping:

```
VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
3        ----         L2       NW   ----      ----                           ----      ----      11.11.11.1       10.10.10.1
3        ----         --       AC   xe34      --- Single Homed Port ---      0         ----      ----             ----

Total number of entries are 2
VTEP2#
```

Verify the local and remote MAC address configured:

```
VTEP2#
VTEP2#show nvo vxlan mac-table
=============================================================================================================================================
                                                    VXLAN MAC Entries
=============================================================================================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI      Type        Status     MAC move AccessPortDesc LeafFlag
_____________________________________________________________________________________________________________________________________

3          xe34      ----      ----      0000.0000.aaaa 11.11.11.1       Dynamic Local   -------    0        -------          ----
3          ----      ----      ----      0000.0000.bbbb 10.10.10.1       Static Remote   -------    0        -------          ----

Total number of entries are : 2

VTEP2#
VTEP2#show nvo vxlan mac-table vnid 3
=============================================================================================================================================
                                                    VXLAN MAC Entries
=============================================================================================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI      Type        Status     MAC move AccessPortDesc LeafFlag
_____________________________________________________________________________________________________________________________________

3           xe34      ----      ----      0000.0000.aaaa 11.11.11.1       Dynamic Local   -------    0    ------      ----
3           ----      ----      ----      0000.0000.bbbb 10.10.10.1       Static Remote   -------    0    -------     ----

Total number of entries are : 2

VTEP2#
```

CHAPTER 2 VLAN to VNID Mapping

Overview

OcNOS supports mapping Virtual Local Area Network Identifier (VLAN ID) to Virtual Extensible Local Area Network Identifier (VNID) to extend the Layer 2 VLAN over to the Layer 3 VxLAN. The VLAN ID is a unique number assigned to a specific VLAN, and the VNID is the same for a specific VxLAN. The Virtual Tunnel Endpoint (VTEP) is a network device or a software component that encapsulates the ethernet frames from a specific VLAN ID into a VxLAN header that contains the VNID. This encapsulated frame is the VxLAN packet, which is again encapsulated in the UDP packet and wrapped in an outer IP header to transport it to the IP network.

Feature Characteristics

- The VLAN ID to VNID communication is regardless the physical and logical port.
- Each VNID is allowed to map with a single VLAN ID.
- VLAN is removed over the tunnel and re-added when egressing out of the host-connected port.

Note: VNI to VLAN mapping for VxLAN is only available for TPID 0x8100.

Benefits

The feature enhances the scalability and flexibility by extending the VLAN to VxLAN boundaries.

Configuration

This section shows the procedure for configuring VxLAN using VLAN to VNI mapping.

Topology

This topology demonstrates the configuration necessary to enable VxLAN VLAN-to-VNID mapping. Leaf1 and Leaf2 are VxLAN VTEPs in the Autonomous System (ASN) 100, and the Spine1 resides in a different Autonomous System (ASN) 200. The setup involves establishing eBGP IPv4 (underlay) neighborship between interface addresses and L2VPN EVPN (overlay) neighborship between loopback interfaces.

Once the underlay and overlay neighborships are established, configure the VxLAN tunnel between Leaf1 and Leaf2. The global VTEP IP is the loopback address.

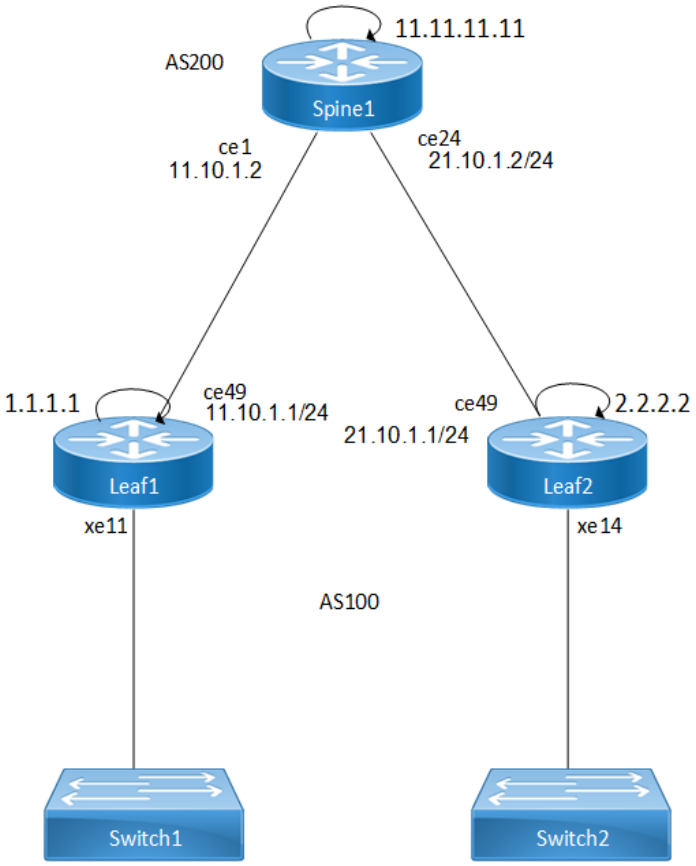

Figure 2-3: VLAN to VNID mapping

Configuring VLAN to VNID mapping

Follow the steps to configure the VxLAN EVPN using VLAN to VNI mapping. The detailed configuration procedure demonstrates how `access-if` VxLAN is configured on a switchport in trunk mode.

Note: The configuration of `access-if` VxLAN is allowed in the access mode also. The example is as follows:

```
(config)#interface xe12
(config-if)# switchport
(config-if)# bridge-group 1 spanning-tree disable
(config-if)# switchport mode access
(config-if)# switchport access vlan 20
(config-if)# access-if-vxlan
(config-if)#exit
```

Configure the Leaf1 and Leaf2:

The parameters used in the configuration procedure are as present for Leaf1. Use the same commands with the Leaf2 parameters to configure the Leaf2.

1. Configure the hostname of the VTEP.

   ```
   (config)#hostname Leaf1
   (config)#commit
   ```

2. Configure the bridge type using the `bridge 1 protocol rstp vlan-bridge` command. This command enables the RSTP VLAN bridge type.

   ```
   (config)#bridge 1 protocol rstp vlan-bridge
   ```

3. Enter the VLAN database using the `vlan database` command and associate the VLAN with a bridge.

```
(config)#vlan database
(config-vlan)#vlan 10 bridge 1 state enable
```

4. Enable the VxLAN globally on the VTEP.

```
(config)#nvo VxLAN enable
```

Note: Save and reboot the system to enable the VxLAN in the hardware.

5. Create the MAC VRF and name it using `mac vrf vrf10`. Configure the VLAN-based service type using `evpn-vlan-service vlan-based` command, assign a unique route distinguisher and route target value using `rd 1.1.1.1:10` and `route-target both 100:10` commands respectively.

```
(config)#mac vrf vrf10
(config-vrf)#evpn-vlan-service vlan-based
(config-vrf)#rd 1.1.1.1:10
(config-vrf)#route-target both 100:10
```

6. Configure a global IP to the VTEP. This IP address uniquely identifies the VTEP.

```
(config)#nvo VxLAN vtep-ip-global 1.1.1.1
```

7. Create a VNID and map it with the bridge VLAN using `nvo VxLAN id 10 ingress-replication bridge-vlan 10`. Configure the host-reachability-protocol as BGP-EVPN and associate the MAC VRF.

```
(config)#nvo VxLAN id 10 ingress-replication bridge-vlan 10
(config-nvo)#VxLAN host-reachability-protocol evpn-bgp vrf10
```

8. Assign the IP addresses to the physical and loopback interfaces of the Leaf1 to connect to the Spine.

```
(config)#interface ce49
(config-if)#ip address 11.10.1.1/24
(config-if)#exit
(config)#interface lo
(config-if)#ip address 1.1.1.1/32 secondary
```

9. Configure the interface `xe11` as a switchport. Use the command `bridge-group 1 spanning-tree disable` to associate the bridge group to this interface and disable the Spanning Tree Protocol (STP) to avoid the port block. Configure the switching characteristic of this interface to trunk mode using the `switchport mode trunk` command and this allows multiple VLANs to run in the interface. Use the command `switchport trunk allowed vlan add 10` to enable VLAN 10 through this interface. Map this interface with the VxLAN using `access-if-`VxLAN command.

```
(config)#interface xe11
(config-if)#switchport
(config-if)# bridge-group 1 spanning-tree disable
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add 10
(config-if)#access-if-vxlan
```

10. Configure the BGP and specify the autonomous number (ASN).

```
(config)#router bgp 100
```

11. Configure the router ID.

```
(config-router)#bgp router-id 1.1.1.1
```

12. Configure the neighboring eBGP peers in a different ASN.

```
(config-router)#neighbor 11.10.1.2 remote-as 200
(config-router)#neighbor 11.11.11.11 remote-as 200
```

13. Configure eBGP multihop as the neighboring peer is not directly connected.

```
(config-router)#neighbor 11.11.11.11 ebgp-multihop
```

14. Configure the source loopback address.

```
(config-router)#neighbor 11.11.11.11 update-source lo
```

15. Configure the IPv4 address family and activate the neighbor.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 1.1.1.1/32
(config-router)#neighbor 11.10.1.2 activate
(config-router-af)#neighbor 11.10.1.2 allowas-in 1
(config-router-af)#exit-address-family
```

16. Configure the Layer 2 VPN address family and activate the neighbor.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor 11.11.11.11 activate
(config-router-af)#neighbor 11.11.11.11 allowas-in 1
(config-router-af)#exit-address-family
(config-router)#exit
(config)#commit
```

Configure the Spine1:

1. Configure the hostname of the Spine.

```
(config)#hostname Spine1
(config)#commit
```

2. Assign the IP addresses to the physical and loopback interfaces of the Spine.

```
(config)#interface ce1
(config-if)#ip address 11.10.1.2/24
(config-if)#exit
(config)#interface ce24
(config-if)#ip address 21.10.1.2/24
(config-if)#exit
(config)#interface lo
(config-if)#ip address 11.11.11.11/32 secondary
(config-if)#exit
```

3. Configure the BGP and specify the ASN.

```
(config)#router bgp 200
```

4. Configure the router ID.

```
(config-router)#bgp router-id 11.11.11.11
```

5. Disable the inbound route filter.

```
(config-router)#no bgp inbound-route-filter
```

6. Configure the neighboring eBGP neighbor in a different ASN.

```
(config-router)#neighbor 11.10.1.1 remote-as 100
(config-router)#neighbor 21.10.1.1 remote-as 100
(config-router)#neighbor 1.1.1.1 remote-as 100
(config-router)#neighbor 2.2.2.2 remote-as 100
```

7. Configure eBGP multihop as the neighboring peer might not be directly connected.

```
(config-router)#neighbor 1.1.1.1 ebgp-multihop
(config-router)#neighbor 2.2.2.2 ebgp-multihop
```

8. Configure the source loopback address.

```
(config-router)#neighbor 1.1.1.1 update-source
(config-router)#neighbor 2.2.2.2 update-source
```

9. Configure the IPv4 address family and activate the neighbor.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 11.11.11.11/32
(config-router-af)#neighbor 11.10.1.1 activate
(config-router-af)#neighbor 21.10.1.1 activate
(config-router-af)#exit-address-family
```

10. Configure the Layer 2 VPN address family and activate the neighbor.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor 1.1.1.1 activate
(config-router-af)#neighbor 2.2.2.2 activate
(config-router-af)#exit-address-family
(config-router)#exit
(config)#commit
```

Running configurations

The running configuration for the Leaf1 is as follows:

```
hostname Leaf1
bridge 1 protocol rstp vlan-bridge
tfo Disable
!
vlan database
 vlan 10 bridge 1 state enable
!
nvo VxLAN enable
!
mac vrf vrf10
 evpn-vlan-service vlan-based
 rd 1.1.1.1:10
 route-target both 100:10
!
nvo VxLAN vtep-ip-global 1.1.1.1
!
nvo VxLAN id 10 ingress-replication bridge-vlan 10
 VxLAN host-reachability-protocol evpn-bgp vrf10
!
interface ce49
 ip address 11.10.1.1/24
!
interface lo
 ip address 1.1.1.1/32 secondary
 !
interface xe11
 switchport
 bridge-group 1 spanning-tree disable
 switchport mode trunk
 switchport trunk allowed vlan add 10
 access-if-vxlan
 !
 exit
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 11.10.1.2 remote-as 200
 neighbor 11.11.11.11 remote-as 200
 neighbor 11.11.11.11 ebgp-multihop
```

```
     neighbor 11.11.11.11 update-source lo
     !
     address-family ipv4 unicast
     network 1.1.1.1/32
     neighbor 11.10.1.2 activate
     neighbor 11.10.1.2 allowas-in 1
     exit-address-family
     !
     address-family l2vpn evpn
     neighbor 11.11.11.11 activate
     neighbor 11.11.11.11 allowas-in 1
     exit-address-family
     !
     exit
    !
    end
```

The running configuration for the Spine1 is as follows:

```
hostname Spine1
!
interface ce1
 ip address 11.10.1.2/24
!
interface ce24
 ip address 21.10.1.2/24
 lldp-agent
!
interface lo
  ip address 11.11.11.11/32 secondary
!
router bgp 200
 bgp router-id 11.11.11.11
 no bgp inbound-route-filter
 neighbor 1.1.1.1 remote-as 100
 neighbor 2.2.2.2 remote-as 100
 neighbor 11.10.1.1 remote-as 100
 neighbor 21.10.1.1 remote-as 100
 neighbor 1.1.1.1 ebgp-multihop
 neighbor 1.1.1.1 update-source lo
 neighbor 2.2.2.2 ebgp-multihop
 neighbor 2.2.2.2 update-source lo
 !
 address-family ipv4 unicast
 network 11.11.11.11/32
 neighbor 11.10.1.1 activate
 neighbor 21.10.1.1 activate
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 2.2.2.2 activate
 exit-address-family
```

```
 !
 exit
end
```

Validation

Validate the show output after configuration as shown below.

Leaf1:

```
Leaf1#show ip bgp summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 7
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down
State/PfxRcd   Desc
11.10.1.2              4   200       266       264      7     0     0  01:37:50
2

Total number of neighbors 1

Total number of Established sessions 1
Leaf1#show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 8
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down
State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
11.11.11.11            4   200       237       235      8     0     0  01:36:10
3      0      1      2      0         0

Total number of neighbors 1

Total number of Established sessions 1
Leaf1#show nvo VxLAN tunnel
VxLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update        Redund
Description
================================================================================================
==================
1.1.1.1          2.2.2.2          Installed     01:35:22      01:35:22      ----      ----

Total number of entries are 1
Leaf1#show nvo VxLAN vlan-vnid bridge-vlan 10
VLAN     VNID      Interface

_____________________________________________________
10       10        xe11
```

```
Total number of entries are 1

Leaf1#show nvo VxLAN vlan-vnid vnid 10
VLAN     VNID      Interface
_____________________________________________________________
10       10        xe11

Total number of entries are 1

Leaf1#show nvo VxLAN vlan-vnid
VLAN     VNID      Interface
_____________________________________________________________
10       10        xe11

Total number of entries are 2
Leaf1#

Leaf1#show nvo vxlan
VxLAN Information
=================
   Codes: NW - Network Port
           AC - Access Port
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN      DF-
Status Src-Addr         Dst-Addr         Redund    Description
_____________________________________________________________________________________
_____________________________________________________
10       ----         L2       NW   ----      ----                          ----      -
---      1.1.1.1          2.2.2.2          ----      ----
10       ----         --       AC   xe11      --- Single Homed Port ---     10        -
---      ----             ----

Total number of entries are 2
```

Spine1:

```
Spine1#show ip bgp summary
BGP router identifier 11.11.11.11, local AS number 200
BGP table version is 3
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down
State/PfxRcd   Desc
11.10.1.1              4   100       247       246      3     0     0  01:44:11
1
21.10.1.1              4   100       249       247      3     0     0  01:44:11
1

Total number of neighbors 2
```

```
Total number of Established sessions 2
Spine1#
Spine1#show bgp l2vpn evpn summary
BGP router identifier 11.11.11.11, local AS number 200
BGP table version is 9
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                  V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down
State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4   100       249       252      9     0     0  01:42:31
3      0      1      2      0         0
2.2.2.2                 4   100       251       250      9     0     0  01:42:01
3      0      1      2      0         0

Total number of neighbors 2

Total number of Established sessions 2
```

VLAN VNID Mapping Command

The VLAN VNID mapping feature introduces the following configuration command.

access-if-vxlan

Use this command to create a logical port per parent interface for egress VLAN translation.

Use `no` form of this command to unconfigure the `access-if-vxlan`.

Command Syntax

```
access-if-vxlan
no access-if-vxlan
```

Parameters

None

Default

None

Command Mode

Interface mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The below example shows how to configure a logical port per parent interface:

```
OcNOS#configure terminal
```

```
OcNOS(config)#interface xe1
OcNOS(config-if)#access-if-vxlan
```

show nvo VxLAN vlan-vnid

Use this command to display the VLAIN ID to VNID mapping.

Command Syntax

```
show nvo VxLAN vlan-vnid (bridge-vlan <VLAN ID> (summary |) | vnid <VNID> | summary)
```

Parameters

`vnid <VNID>`	(Optional) Displays all the VLAN ID to VNID mapping.
`bridge-vlan <VLAN ID>`	(Optional) Displays all the VLAN ID to VNID mapping.
`summary`	(Optional) Displays the total count of VLAN to VNID mapping.

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The below examples show the output of VLAN to VNID mapping:

```
OcNOS#show nvo VxLAN vlan-vnid
VLAN     VNID      Interface
_____________________________________________________________
10       10        xe11

Total number of entries are 1
OcNOS#

OcNOS#show nvo VxLAN vlan-vnid bridge-vlan 10
VLAN     VNID      Interface
_____________________________________________________________
10       10        xe11

Total number of entries are 1
OcNOS#
```

Table P-2-1 explains the output fields.

Table 2-1: VLAN VNID fields

Field	Description
VLAN	VLAN Identifier.
VNID	VxLAN Identifier.
Interface	Name of the interface.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Virtual Local Area Network Identifier (VLAN ID)	Virtual Local Area Network Identifier is a 12-bit unique identifier assigned to a VLAN to identify it in a network.
Virtual Extensible Local Area Network Identifier (VNID)	Virtual Extensible Local Area Network Identifier is a unique 24-bit identifier assigned to a VxLAN to identify it in a network.
Virtual Local Area Network (VLAN)	Virtual Local Area Network in a network configuration creates a separate and isolated virtual network with other virtual networks over a single physical interface.
Virtual Extensible Local Area Network (VxLAN)	Virtual Extensible Local Area Network (VxLAN) enables the creation of a virtualized Layer 2 network over the Layer 3 infrastructure. This is an overlay network on Layer 3 designed to overcome the limitations of VLANs.
Virtual Tunnel Endpoint (VTEP)	Virtual Tunnel Endpoint is a significant component in VxLAN that encapsulates or decapsulates the VxLAN traffic as it enters or leaves the VxLAN overlay network respectively.

VxLAN - Ethernet Virtual Private Network

CHAPTER 1 EVPN-VxLAN Configuration

This section contains basic EVPN-VxLAN configuration examples.

VxLAN (Virtual eXtended LAN) creates LAN segments using a MAC-in-IP encapsulation. The encapsulation carries the original L2 frame received from a host to the destination in another host using IP tunnels. The endpoints of the virtualized tunnel formed using VxLAN are called VTEPs (VxLAN Tunnel End Points). The VTEPs carry tenant data in L3 tunnels over the network which permits the network to support multiple tenants.The tenant data is not used in routing or switching. This aids in tenant machine movement and allows the tenants to have same IP/MAC addresses.

Information about the given VM to get to the VTEP is crucial in VxLAN protocol; therefore MP-BGP based EVPN is used to carry this information across VTEPs.

Topology

The procedures in this section use the topology in Figure 1-4.

The topology uses the following nodes:

- Two single-homed hosts, SH1 and SH2, are connected to LEAF1 (hosting VTEP1) and LEAF4 (hosting VTEP4) respectively.
- A multihomed host, MH1, is connected to both LEAF1 (VTEP1) and LEAF2 (VTEP2).
- Another multihomed host, MH2, is connected to both LEAF3 (VTEP3) and LEAF4 (VTEP4).
- All leaf switches (LEAF1–LEAF4) are connected to SPINE1 and SPINE2, forming an eBGP-based IP underlay fabric with ECMP (Equal-Cost Multi-Path) routing enabled.

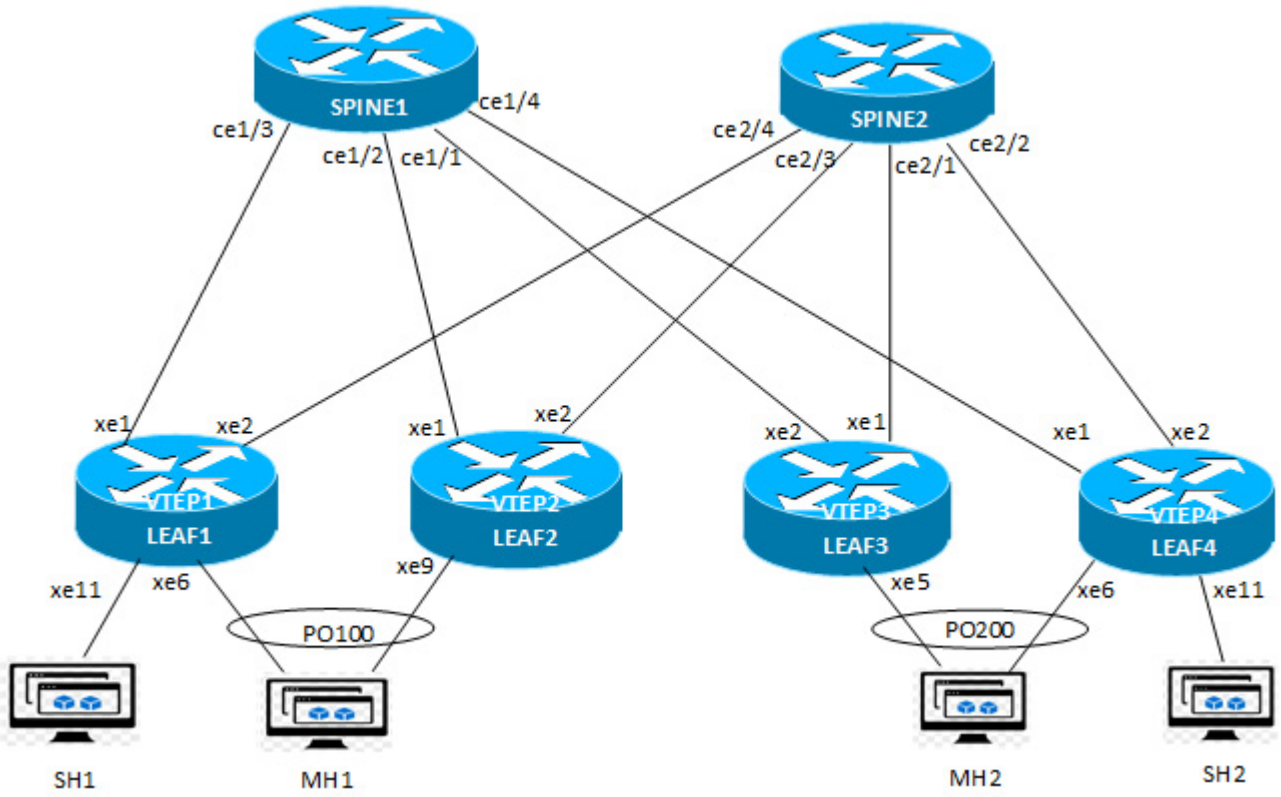

Figure 1-4: EVPN VxLAN with LAG and ECMP

Configuration

This section list the configuration required to set up EVPN VxLAN with SH and MH access port with Dynamic LAG port for multihoming and ECMP routing on the network side. It includes the configuration on following nodes:

1. LEAFs Configuration
2. SPINEs Configuration.

LEAFs Configuration

Following are the sample configuration on LEAF1. Perform the similar configurations on other LEAFs.

1. Login to LEAF1 and configure a host name and enable load balancing configurations and activate VXLAN, VxLAN multihoming at the global level.

```
load-balance rtag7
load-balance rtag7 macro-flow
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port next-hdr
load-balance rtag7 vxlan inner-l3 dest-ip src-ip destl4-port srcl4-port protocol-
id

hostname VTEP1
nvo vxlan enable
!
evpn esi holdtime 90
!
evpn vxlan multihoming enable
```

2. Configure the loopback IP address and assign a global IP to the VTEP1.

```
interface lo
ip address 10.10.10.1/32 secondary
 !
nvo vxlan vtep-ip-global 10.10.10.1
```

3. Create a user-defined MAC VRF (Virtual Routing and Forwarding).

```
mac vrf vpn_blue
     rd 10.10.10.1:200
     route-target both 200:200
mac vrf vpn_green
     rd 10.10.10.1:100
     route-target both 100:100
```

4. Configure a VNID on VTEP and enter into NVO mode. Configure host-reachability-protocol as EVPN-BGP and associate the VNID with VRFs `vpn_green` and `vpn_blue`.

```
nvo vxlan id 100 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vpn_green
!
nvo vxlan id 100 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vpn_blue
!
```

5. Set the dynamic MH LAG access port and assign a `system-mac` address and associate it to a physical port `xe6` and also configure switchport configs on SH access interface.

```
!
interface po100
description multihome_access_interface switchport
load-interval 30
evpn multi-homed system-mac 0000.aaaa.bbbb
!
interface xe6
 description mh_access_member_interface
channel-group 100 mode active
!
interface xe11
 description sh_access_interface
switchport
load-interval 30
!
```

5. Configure the network interfaces as un-numbered and add load interval configurations.

```
interface xe1
    description network_to_spine1
     load-interval 30
interface xe2
     description network_to_spine2
     load-interval 30
```

6. Configure a IP routing between SPINES and VTEPs using eBGP unnumbered router for network interfaces `xe1` and `xe2` as below.

```
router bgp 65001
 bgp router-id 10.10.10.1
 !
 bgp unnumbered-mode
 neighbor xe2 remote-as external
 neighbor xe2 as-origination-interval 1
 neighbor xe2 advertisement-interval 0
 neighbor xe2 fall-over bfd
 neighbor xe1 remote-as external
 neighbor xe1 as-origination-interval 1
 neighbor xe1 advertisement-interval 0
 neighbor xe1 fall-over bfd
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 4
 redistribute connected
 redistribute static
 !
 bgp v4-unnumbered-mode
 neighbor xe2 activate
```

```
 neighbor xe2 allowas-in 1
 neighbor xe1 activate
 neighbor xe1 allowas-in 1
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 !
 bgp l2vpn-unnumbered-mode
 neighbor xe2 activate
 neighbor xe2 allowas-in 1
 neighbor xe1 activate
 neighbor xe1 allowas-in 1
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
 exit
!
```

7. Configure VxLAN access-port for SH `xe11, MH po100` and map to VNI.

```
nvo vxlan access-if port-vlan xe11 51
 description access_if_sh_vlan51_vpn_green
 map vnid 100
!
nvo vxlan access-if port-vlan xe11 151
 description access_if_sh_vlan151_vpn_blue
 map vnid 200
!
nvo vxlan access-if port-vlan po100 101
 description access_if_mh_vlan101_vpn_green
 map vnid 100
!
nvo vxlan access-if port-vlan po100 201
 description access_if_mh_vlan201_vpn_blue
 map vnid 200
```

Sample Configuration for LEAF2:

```
hostname VTEP2
load-balance rtag7
load-balance rtag7 macro-flow
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port next-hdr
load-balance rtag7 vxlan inner-l3 dest-ip src-ip destl4-port srcl4-port protocol-id
 !
 nvo vxlan enable
 !
 evpn esi hold-time 90
```

```
!
evpn vxlan multihoming enable
!
mac vrf vpn_blue
 rd 10.10.10.2:200
 route-target both 200:200
!
mac vrf vpn_green
 rd 10.10.10.2:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 10.10.10.2
!
nvo vxlan id 100 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vpn_green
!
nvo vxlan id 200 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vpn_blue
!
interface po100
 description multihome_access_interface
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.aaaa.bbbb
!
interface lo
 ip address 10.10.10.2/32 secondary
!
interface xe1
 description network_to_spine1
 load-interval 30
!
interface xe2
 description network_to_spine2
 load-interval 30
!
interface xe9
description mh_access_member_interface
 channel-group 100 mode active
!
 exit
!
router bgp 65001
 bgp router-id 10.10.10.2
 !
 bgp unnumbered-mode
 neighbor xe2 remote-as external
 neighbor xe2 as-origination-interval 1
 neighbor xe2 advertisement-interval 0
 neighbor xe2 fall-over bfd
```

```
 neighbor xe1 remote-as external
 neighbor xe1 as-origination-interval 1
 neighbor xe1 advertisement-interval 0
 neighbor xe1 fall-over bfd
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 4
 redistribute connected
 redistribute static
 !
 bgp v4-unnumbered-mode
 neighbor xe2 activate
 neighbor xe2 allowas-in 1
 neighbor xe1 activate
 neighbor xe1 allowas-in 1
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 !
 bgp l2vpn-unnumbered-mode
 neighbor xe2 activate
 neighbor xe2 allowas-in 1
 neighbor xe1 activate
 neighbor xe1 allowas-in 1
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
 exit

!
nvo vxlan access-if port-vlan po100 101
 description access_if_mh_vlan101_vpn_green
 map vnid 100
!
nvo vxlan access-if port-vlan po100 201
 description access_if_mh_vlan201_vpn_blue
 map vnid 200
!
!
end
```

Sample Configuration for LEAF3:

```
hostname VTEP3
load-balance rtag7
```

```
load-balance rtag7 macro-flow
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port next-hdr
load-balance rtag7 vxlan inner-l3 dest-ip src-ip destl4-port srcl4-port protocol-id
 !
 nvo vxlan enable
 !
 evpn esi hold-time 90
 !
 evpn vxlan multihoming enable
 !
 mac vrf vpn_blue
  rd 10.10.10.5:200
  route-target both 200:200
 !
 mac vrf vpn_green
  rd 10.10.10.5:100
  route-target both 100:100
 !
 nvo vxlan vtep-ip-global 10.10.10.5
 !
 nvo vxlan id 100 ingress-replication inner-vid-disabled
  vxlan host-reachability-protocol evpn-bgp vpn_green
 !
 nvo vxlan id 200 ingress-replication inner-vid-disabled
  vxlan host-reachability-protocol evpn-bgp vpn_blue
 !
 interface po200
  description multihome_access_interface
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.cccc.dddd
 !
 interface lo
  ip address 10.10.10.5/32 secondary
 !
 interface xe1
  description network_to_spine1
  load-interval 30
 !
 interface xe2
  description network_to_spine2
  load-interval 30
 !
 interface xe9
  description mh_access_member_interface
  channel-group 200 mode active
 !
  exit
 !
```

```
router bgp 65001
 bgp router-id 10.10.10.5
 !
 bgp unnumbered-mode
 neighbor xe2 remote-as external
 neighbor xe2 as-origination-interval 1
 neighbor xe2 advertisement-interval 0
 neighbor xe2 fall-over bfd
 neighbor xe1 remote-as external
 neighbor xe1 as-origination-interval 1
 neighbor xe1 advertisement-interval 0
 neighbor xe1 fall-over bfd
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 4
 redistribute connected
 redistribute static
 !
 bgp v4-unnumbered-mode
 neighbor xe2 activate
 neighbor xe2 allowas-in 1
 neighbor xe1 activate
 neighbor xe1 allowas-in 1
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 !
 bgp l2vpn-unnumbered-mode
 neighbor xe2 activate
 neighbor xe2 allowas-in 1
 neighbor xe1 activate
 neighbor xe1 allowas-in 1
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan po200 101
 description access_if_mh_vlan101_vpn_green
 map vnid 100
!
nvo vxlan access-if port-vlan po200 201
 description access_if_mh_vlan201_vpn_blue
 map vnid 200
!
!
```

```
 end
```

Sample Configuration for LEAF4:

```
hostname VTEP4
load-balance rtag7
load-balance rtag7 macro-flow
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port next-hdr
load-balance rtag7 vxlan inner-l3 dest-ip src-ip destl4-port srcl4-port protocol-id
 !
 nvo vxlan enable
 !
 evpn esi hold-time 90
 !
 evpn vxlan multihoming enable
 !
 mac vrf vpn_blue
  rd 10.10.10.6:200
  route-target both 200:200
 !
 mac vrf vpn_green
  rd 10.10.10.6:100
  route-target both 100:100
 !
 nvo vxlan vtep-ip-global 10.10.10.6
 !
 nvo vxlan id 100 ingress-replication inner-vid-disabled
  vxlan host-reachability-protocol evpn-bgp vpn_green
 !
 nvo vxlan id 200 ingress-replication inner-vid-disabled
  vxlan host-reachability-protocol evpn-bgp vpn_blue
 !
 interface po200
  description multihome_access_interface
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.cccc.dddd
 !
 interface lo
  ip address 10.10.10.6/32 secondary
 !
 interface xe1
  description network_to_spine1
  load-interval 30
 !
 interface xe2
  description network_to_spine2
  load-interval 30
 !
```

```
interface xe6
description mh_access_member_interface
 channel-group 200 mode active
!
interface xe11
 description sh_access_interface
 switchport
 load-interval 30
!
 exit
!
router bgp 65001
 bgp router-id 10.10.10.6
 !
 bgp unnumbered-mode
 neighbor xe2 remote-as external
 neighbor xe2 as-origination-interval 1
 neighbor xe2 advertisement-interval 0
 neighbor xe2 fall-over bfd
 neighbor xe1 remote-as external
 neighbor xe1 as-origination-interval 1
 neighbor xe1 advertisement-interval 0
 neighbor xe1 fall-over bfd
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 4
 redistribute connected
 redistribute static
 !
 bgp v4-unnumbered-mode
 neighbor xe2 activate
 neighbor xe2 allowas-in 1
 neighbor xe1 activate
 neighbor xe1 allowas-in 1
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 !
 bgp l2vpn-unnumbered-mode
 neighbor xe2 activate
 neighbor xe2 allowas-in 1
 neighbor xe1 activate
 neighbor xe1 allowas-in 1
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
```

```
 exit
!
nvo vxlan access-if port-vlan xe11 51
 description access_if_sh_vlan51_vpn_green
 map vnid 100
!
nvo vxlan access-if port-vlan xe11 151
 description access_if_sh_vlan151_vpn_blue
 map vnid 200
!
nvo vxlan access-if port-vlan po200 101
 description access_if_mh_vlan101_vpn_green
 map vnid 100
!
nvo vxlan access-if port-vlan po200 201
 description access_if_mh_vlan201_vpn_blue
 map vnid 200
!
!
```

SPINEs Configuration

Following are the sample configuration of SPINE1. Perform the similar configurations on other SPINE2.

1. Login to SPINE1. Set host name.

   ```
   hostname SPINE1
   ```

2. Configure the loop back ip address on SPINE

   ```
   interface lo
   ip address 10.10.10.3/32 secondary
   ```

3. Configure network interfaces towards VTEPs as un-numbered and add load interval configs. .

   ```
   interface ce1/1
    port breakout enable 4X10g
    description network_to_vtep3
    load-interval 30
   !
   interface ce1/2
    description network_to_vtep2
    load-interval 30
   !
   interface ce1/3
    description network_to_vtep1
    load-interval 30
   !
   interface ce1/4
    description network_to_vtep4
    load-interval 30
   !
   ```

4. Configure the eBGP with `address-family IPv4 unicast`, `bgp v4-unnumbered-mode` and `bgp l2vpn-unnumbered-mode` configurations and advertise the network using following configurations.

```
router bgp 65009
 bgp router-id 10.10.10.3
 no bgp inbound-route-filter
 !
 bgp unnumbered-mode
 neighbor ce1/2 remote-as external
 neighbor ce1/2 as-origination-interval 1
 neighbor ce1/2 advertisement-interval 0
 neighbor ce1/2 fall-over bfd
 neighbor ce1/4 remote-as external
 neighbor ce1/4 as-origination-interval 1
 neighbor ce1/4 advertisement-interval 0
 neighbor ce1/4 fall-over bfd
 neighbor ce1/1 remote-as external
 neighbor ce1/1 as-origination-interval 1
 neighbor ce1/1 advertisement-interval 0
 neighbor ce1/1 fall-over bfd
 neighbor ce1/3 remote-as external
 neighbor ce1/3 as-origination-interval 1
 neighbor ce1/3 advertisement-interval 0
 neighbor ce1/3 fall-over bfd
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 4
 redistribute connected
 redistribute static
 !
 bgp v4-unnumbered-mode
 neighbor ce1/2 activate
 neighbor ce1/4 activate
 neighbor ce1/1 activate
 neighbor ce1/3 activate
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 !
 bgp l2vpn-unnumbered-mode
 neighbor ce1/2 activate
 neighbor ce1/4 activate
 neighbor ce1/1 activate
 neighbor ce1/3 activate
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
```

```
 exit
```

Sample configuration on Spine2:

```
hostname SPINE2
!
interface ce2/1
 port breakout enable 4X10g
 description network_to_vtep3
 load-interval 30
!
interface ce2/2
 description network_to_vtep4
 load-interval 30
!
interface ce2/3
 description network_to_vtep2
 load-interval 30
!
interface ce2/4
 description network_to_vtep1
 load-interval 30
!
interface lo
 ip address 10.10.10.4/32 secondary
!
 exit
!
router bgp 65009
 bgp router-id 10.10.10.4
 no bgp inbound-route-filter
 !
 bgp unnumbered-mode
 neighbor ce2/3 remote-as external
 neighbor ce2/3 as-origination-interval 1
 neighbor ce2/3 advertisement-interval 0
 neighbor ce2/3 fall-over bfd
 neighbor ce2/2 remote-as external
 neighbor ce2/2 as-origination-interval 1
 neighbor ce2/2 advertisement-interval 0
 neighbor ce2/2 fall-over bfd
 neighbor ce2/1 remote-as external
 neighbor ce2/1 as-origination-interval 1
 neighbor ce2/1 advertisement-interval 0
 neighbor ce2/1 fall-over bfd
 neighbor ce2/4 remote-as external
 neighbor ce2/4 as-origination-interval 1
 neighbor ce2/4 advertisement-interval 0
 neighbor ce2/4 fall-over bfd
 exit-unnumbered-mode
```

```
 !
 address-family ipv4 unicast
 max-paths ebgp 4
 redistribute connected
 redistribute static
 !
 bgp v4-unnumbered-mode
 neighbor ce2/3 activate
 neighbor ce2/2 activate
 neighbor ce2/1 activate
 neighbor ce2/4 activate
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 !
 bgp l2vpn-unnumbered-mode
 neighbor ce2/3 activate
 neighbor ce2/2 activate
 neighbor ce2/1 activate
 neighbor ce2/4 activate
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
 exit
!
```

Validation

Execute the following show command to display information about BGP in VTEP1.

VTEP1

```
VTEP1#show ip bgp summary
BGP router identifier 10.10.10.1, local AS number 65001
BGP table version is 6
3 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe2(fe80::5e17:83ff:fece:3079)          4 65009        432        344       6      0      0   00:08:55             4
xe1(fe80::5e17:83ff:fecf:3174)          4 65009        459        338       6      0      0   00:08:55             4

Total number of neighbors 2

Total number of Established sessions 2
VTEP1#
VTEP1#
```

Execute the following show command to verify route and ECMP information between VTEPs.

```
VTEP1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.1/32 is directly connected, lo, installed 03:20:18, last update 03:20:18 ago
B            10.10.10.2/32 [20/0] via fe80::5e17:83ff:fecf:3174, xe1, installed 00:08:53, last update 00:08:53 ago
                           [20/0] via fe80::5e17:83ff:fece:3079, xe2
B            10.10.10.3/32 [20/0] via fe80::5e17:83ff:fecf:3174, xe1, installed 00:08:53, last update 00:08:53 ago
B            10.10.10.4/32 [20/0] via fe80::5e17:83ff:fece:3079, xe2, installed 00:08:53, last update 00:08:53 ago
B            10.10.10.5/32 [20/0] via fe80::5e17:83ff:fecf:3174, xe1, installed 00:08:53, last update 00:08:53 ago
                           [20/0] via fe80::5e17:83ff:fece:3079, xe2
B            10.10.10.6/32 [20/0] via fe80::5e17:83ff:fecf:3174, xe1, installed 00:08:53, last update 00:08:53 ago
                           [20/0] via fe80::5e17:83ff:fece:3079, xe2
C            127.0.0.0/8 is directly connected, lo, installed 03:20:36, last update 03:20:36 ago

Gateway of last resort is not set
VTEP1#
```

Execute the following show command to view the BGP EVPN entries.

```
VTEP1#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.1, local AS number 65001
BGP table version is 37
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                                 V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD
MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
xe2(fe80::5e17:83ff:fece:3079)           4 65009        432        344      31      0      0  00:08:55           20      9
4      6      1      0
xe1(fe80::5e17:83ff:fecf:3174)           4 65009        459        338      36      0      0  00:08:55           20      9
4      6      1      0

Total number of neighbors 2

Total number of Established sessions 2
VTEP1#
```

Execute the following show command to view the BGP EVPN entries specific to VRF.

```
VTEP1#show bgp l2vpn evpn vrf vpn_green
BGP table version is 1, local router ID is 10.10.10.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best,
i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
```

```
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap
*  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[100]:[100]
                    10.10.10.2          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN
*>                  10.10.10.1          0         100       32768  i  ----------      VXLAN
*  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294967295]:[0]
                    10.10.10.2          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN
*  [1]:[00:00:00:cc:cc:dd:dd:00:00:00]:[100]:[100]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN
*                   10.10.10.5          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN
*  [1]:[00:00:00:cc:cc:dd:dd:00:00:00]:[4294967295]:[0]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN
*                   10.10.10.5          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN
*  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[100]:[48,0010:9410:0101]:[0]:[100]
                    10.10.10.2          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3079 VXLAN
*>                  10.10.10.1          0         100       32768  i  ----------      VXLAN
*  [2]:[00:00:00:cc:cc:dd:dd:00:00:00]:[100]:[48,0010:9410:0201]:[0]:[100]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3079 VXLAN
*                   10.10.10.5          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3079 VXLAN
*> [2]:[0]:[100]:[48,0010:9450:0101]:[0]:[100]
                    10.10.10.1          0         100       32768  i  ----------      VXLAN
*  [2]:[0]:[100]:[48,0010:9450:0201]:[0]:[100]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3079 VXLAN
*> [3]:[100]:[32,10.10.10.1]
                    10.10.10.1          0         100       32768  i  ----------      VXLAN
*  [3]:[100]:[32,10.10.10.2]
                    10.10.10.2          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN
*  [3]:[100]:[32,10.10.10.5]
                    10.10.10.5          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN
*  [3]:[100]:[32,10.10.10.6]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3174 VXLAN

Total number of prefixes 12
VTEP1#
```

Execute the following show command to view the source, destination, and status of the VxLAN tunnel entries.

```
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
10.10.10.1       10.10.10.5       Installed     00:30:09      00:30:09
10.10.10.1       10.10.10.2       Installed     00:30:09      00:30:09
10.10.10.1       10.10.10.6       Installed     00:30:09      00:30:09

Total number of entries are 3
VTEP1#
```

Execute the following show command to view the VxLAN information.

```
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
```

```
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN      DF-Status Src-Addr        Dst-Addr
Router-Mac
_____________________________________________________________________________________________________________________
_________________
100      ----         L2       NW   ----      ----                         ----      ----      10.10.10.1      10.10.10.5
100      ----         L2       NW   ----      ----                         ----      ----      10.10.10.1      10.10.10.2
100      ----         L2       NW   ----      ----                         ----      ----      10.10.10.1      10.10.10.6
100      ----         --       AC   xe11      --- Single Homed Port ---    51        ----      ----            ----
100      ----         --       AC   po100     00:00:00:aa:aa:bb:bb:00:00:00 101      NON-DF    ----            ----
200      ----         L2       NW   ----      ----                         ----      ----      10.10.10.1      10.10.10.5
200      ----         L2       NW   ----      ----                         ----      ----      10.10.10.1      10.10.10.2
200      ----         L2       NW   ----      ----                         ----      ----      10.10.10.1      10.10.10.6
200      ----         --       AC   xe11      --- Single Homed Port ---    151       ----      ----            ----
200      ----         --       AC   po100     00:00:00:aa:aa:bb:bb:00:00:00 201      NON-DF    ----            ----

Total number of entries are 10
VTEP1#
```

Execute the following show command to view the EVPN multihoming status information.

```
VTEP1#show evpn multihoming-status
Multihoming is ACTIVE in Hardware
VxLAN Multihoming is enabled
Total configured ESI count = 1

      VTEP1#show evpn multi-homing all
      ESI                               Access-IF      PE-IP-ADDRESS
      ============================================================
      00:00:00:aa:aa:bb:bb:00:00:00   po100           1.1.1.1
      00:00:00:aa:aa:bb:bb:00:00:00   ----            2.2.2.2
      00:00:00:cc:cc:dd:dd:00:00:00   ----            5.5.5.5
      00:00:00:cc:cc:dd:dd:00:00:00   ----            6.6.6.6

      Total number of entries are 4
```

Execute the following show command to display the configuration of the access interfaces.

```
VTEP1#show nvo vxlan access-if brief

                     Inner                          Admin   Link
Interface  Vlan       vlan  Ifindex    Vnid        status  status
-------------------------------------------------------------------
xe11       51         ---   0x7a120    100         up      up
xe11       151        ---   0x7a121    200         up      up
po100      101        ---   0x7a122    100         up      up
po100      201        ---   0x7a123    200         up      up

Total number of entries are 4
VTEP1#
```

Execute the following show command to display the VxLAN MAC address table.

```
    VTEP1#show nvo vxlan mac-table
```

```
================================================================================================================
================================================================================
                                        VXLAN MAC Entries
================================================================================================================
=====================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                                Type           Status
MAC move AccessPortDesc LeafFlag
----------------------------------------------------------------------------------------------------------------
-------------------------------------

100 po100 101 ----0010.9410.0101 00:00:00:aa:aa:bb:bb:00:00:00 Dynamic Local------- 0 access_if_mh_vlan101_vpn_green  ----
100 ---- ---- ----0010.9410.0201 00:00:00:cc:cc:dd:dd:00:00:00 Dynamic Remote -------    0        -------              ----
100 xe11 51   ----0010.9450.0101 10.10.10.1            Dynamic Local -------    0   access_if_sh_vlan51_vpn_green     ----
100 ---- ---- ----0010.9450.0201 10.10.10.6            Dynamic Remote-------    0   -------          ----

Total number of entries are : 4

VTEP1#
VTEP1#show nvo vxlan mac-table hardware
================================================================================================================
================================================================================
                                                                VXLAN MAC Entries
================================================================================================================
=====================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                                Type           Status
Time-out AccessPortDesc
----------------------------------------------------------------------------------------------------------------
--------------------------

100 po100 101      ----   0010.9410.0101 00:00:00:aa:aa:bb:bb:00:00:00      Dynamic Local   -------    300
access_if_mh_vlan101_vpn_green
100 ---   ----     ----   0010.9410.0201 00:00:00:cc:cc:dd:dd:00:00:00      Remote          -------    300      -------
100 xe11 51        ----   0010.9450.0101 10.10.10.1                         Dynamic Local   -------    300
access_if_sh_vlan51_vpn_green
100 --- ----      ----   0010.9450.0201 10.10.10.6                         Remote          -------    300      -------

Total number of entries are 4
VTEP1#
VTEP1#show nvo vxlan mac-table summary
============================================================
                   VXLAN MAC Summary
============================================================

Total number of entries are : 4

VTEP1#
```

VTEP2

Execute the following show command to display information about BGP in VTEP2.

```
VTEP2#show ip bgp summary
BGP router identifier 10.10.10.2, local AS number 65001
BGP table version is 7
3 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe2(fe80::5e17:83ff:fece:3078)          4 65009        385        286       7      0      0  00:09:09             4
xe1(fe80::5e17:83ff:fecf:3173)          4 65009        415        339       7      0      0  00:09:09             4
```

```
Total number of neighbors 2

Total number of Established sessions 2
VTEP2#
VTEP2#
```

Execute the following show command to verify route and ECMP information between VTEPs.

```
VTEP2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::5e17:83ff:fecf:3173, xe1, installed 00:09:07, last update 00:09:07 ago
                           [20/0] via fe80::5e17:83ff:fece:3078, xe2
C            10.10.10.2/32 is directly connected, lo, installed 03:20:31, last update 03:20:31 ago
B            10.10.10.3/32 [20/0] via fe80::5e17:83ff:fecf:3173, xe1, installed 00:09:07, last update 00:09:07 ago
B            10.10.10.4/32 [20/0] via fe80::5e17:83ff:fece:3078, xe2, installed 00:09:07, last update 00:09:07 ago
B            10.10.10.5/32 [20/0] via fe80::5e17:83ff:fecf:3173, xe1, installed 00:09:07, last update 00:09:07 ago
                           [20/0] via fe80::5e17:83ff:fece:3078, xe2
B            10.10.10.6/32 [20/0] via fe80::5e17:83ff:fecf:3173, xe1, installed 00:09:07, last update 00:09:07 ago
                           [20/0] via fe80::5e17:83ff:fece:3078, xe2
C            127.0.0.0/8 is directly connected, lo, installed 03:20:50, last update 03:20:50 ago

Gateway of last resort is not set
VTEP2
```

Execute the following show command to view the BGP EVPN entries.

```
VTEP2#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.2, local AS number 65001
BGP table version is 24
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD
MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
xe2(fe80::5e17:83ff:fece:3078)          4 65009        385        286      24      0      0  00:09:09           21      9
5      6      1      0
xe1(fe80::5e17:83ff:fecf:3173)          4 65009        415        339      24      0      0  00:09:09           21      9
5      6      1      0

Total number of neighbors 2

Total number of Established sessions 2
VTEP2#
```

Execute the following show command to view the BGP EVPN entries specific to VRF.

```
VTEP2#show bgp l2vpn evpn vrf vpn_green
BGP table version is 1, local router ID is 10.10.10.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best,
i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap
*  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[100]:[100]
                     10.10.10.1          0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3173 VXLAN
*>                   10.10.10.2          0        100       32768  i  ----------      VXLAN
*  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294967295]:[0]
                     10.10.10.1          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*  [1]:[00:00:00:cc:cc:dd:dd:00:00:00]:[100]:[100]
                     10.10.10.6          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*                    10.10.10.5          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*  [1]:[00:00:00:cc:cc:dd:dd:00:00:00]:[4294967295]:[0]
                     10.10.10.6          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*                    10.10.10.5          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*> [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[100]:[48,0010:9410:0101]:[0]:[100]
                     10.10.10.2          0        100       32768  i  ----------      VXLAN
*                    10.10.10.1          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*  [2]:[00:00:00:cc:cc:dd:dd:00:00:00]:[100]:[48,0010:9410:0201]:[0]:[100]
                     10.10.10.6          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*                    10.10.10.5          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*  [2]:[0]:[100]:[48,0010:9450:0101]:[0]:[100]
                     10.10.10.1          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*  [2]:[0]:[100]:[48,0010:9450:0201]:[0]:[100]
                     10.10.10.6          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*  [3]:[100]:[32,10.10.10.1]
                     10.10.10.1          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*> [3]:[100]:[32,10.10.10.2]
                     10.10.10.2          0        100       32768  i  ----------      VXLAN
*  [3]:[100]:[32,10.10.10.5]
                     10.10.10.5          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN
*  [3]:[100]:[32,10.10.10.6]
                     10.10.10.6          0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3078 VXLAN

Total number of prefixes 12
VTEP2#
```

Execute the following show command to view the source, destination, and status of the VxLAN tunnel entries.

```
VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=========================================================================
10.10.10.2       10.10.10.5       Installed     00:31:22      00:31:22
10.10.10.2       10.10.10.1       Installed     00:31:22      00:31:22
10.10.10.2       10.10.10.6       Installed     00:31:22      00:31:22

Total number of entries are 3
VTEP2#
```

Execute the following show command to view the VxLAN information.

```
VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
```

```
          (u) - Untagged

VNID    VNI-Name    VNI-Type Type Interface ESI                          VLAN     DF-Status Src-Addr        Dst-Addr
Router-Mac
_____________________________________________________________________________________________________________________
_________________
100     ----        L2       NW   ----      ----                         ----     ----      10.10.10.2      10.10.10.5
100     ----        L2       NW   ----      ----                         ----     ----      10.10.10.2      10.10.10.1
100     ----        L2       NW   ----      ----                         ----     ----      10.10.10.2      10.10.10.6
100     ----        --       AC   po100     00:00:00:aa:aa:bb:bb:00:00:00 101      DF        ----            ----
200     ----        L2       NW   ----      ----                         ----     ----      10.10.10.2      10.10.10.5
200     ----        L2       NW   ----      ----                         ----     ----      10.10.10.2      10.10.10.1
200     ----        L2       NW   ----      ----                         ----     ----      10.10.10.2      10.10.10.6
200     ----        --       AC   po100     00:00:00:aa:aa:bb:bb:00:00:00 201      DF        ----            ----

Total number of entries are 8
VTEP2#
```

VTEP 3

Execute the following show command to view the EVPN multihoming status information.

```
VTEP2#show evpn multihoming-status
Multihoming is ACTIVE in Hardware
VxLAN Multihoming is enabled
Total configured ESI count = 1
VTEP2#
VTEP2#show evpn multi-homing all
ESI                             Access-IF    PE-IP-ADDRESS
============================================================
00:00:00:aa:aa:bb:bb:00:00:00  ----         10.10.10.1
00:00:00:aa:aa:bb:bb:00:00:00  po100        10.10.10.2
00:00:00:cc:cc:dd:dd:00:00:00  ----         10.10.10.5
00:00:00:cc:cc:dd:dd:00:00:00  ----         10.10.10.6
Total number of entries are 4
VTEP2#
```

Execute the following show command to display the configuration of the access interfaces.

```
VTEP2#show nvo vxlan access-if brief

                      Inner                          Admin   Link
Interface  Vlan       vlan  Ifindex    Vnid         status  status
-------------------------------------------------------------
po100      101        ---   0x7a120    100          up      up
po100      201        ---   0x7a121    200          up      up

Total number of entries are 2
VTEP2#
```

Execute the following show command to display the VxLAN MAC address table.

```
VTEP2#show nvo vxlan mac-table
=============================================================================================
=============================================================================
```

```
                                                    VXLAN MAC Entries
================================================================================================
================================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                              Type           Status
MAC move AccessPortDesc LeafFlag
------------------------------------------------------------------------------------------------------------------------
------------------------------------

100 po100 101 ---- 0010.9410.0101 00:00:00:aa:aa:bb:bb:00:00:00 Dynamic Local ------- 0 access_if_mh_vlan101_vpn_green   ----
100 ---- ---- ----0010.9410.0201 00:00:00:cc:cc:dd:dd:00:00:00  Dynamic Remote ------- 0  -------          ----
100 ---- ---- ---- 0010.9450.0101 10.10.10.1                    Dynamic Remote  ------- 0  -------          ----
100 ---- ---- ---- 0010.9450.0201 10.10.10.6                    Dynamic Remote  -------    0        -------          ----

Total number of entries are : 4

VTEP2#
VTEP2#show nvo vxlan mac-table hardware
================================================================================================
================================================================================
                                                    VXLAN MAC Entries
================================================================================================================================
==================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                              Type           Status
Time-out AccessPortDesc
--------------------------------------------------------------------------------------------------------------------------------
--------------------------

100 po100 101 ----  0010.9410.0101 00:00:00:aa:aa:bb:bb:00:00:00 Dynamic Local ------- 300     access_if_mh_vlan101_vpn_green
100 ---   ---- ---- 0010.9410.0201 00:00:00:cc:cc:dd:dd:00:00:00 Remote       ------- 300     -------
100 ---   ---- ---- 0010.9450.0101 10.10.10.1                     Remote       ------- 300     -------
100 ---   ---- ---- 0010.9450.0201 10.10.10.6                     Remote       ------- 300     -------

Total number of entries are 4
VTEP2#
VTEP2#show nvo vxlan mac-table summary
============================================================
                    VXLAN MAC Summary
============================================================

Total number of entries are : 4

VTEP2#
```

VTEP3

Execute the following show command to display information about BGP in VTEP3.

```
VTEP3#show ip bgp summary
BGP router identifier 10.10.10.5, local AS number 65001
BGP table version is 7
3 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe2(fe80::5e17:83ff:fece:3076)          4 65009        390        290       7      0      0  00:09:22             4
xe1(fe80::5e17:83ff:fecf:3172)          4 65009        417        336       7      0      0  00:09:22             4
```

```
Total number of neighbors 2

Total number of Established sessions 2
VTEP3#
```

Execute the following show command to verify route and ECMP information between VTEPs

```
VTEP3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B           10.10.10.1/32 [20/0] via fe80::5e17:83ff:fecf:3172, xe1, installed 00:09:20, last update 00:09:20 ago
                          [20/0] via fe80::5e17:83ff:fece:3076, xe2
B           10.10.10.2/32 [20/0] via fe80::5e17:83ff:fecf:3172, xe1, installed 00:09:20, last update 00:09:20 ago
                          [20/0] via fe80::5e17:83ff:fece:3076, xe2
B           10.10.10.3/32 [20/0] via fe80::5e17:83ff:fecf:3172, xe1, installed 00:09:20, last update 00:09:20 ago
B           10.10.10.4/32 [20/0] via fe80::5e17:83ff:fece:3076, xe2, installed 00:09:20, last update 00:09:20 ago
C           10.10.10.5/32 is directly connected, lo, installed 03:20:45, last update 03:20:45 ago
B           10.10.10.6/32 [20/0] via fe80::5e17:83ff:fecf:3172, xe1, installed 00:09:20, last update 00:09:20 ago
                          [20/0] via fe80::5e17:83ff:fece:3076, xe2
C           127.0.0.0/8 is directly connected, lo, installed 03:21:03, last update 03:21:03 ago

Gateway of last resort is not set
VTEP3
```

Execute the following show command to view the BGP EVPN entries.

```
VTEP3#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.5, local AS number 65001
BGP table version is 32
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                              V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd  AD  MACIP
MCAST    ESI  PREFIX-ROUTE  Desc
xe2(fe80::5e17:83ff:fece:3076)          4 65009        390        290      25      0      0  00:09:22           21      9
5      6      1      0
xe1(fe80::5e17:83ff:fecf:3172)          4 65009        417        336      31      0      0  00:09:22           21      9
5      6      1      0

Total number of neighbors 2

Total number of Established sessions 2
VTEP3#
VTEP3
```

Execute the following show command to view the BGP EVPN entries specific to VRF.

```
VTEP3#show bgp l2vpn evpn vrf vpn_green
BGP table version is 1, local router ID is 10.10.10.5
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best,
i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap
*  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[100]:[100]
                    10.10.10.1          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3172 VXLAN
*                   10.10.10.2          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294967295]:[0]
                    10.10.10.1          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3172 VXLAN
*                   10.10.10.2          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*  [1]:[00:00:00:cc:cc:dd:dd:00:00:00]:[100]:[100]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*>                  10.10.10.5          0         100       32768  i  ----------      VXLAN
*  [1]:[00:00:00:cc:cc:dd:dd:00:00:00]:[4294967295]:[0]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[100]:[48,0010:9410:0101]:[0]:[100]
                    10.10.10.2          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*                   10.10.10.1          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*  [2]:[00:00:00:cc:cc:dd:dd:00:00:00]:[100]:[48,0010:9410:0201]:[0]:[100]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*>                  10.10.10.5          0         100       32768  i  ----------      VXLAN
*  [2]:[0]:[100]:[48,0010:9450:0101]:[0]:[100]
                    10.10.10.1          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*  [2]:[0]:[100]:[48,0010:9450:0201]:[0]:[100]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*  [3]:[100]:[32,10.10.10.1]
                    10.10.10.1          0         100       0   65009 65001 i  fe80::5e17:83ff:fecf:3172 VXLAN
*  [3]:[100]:[32,10.10.10.2]
                    10.10.10.2          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN
*> [3]:[100]:[32,10.10.10.5]
                    10.10.10.5          0         100       32768  i  ----------      VXLAN
*  [3]:[100]:[32,10.10.10.6]
                    10.10.10.6          0         100       0   65009 65001 i  fe80::5e17:83ff:fece:3076 VXLAN

Total number of prefixes 12
```

Execute the following show command to view the source, destination, and status of the VxLAN tunnel entries.

```
VTEP3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status         Up/Down       Update
=========================================================================
10.10.10.5       10.10.10.2       Installed      00:31:30      00:31:30
10.10.10.5       10.10.10.1       Installed      00:31:30      00:31:30
10.10.10.5       10.10.10.6       Installed      00:31:30      00:31:30

Total number of entries are 3
```

```
VTEP3#
```

Execute the following show command to view the VxLAN information.

```
VTEP3#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID    VNI-Name    VNI-Type Type Interface ESI                         VLAN     DF-Status Src-Addr        Dst-Addr
Router-Mac
_____________________________________________________________________________________________________________
_________________
100     ----        L2       NW   ----      ----                        ----     ----      10.10.10.5      10.10.10.2
100     ----        L2       NW   ----      ----                        ----     ----      10.10.10.5      10.10.10.1
100     ----        L2       NW   ----      ----                        ----     ----      10.10.10.5      10.10.10.6
100     ----        --       AC   po200     00:00:00:cc:cc:dd:dd:00:00:00 101      NON-DF    ----            ----
200     ----        L2       NW   ----      ----                        ----     ----      10.10.10.5      10.10.10.2
200     ----        L2       NW   ----      ----                        ----     ----      10.10.10.5      10.10.10.1
200     ----        L2       NW   ----      ----                        ----     ----      10.10.10.5      10.10.10.6
200     ----        --       AC   po200     00:00:00:cc:cc:dd:dd:00:00:00 201      NON-DF    ----            ----

Total number of entries are 8
VTEP3#
```

Execute the following show command to view the EVPN multihoming status information.

```
VTEP3#show evpn multihoming-status
Multihoming is ACTIVE in Hardware
VxLAN Multihoming is enabled
Total configured ESI count = 1
VTEP3#
VTEP3#show evpn multi-homing all
ESI                             Access-IF    PE-IP-ADDRESS
============================================================
00:00:00:aa:aa:bb:bb:00:00:00   ----         10.10.10.1
00:00:00:aa:aa:bb:bb:00:00:00   ----         10.10.10.2
00:00:00:cc:cc:dd:dd:00:00:00   po200        10.10.10.5
00:00:00:cc:cc:dd:dd:00:00:00   ----         10.10.10.6
Total number of entries are 4
VTEP3#
```

Execute the following show command to display the configuration of the access interfaces.

```
VTEP3#show nvo vxlan access-if brief

                      Inner                          Admin   Link
Interface  Vlan       vlan  Ifindex    Vnid         status  status
-------------------------------------------------------------------
po200      101        ---   0x7a120    100          up      up
po200      201        ---   0x7a121    200          up      up
```

```
Total number of entries are 2
VTEP3#
```

Execute the following show command to display the VxLAN MAC address table.

```
VTEP3#show nvo vxlan mac-table
================================================================================================================================================
====================================
                                                    VXLAN MAC Entries
================================================================================================================================================
====================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                                  Type            Status
MAC move AccessPortDesc LeafFlag
------------------------------------------------------------------------------------------------------------------------------------------------
------------------------------------

100 ---- --- ---- 0010.9410.0101 00:00:00:aa:aa:bb:bb:00:00:00  Dynamic Remote ------- 0 -------        ----
100 po200 101 ---- 0010.9410.0201 00:00:00:cc:cc:dd:dd:00:00:00 Dynamic Local  -------  0 access_if_mh_vlan101_vpn_green ----
100 ---- ---- ---- 0010.9450.0101 10.10.10.1                    Dynamic Remote  -------    0        -------          ----
100 ---- ---- ---- 0010.9450.0201 10.10.10.6                    Dynamic Remote  -------    0        -------          ----

Total number of entries are : 4

VTEP3#


VTEP3#show nvo vxlan mac-table hardware
================================================================================================================================================
====================================
                                                    VXLAN MAC Entries
================================================================================================================================================
====================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                                  Type            Status
Time-out AccessPortDesc
------------------------------------------------------------------------------------------------------------------------------------------------
------------------------------

100 --- ---- ---- 0010.9410.0101 00:00:00:aa:aa:bb:bb:00:00:00 Remote        -------   300      -------
100 po200 101 ---- 0010.9410.0201 00:00:00:cc:cc:dd:dd:00:00:00 Dynamic Local  -------   300 access_if_mh_vlan101_vpn_green
100 --- ---- ---- 0010.9450.0101 10.10.10.1                    Remote        -------   300      -------
100 --- ---- ---- 0010.9450.0201 10.10.10.6                    Remote        -------   300      -------

Total number of entries are 4
VTEP3#
VTEP3#show nvo vxlan mac-table summary
============================================================
                    VXLAN MAC Summary
============================================================

Total number of entries are : 4
VTEP3#
```

VTEP4

Execute the following show command to display information about BGP in VTEP4.

```
VTEP4#show ip bgp summary
BGP router identifier 10.10.10.6, local AS number 65001
BGP table version is 6
3 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe2(fe80::5e17:83ff:fece:3077)          4 65009        355        274       6      0      0  00:09:34             4
xe1(fe80::5e17:83ff:fecf:3175)          4 65009        356        286       6      0      0  00:09:34             4
```

```
Total number of neighbors 2

Total number of Established sessions 2
VTEP4#
```

Execute the following show command to verify route and ECMP information between VTEPs

```
VTEP4#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::5e17:83ff:fecf:3175, xe1, installed 00:09:32, last update 00:09:32 ago
                           [20/0] via fe80::5e17:83ff:fece:3077, xe2
B            10.10.10.2/32 [20/0] via fe80::5e17:83ff:fecf:3175, xe1, installed 00:09:32, last update 00:09:32 ago
                           [20/0] via fe80::5e17:83ff:fece:3077, xe2
B            10.10.10.3/32 [20/0] via fe80::5e17:83ff:fecf:3175, xe1, installed 00:09:32, last update 00:09:32 ago
B            10.10.10.4/32 [20/0] via fe80::5e17:83ff:fece:3077, xe2, installed 00:09:32, last update 00:09:32 ago
B            10.10.10.5/32 [20/0] via fe80::5e17:83ff:fecf:3175, xe1, installed 00:09:32, last update 00:09:32 ago
                           [20/0] via fe80::5e17:83ff:fece:3077, xe2
C            10.10.10.6/32 is directly connected, lo, installed 03:20:57, last update 03:20:57 ago
C            127.0.0.0/8 is directly connected, lo, installed 03:21:15, last update 03:21:15 ago

Gateway of last resort is not set
VTEP4#
```

Execute the following show command to view the BGP EVPN entries.

```
VTEP4#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.6, local AS number 65001
BGP table version is 15
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD
MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
xe2(fe80::5e17:83ff:fece:3077)          4 65009        355        274       9      0      0  00:09:34           20      9
4      6      1      0
xe1(fe80::5e17:83ff:fecf:3175)          4 65009        356        286      14      0      0  00:09:34           20      9
4      6      1      0

Total number of neighbors 2

Total number of Established sessions 2
VTEP4#
```

Execute the following show command to view the BGP EVPN entries specific to VRF.

```
VTEP4#show bgp l2vpn evpn vrf vpn_green
BGP table version is 1, local router ID is 10.10.10.6
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
```

```
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap
*  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[100]:[100]
                    10.10.10.1           0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*                   10.10.10.2           0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294967295]:[0]
                    10.10.10.2           0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*                   10.10.10.1           0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*  [1]:[00:00:00:cc:cc:dd:dd:00:00:00]:[100]:[100]
                    10.10.10.5           0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*>                  10.10.10.6           0        100       32768  i  ----------      VXLAN
*  [1]:[00:00:00:cc:cc:dd:dd:00:00:00]:[4294967295]:[0]
                    10.10.10.5           0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[100]:[48,0010:9410:0101]:[0]:[100]
                    10.10.10.2           0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3077 VXLAN
*                   10.10.10.1           0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3077 VXLAN
*  [2]:[00:00:00:cc:cc:dd:dd:00:00:00]:[100]:[48,0010:9410:0201]:[0]:[100]
                    10.10.10.5           0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3077 VXLAN
*>                  10.10.10.6           0        100       32768  i  ----------      VXLAN
*  [2]:[0]:[100]:[48,0010:9450:0101]:[0]:[100]
                    10.10.10.1           0        100       0   65009 65001 i  fe80::5e17:83ff:fece:3077 VXLAN
*> [2]:[0]:[100]:[48,0010:9450:0201]:[0]:[100]
                    10.10.10.6           0        100       32768  i  ----------      VXLAN
*  [3]:[100]:[32,10.10.10.1]
                    10.10.10.1           0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*  [3]:[100]:[32,10.10.10.2]
                    10.10.10.2           0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*  [3]:[100]:[32,10.10.10.5]
                      10.10.10.5         0        100       0   65009 65001 i  fe80::5e17:83ff:fecf:3175 VXLAN
*> [3]:[100]:[32,10.10.10.6]
                      10.10.10.6         0        100       32768  i  ----------      VXLAN

Total number of prefixes 12
VTEP4#
```

Execute the following show command to view the source, destination, and status of the VxLAN tunnel entries..

```
VTEP4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=========================================================================
10.10.10.6       10.10.10.5       Installed     00:31:37      00:31:37
10.10.10.6       10.10.10.2       Installed     00:31:37      00:31:37
10.10.10.6       10.10.10.1       Installed     00:31:37      00:31:37

Total number of entries are 3
VTEP4#
```

Execute the following show command to view the VxLAN information.

```
VTEP4#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN      DF-Status Src-Addr        Dst-Addr
Router-Mac
_____________________________________________________________________________________________________________________
_____________________
100      ----         L2       NW   ----      ----                         ----      ----      10.10.10.6      10.10.10.5
100      ----         L2       NW   ----      ----                         ----      ----      10.10.10.6      10.10.10.2
100      ----         L2       NW   ----      ----                         ----      ----      10.10.10.6      10.10.10.1
100      ----         --       AC   xe11      --- Single Homed Port ---    51        ----      ----            ----
100      ----         --       AC   po200     00:00:00:cc:cc:dd:dd:00:00:00 101      DF        ----            ----
200      ----         L2       NW   ----      ----                         ----      ----      10.10.10.6      10.10.10.5
200      ----         L2       NW   ----      ----                         ----      ----      10.10.10.6      10.10.10.2
200      ----         L2       NW   ----      ----                         ----      ----      10.10.10.6      10.10.10.1
200      ----         --       AC   xe11      --- Single Homed Port ---    151       ----      ----            ----
```

```
200      ----         --       AC   po200     00:00:00:cc:cc:dd:dd:00:00:00 201       DF        ----            ----

Total number of entries are 10
VTEP4#
```

Execute the following show command to view the EVPN multihoming status information.

```
VTEP4#show evpn multihoming-status
Multihoming is ACTIVE in Hardware
VxLAN Multihoming is enabled
Total configured ESI count = 1
VTEP4#
VTEP4#show evpn multi-homing all
ESI                           Access-IF    PE-IP-ADDRESS
===========================================================
00:00:00:aa:aa:bb:bb:00:00:00  ----         10.10.10.1
00:00:00:aa:aa:bb:bb:00:00:00  ----         10.10.10.2
00:00:00:cc:cc:dd:dd:00:00:00  ----         10.10.10.5
00:00:00:cc:cc:dd:dd:00:00:00  po200        10.10.10.6
Total number of entries are 4
VTEP4#
```

Execute the following show command to display the configuration of the access interfaces.

```
VTEP4#show nvo vxlan access-if brief
                     Inner                     Admin   Link
Interface  Vlan      vlan  Ifindex    Vnid      status  status
---------------------------------------------------------------
xe11       51        ---   0x7a120    100       up      up
xe11       151       ---   0x7a121    200       up      up
po200      101       ---   0x7a122    100       up      up
po200      201       ---   0x7a123    200       up      up

Total number of entries are 4
VTEP4#
```

Execute the following show command to display the VxLAN MAC address table.

```
VTEP4#show nvo vxlan mac-table
=================================================================================================================================
======================================
                                                       VXLAN MAC Entries
=================================================================================================================================
======================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                                 Type            Status
MAC move AccessPortDesc LeafFlag
_________________________________________________________________________________________________________________________________
______________________________________

100 ---- ---- ---- 0010.9410.0101 00:00:00:aa:aa:bb:bb:00:00:00Dynamic Remote -------  0 -------          ----
100 po200 101 ---- 0010.9410.0201 00:00:00:cc:cc:dd:dd:00:00:00Dynamic Local -------   0 access_if_mh_vlan101_vpn_green ----
100 ---- ---- ---- 0010.9450.0101 10.10.10.1         Dynamic Remote -------   0       -------                         ----
100 xe11 51   ---- 0010.9450.0201 10.10.10.6         Dynamic Local  -------   0      access_if_sh_vlan51_vpn_green ----

Total number of entries are : 4
```

```
VTEP4#

VTEP4#show nvo vxlan mac-table hardware
================================================================================================================================
=========================                    VXLAN MAC Entries
================================================================================================================================
VNID  Interface VlanId  In-VlanId Mac-Addr   VTEP-Ip/ESI         Type   Status              Time-out   AccessPortDesc
________________________________________________________________________________________________________________________________

100  ---  ---- ---- 0010.9410.0101 00:00:00:aa:aa:bb:bb:00:00:00 Remote        -------  300   -------
100  po200 101 ---- 0010.9410.0201 00:00:00:cc:cc:dd:dd:00:00:00 Dynamic Local -------  300    access_if_mh_vlan101_vpn_green
100  ---  ---- ---- 0010.9450.0101 10.10.10.1                    Remote        -------  300    -------
100  xe11 51   ---- 0010.9450.0201 10.10.10.6                    Dynamic Local -------  300    access_if_sh_vlan51_vpn_green

Total number of entries are 4
VTEP4#
VTEP4#show nvo vxlan mac-table summary
============================================================
                    VXLAN MAC Summary
============================================================

Total number of entries are : 4

VTEP4#
```

SPINE1

Execute the following show command to display information about BGP in SPINE1.

```
SPINE1#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.3, local AS number 65009
BGP table version is 20
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD
MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
ce1/2(fe80::7e8d:9cff:fec3:6932)        4 65001        330        392      18      0      0  00:09:44            7      3
1      2      1      0
ce1/4(fe80::7e8d:9cff:fec9:6332)        4 65001        317        400      20      0      0  00:09:44            8      3
2      2      1      0
ce1/1(fe80::7e8d:9cff:fed4:7e32)        4 65001        326        393      20      0      0  00:09:44            7      3
1      2      1      0
ce1/3(fe80::7e8d:9cff:fed5:7f32)        4 65001        329        436      20      0      0  00:09:44            8      3
2      2      1      0

Total number of neighbors 4

Total number of Established sessions 4
SPINE1#
```

Execute the following show command to verify route informations.

```
SPINE1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B           10.10.10.1/32 [20/0] via fe80::7e8d:9cff:fed5:7f32, ce1/3, installed 00:09:41, last update 00:09:41 ago
B           10.10.10.2/32 [20/0] via fe80::7e8d:9cff:fec3:6932, ce1/2, installed 00:09:41, last update 00:09:41 ago
C           10.10.10.3/32 is directly connected, lo, installed 03:20:50, last update 03:20:50 ago
B           10.10.10.5/32 [20/0] via fe80::7e8d:9cff:fed4:7e32, ce1/1, installed 00:09:41, last update 00:09:41 ago
B           10.10.10.6/32 [20/0] via fe80::7e8d:9cff:fec9:6332, ce1/4, installed 00:09:41, last update 00:09:41 ago
C           127.0.0.0/8 is directly connected, lo, installed 03:21:17, last update 03:21:17 ago

Gateway of last resort is not set
SPINE1#

SPINE2:
```

Execute the following show command to display information about BGP in SPINE2.

```
SPINE2#show ip bgp summary
BGP router identifier 10.10.10.4, local AS number 65009
BGP table version is 11
2 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
ce2/3(fe80::7e8d:9cff:fec3:6933)        4 65001        285        389      11      0      0  00:09:53             1
ce2/2(fe80::7e8d:9cff:fec9:6333)        4 65001        315        393      11      0      0  00:09:53             1
ce2/1(fe80::7e8d:9cff:fed4:7e33)        4 65001        289        391      11      0      0  00:09:53             1
ce2/4(fe80::7e8d:9cff:fed5:7f33)        4 65001        344        434      11      0      0  00:09:53             1

Total number of neighbors 4

Total number of Established sessions 4
SPINE2#
SPINE2#
```

Execute the following show command to verify route informations.

```
SPINE2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
```

```
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::7e8d:9cff:fed5:7f33, ce2/4, installed 00:09:52, last update 00:09:52 ago
B            10.10.10.2/32 [20/0] via fe80::7e8d:9cff:fec3:6933, ce2/3, installed 00:09:52, last update 00:09:52 ago
C            10.10.10.4/32 is directly connected, lo, installed 03:21:02, last update 03:21:02 ago
B            10.10.10.5/32 [20/0] via fe80::7e8d:9cff:fed4:7e33, ce2/1, installed 00:09:52, last update 00:09:52 ago
B            10.10.10.6/32 [20/0] via fe80::7e8d:9cff:fec9:6333, ce2/2, installed 00:09:52, last update 00:09:52 ago
C            127.0.0.0/8 is directly connected, lo, installed 03:21:28, last update 03:21:28 ago

Gateway of last resort is not set
SPINE2#
```

CHAPTER 2 VXLAN Multi-homing Configuration

This chapter contains the configurations for VXLAN Multi-homing feature.

Overview

EVPN VxLAN Multi-homing features enables to connect a CE/Host node to two VTEPs with all-active redundancy mode.EVPN Multi-homing helps in VTEP to host failure and VTEP failure. If one VTEP goes down, other will forward the entire traffic.

Below are Multi-homing concepts:

- Ethernet Segment: Set of links which connect host/CE to two active-active multi-homed VTEP (only two VTEPs are supported) which appears as LACP link for host.
- Ethernet Segment Identifier: Ethernet Segment Identifier (ESI) which is an 10 octet-value, which can be configured in two ways, system mac is configured as esi in case of Dynamic Lag and 10-octet ESI format config is used on physical interface ES.
- Ethernet Segment Route (ES route): When a multi-homed CE is configured as an VXLAN access-port, Ethernet segment route is sent. The main purpose of this route is to discover other VTEPs which share the ES and to perform DF election.
- Ethernet A-D route per ESI: This route is used for Fast Convergence and Split Horizon.
- Ethernet A-D route per EVI: This route is used for load sharing between DF and NON-DF by the remote VTEPs

Topology

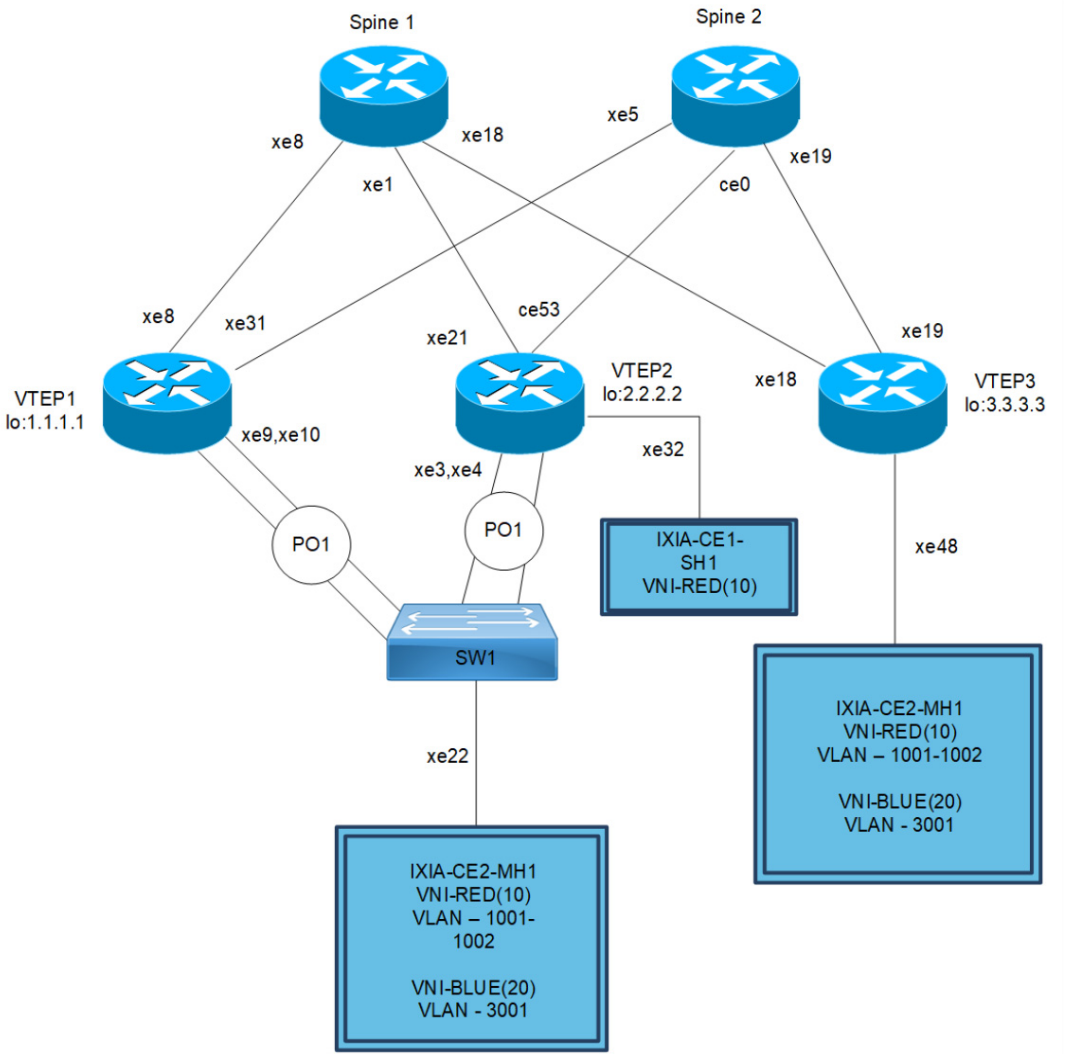

Figure 2-5: VxLAN-Multihoming

Note: Enable VXLAN MUTIHOMING before executing any configurations.

EVPN-VxLAN MH Configuration

ESI can be configured in below two ways

Ethernet Segment through Dynamic Lag interface

`#configure terminal`	Enter configure mode.
`(config)#interface po1`	Enter interface mode for po1
`(config-if)#switchport`	Make it L2 interface
`(config-if)#evpn multi-homed system-mac 8899.4400.6745`	Configure system mac as ESI value for Lag (po1) interface

`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

OR

Ethernet Segment through Physical interface

`#configure terminal`	Enter configure mode.
`(config)#interface xe41`	Enter interface mode for xe41
`(config-if)#switchport`	Make it L2 interface
`(config-if)# evpn multi-homed esi 00:01:02:03:04:05:06:07:08`	Configure 9-octet ESI value for xe41 interface (in static config, out of 10-octet ESI value, first octet is reserved)
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

VTEP1

(Multi-homed group1) – Part of both Multi-homed with po1 (MH1)

Generic Configuration

`#configure terminal`	Enter Configure mode.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#qos enable`	Enabling QoS
`(config)#commit`	Commit the candidate configuration to the running configuration

Interface and Loopback Configuration

`(config)#interface po1`	Enter Interface mode for po1 (MH1)
`(config-if)#switchport`	Make it L2 interface
`(config-if)# evpn multi-homed system-mac 0000.0000.1111`	Configure system MAC as ESI value for LAG (po1) interface
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe9`	Enter Interface mode for xe9
`(config-if)#channel-group 1 mode active`	Make it member port of po1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe10`	Enter Interface mode for xe10
`(config-if)#channel-group 1 mode active`	Make it member port of po1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure loopback ip address as 1.1.1.1 for VTEP1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe8`	Enter Interface mode for xe8
`(config-if)#ip address 10.10.10.1/24`	Configure IP address as 10.10.10.1 on network side of Spine1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe31`	Enter Interface mode for xe31
`(config-if)#ip address 20.20.20.1/24`	Configure IP address as 20.20.20.1 on network side of Spine2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

OSPF Configuration

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 1.1.1.1`	Configure router-id as 1.1.1.1 (lo ip address)
`(config-router)#network 1.1.1.1/32 area 0.0.0.0`	Add 1.1.1.1 (lo IP address) network into area 0
`(config-router)#network 10.10.10.0/24 area 0.0.0.0`	Add 10.10.10.0 (Spine1) network into area 0
`(config-router)#network 20.20.20.0/24 area 0.0.0.0`	Add 20.20.20.0 (Spine2) network into area 0

(config-router)#bfd all-interfaces	Enabling BFD on all OSPF interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

BGP Configuration

(config)#router bgp 500	Enter into Router BGP mode
(config-router)#bgp router-id 1.1.1.1	Configure router-id as 1.1.1.1 (lo IP address)
(config-router)#neighbor 2.2.2.2 remote-as 500	Specify a VTEP2 loopback IP address and remote-as defined
(config-router)#neighbor 2.2.2.2 update-source lo	Configure update as loopback for VTEP2
(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP2
(config-router)#neighbor 3.3.3.3 remote-as 500	Specify a VTEP3 loopback IP address and remote-as defined
(config-router)#neighbor 3.3.3.3 update-source lo	Configure update as loopback for VTEP3
(config-router)#neighbor 3.3.3.3 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP3
(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family mode
(config-router-af)#network 1.1.1.1/32	Advertise loopback network into BGP for VTEP ID reachability
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2 (VTEP2) into ipv4 unicast address family mode
(config-router-af)#neighbor 3.3.3.3 activate	Activate 3.3.3.3 (VTEP2) into ipv4 unicast address family mode
(config-router-af)#exit-address-family	Exit from ipv4 unicast address family mode
(config-router)#address-family l2vpn evpn	Enter into L2VPN EVPN address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2 (VTEP2) into L2VPN evpn address family mode
(config-router-af)#neighbor 3.3.3.3 activate	Activate 3.3.3.3 (VTEP3) into L2VPN evpn address family mode
(config-router-af)#exit-address-family	Exit from L2VPN address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to the running configuration

VRF Configuration

`(config)#mac vrf VRF1`	Create mac routing/forwarding instance with VRF1 name and enter into VRF mode
`(config-vrf)#rd 1.1.1.1:11`	Assign RD value
`(config-vrf)#route-target both 9.9.9.9:100`	Assign route-target value for same for import and export. Should be same on all node for VRF1
`(config-vrf)#exit`	Exit from VRF mode
`(config)#mac vrf VRF2`	Create MAC routing/forwarding instance with VRF1 name and enter into VRF mode
`(config-vrf)#rd 1.1.1.1:21`	Assign RD value
`(config-vrf)#route-target both 90.90.90.90:100`	Assign route-target value for same for import and export
`(config-vrf)#exit`	Exit from VRF mode
`(config)#commit`	Commit the candidate configuration to the running configuration

VxLAN Configuration

(config)#nvo vxlan enable	Enable VxLAN
(config)#evpn esi hold-time 90	Configure ESI hold time to allow tunnel to come up at the time of VxLAN initialization before making the ESI up
(config)#nvo vxlan vtep-ip-global 1.1.1.1	Configure Source VTEP-IP-global configuration
(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#vni-name VNI-RED	Configure VNI-name as VNI-RED
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 20 ingress-replication inner-vid-disabled	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#vni-name VNI-BLUE	Configure VNI-name as VNI-BLUE
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan access-if port-vlan po1 1001	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan po1 1002	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan po1 3001	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vni-name VNI-BLUE	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit from configuration mode

VTEP2

(Multi-homed group1) – Part of both Multi-homed with p01. And it has xe32 as single home access-if port (SH2)

Generic Configuration

`#configure terminal`	Enter Configure mode.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#qos enable`	Enabling QoS
`(config)#commit`	Commit the candidate configuration to the running configuration

Interface and Loopback Configuration

`(config)#interface po1`	Enter Interface mode for po1 (MH1)
`(config-if)#switchport`	Make it L2 interface
`(config-if)# evpn multi-homed system-mac 0000.0000.1111`	Configure system MAC as ESI value for LAG (po1) interface
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe3`	Enter Interface mode for xe3
`(config-if)#channel-group 1 mode active`	Make it member port of po1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe4`	Enter Interface mode for xe4
`(config-if)#channel-group 1 mode active`	Make it member port of po1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe32`	Enter Interface mode for xe32 (SH2)
`(config-if)#switchport`	Make it L2 interface
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 2.2.2.2/32 secondary`	Configure loopback IP address as 2.2.2.2 for VTEP2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe21`	Enter Interface mode for xe21
`(config-if)#ip address 30.30.30.1/24`	Configure IP address as 30.30.30.1 on network side of Spine1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface ce53`	Enter Interface mode for ce53
`(config-if)#ip address 40.40.40.1/24`	Configure IP address as 40.40.40.1 on network side of Spine2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

OSPF Configuration

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 2.2.2.2`	Configure router-id as 2.2.2.2 (lo IP address)
`(config-router)#network 2.2.2.2/32 area 0.0.0.0`	Add 2.2.2.2 (lo IP address) network into area 0
`(config-router)#network 30.30.30.0/24 area 0.0.0.0`	Add 30.30.30.0 (Spine1) network into area 0
`(config-router)#network 40.40.40.0/24 area 0.0.0.0`	Add 40.40.40.0 (Spine2) network into area 0

(config-router)#bfd all-interfaces	Enabling BFD on all OSPF interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

BGP Configuration

(config)#router bgp 500	Enter into Router BGP mode
(config-router)#bgp router-id 2.2.2.2	Configure router-id as 2.2.2.2 (lo IP address)
(config-router)#neighbor 1.1.1.1 remote-as 500	Specify a VTEP1 loopback IP address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Configure update as loopback for VTEP1
(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP1
(config-router)#neighbor 3.3.3.3 remote-as 500	Specify a VTEP3 loopback IP address and remote-as defined
(config-router)#neighbor 3.3.3.3 update-source lo	Configure update as loopback for VTEP3
(config-router)#neighbor 3.3.3.3 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP3
(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family mode
(config-router-af)#network 2.2.2.2/32	Advertise loopback network into BGP for VTEP ID reachability
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1 (VTEP2) into ipv4 unicast address family mode
(config-router-af)#neighbor 3.3.3.3 activate	Activate 3.3.3.3 (VTEP2) into ipv4 unicast address family mode
(config-router-af)#exit-address-family	Exit from ipv4 unicast address family mode
(config-router)#address-family l2vpn evpn	Enter into L2VPN EVPN address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1(VTEP1) into L2VPN evpn address family mode
(config-router-af)#neighbor 3.3.3.3 activate	Activate 3.3.3.3(VTEP3) into L2VPN evpn address family mode
(config-router-af)#exit-address-family	Exit from L2VPN address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to the running configuration

VRF Configuration

(config)# mac vrf VRF1	Create mac routing/forwarding instance with VRF1 name and enter into VRF mode
(config-vrf)#rd 2.2.2.2:11	Assign RD value
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for VRF1
(config-vrf)#exit	Exit from VRF mode
(config)#mac vrf VRF2	Create MAC routing/forwarding instance with VRF1 name and enter into VRF mode
(config-vrf)#rd 2.2.2.2:21	Assign RD value
(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#exit	Exit from VRF mode
(config)#commit	Commit the candidate configuration to the running configuration

VxLAN Configuration

(config)#nvo vxlan enable	Enable VxLAN
(config)#evpn esi hold-time 90	Configure ESI hold time to allow tunnel to come up at the time of VxLAN initialization before making the ESI up
(config)#nvo vxlan vtep-ip-global 2.2.2.2	Configure Source VTEP-IP-global configuration
(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#vni-name VNI-RED	Configure VNI-name as VNI-RED
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 20 ingress-replication inner-vid-disabled	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#vni-name VNI-BLUE	Configure VNI-name as VNI-BLUE
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan access-if port-vlan po1 1001	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan po1 1002	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan po1 3001	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port

(config-nvo-acc-if)#map vni-name VNI-BLUE	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port xe32	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit from configuration mode

VTEP3

It has xe48 as Single home access-if port (SH2)

Generic Configuration

#configure terminal	Enter Configure mode.
(config)#evpn vxlan multihoming enable	Enable Multihoming, save configs and reboot the board for multihoming to be effective
(config)#qos enable	Enabling QoS
(config)#commit	Commit the candidate configuration to the running configuration

Interface and loopback configuration

(config)#interface xe48	Enter Interface mode for xe48 (SH3)
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 3.3.3.3/32 secondary	Configure loopback IP address as 3.3.3.3 for VTEP3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe18	Enter Interface mode for xe18
(config-if)#ip address 50.50.50.1/24	Configure IP address as 50.50.50.1 on network side of Spine1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe19	Enter Interface mode for xe19
(config-if)#ip address 60.60.60.1/24	Configure IP address as 60.60.60.1 on network side of Spine2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

OSPF Configuration

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 3.3.3.3`	Configure router-ID as 3.3.3.3 (lo IP address)
`(config-router)#network 3.3.3.3/32 area 0.0.0.0`	Add 3.3.3.3 (lo IP address) network into area 0
`(config-router)#network 50.50.50.0/24 area 0.0.0.0`	Add 50.50.50.0 (Spine1) network into area 0
`(config-router)#network 60.60.60.0/24 area 0.0.0.0`	Add 60.60.60.0 (Spine2) network into area 0
`(config-router)#bfd all-interfaces`	Enabling BFD on all OSPF interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

BGP Configuration

`(config)#router bgp 500`	Enter into Router BGP mode
`(config-router)#bgp router-id 3.3.3.3`	Configure router-ID as 3.3.3.3 (lo ip address)
`(config-router)#neighbor 1.1.1.1 remote-as 500`	Specify a VTEP1 loopback IP address and remote-as defined
`(config-router)#neighbor 1.1.1.1 update-source lo`	Configure update as loopback for VTEP1
`(config-router)#neighbor 1.1.1.1 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP1
`(config-router)#neighbor 2.2.2.2 remote-as 500`	Specify a VTEP2 loopback IP address and remote-as defined
`(config-router)#neighbor 2.2.2.2 update-source lo`	Configure update as loopback for VTEP2
`(config-router)#neighbor 2.2.2.2 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP3
`(config-router)#address-family ipv4 unicast`	Enter into ipv4 unicast address family mode
`(config-router-af)#network 3.3.3.3/32`	Advertise loopback network into BGP for VTEP ID reachability
`(config-router-af)#neighbor 1.1.1.1 activate`	Activate 1.1.1.1 (VTEP2) into ipv4 unicast address family mode
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate 2.2.2.2 (VTEP2) into ipv4 unicast address family mode
`(config-router-af)#exit-address-family`	Exit from ipv4 unicast address family mode
`(config-router)#address-family l2vpn evpn`	Enter into L2VPN EVPN address family mode
`(config-router-af)#neighbor 1.1.1.1 activate`	Activate 1.1.1.1 (VTEP1) into L2VPN evpn address family mode
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate 2.2.2.2 (VTEP2) into L2VPN evpn address family mode
`(config-router-af)#exit-address-family`	Exit from L2VPN address family mode
`(config-router)#exit`	Exit from Router BGP mode and enter into config mode
`(config)#commit`	Commit the candidate configuration to the running configuration

VRF Configuration

`(config)# mac vrf VRF1`	Create MAC routing/forwarding instance with VRF1 name and enter into VRF mode
`(config-vrf)#rd 3.3.3.3:11`	Assign RD value
`(config-vrf)#route-target both 9.9.9.9:100`	Assign route-target value for same for import and export. Should be same on all node for VRF1
`(config-vrf)#exit`	Exit from VRF mode
`(config)#mac vrf VRF2`	Create MAC routing/forwarding instance with VRF2 name and enter into VRF mode
`(config-vrf)#rd 3.3.3.3:21`	Assign RD value
`(config-vrf)#route-target both 90.90.90.90:100`	Assign route-target value for same for import and export
`(config-vrf)#exit`	Exit from VRF
`(config)#commit`	Commit the candidate configuration to the running configuration

VxLAN Configuration

`(config)#nvo vxlan enable`	Enable VxLAN
`(config)#nvo vxlan vtep-ip-global 3.3.3.3`	Configure Source VTEP-IP-global configuration
`(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled`	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1`	Assign VRF for EVPN-BGP to carry EVPN route
`(config-nvo)#vni-name VNI-RED`	Configure VNI-name as VNI-RED
`(config-nvo)#exit`	Exit from VxLAN tenant mode and enter into configuration mode.
`(config)#nvo vxlan id 20 ingress-replication inner-vid-disabled`	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF2`	Assign VRF for EVPN-BGP to carry EVPN route
`(config-nvo)#vni-name VNI-BLUE`	Configure VNI-name as VNI-BLUE
`(config-nvo)#exit`	Exit from VxLAN tenant mode and enter into configuration mode.
`(config)#nvo vxlan access-if port-vlan xe48 1001`	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
`(config-nvo-acc-if)#map vni-name VNI-RED`	Map VxLAN Identified to access-port for VxLAN
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-interface mode and enter into configuration mode
`(config)#nvo vxlan access-if port-vlan xe48 1002`	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
`(config-nvo-acc-if)#map vni-name VNI-RED`	Map VxLAN Identified to access-port for VxLAN
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-interface mode and enter into configuration mode
`(config)#nvo vxlan access-if port-vlan xe48 3001`	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
`(config-nvo-acc-if)#map vni-name VNI-BLUE`	Map VxLAN Identified to access-port for VxLAN
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-interface mode and enter into configuration mode
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit from configuration mode

Switch (CE2)

Multihomed to 2-VTEPs (VTEP1 and VTEP2)

`#configure terminal`	Enter Configure mode.
`(config)#bridge 1 protocol ieee vlan-bridge`	Configure IEEE VLAN bridge
`(config)#vlan 1001-1002 bridge 1 state enable`	Configure VLANs from 1001-1002 and associate with bridge 1
`(config)#vlan 3001 bridge 1 state enable`	Configure VLANs from 3001 and associate with bridge 1
`(config)#interface xe22`	Enter Interface mode for xe22
`(config-if)#switchport`	Make xe22 as L2 port by configuring switchport

(config-if)#bridge-group 1	Associate xe22 to bridge 1
(config-if)#switchport mode hybrid	Configure xe22 as hybrid port
(config-if)#switchport hybrid allowed vlan add 1001-1002,3001 egress-tagged enable	Allow 1001-1002 and 3001 configured VLANs on xe22
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po1	Enter Interface mode for po1
(config-if)#switchport	Make po1 as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate po1 to bridge 1
(config-if)#switchport mode hybrid	Configure po1 as hybrid port
(config-if)#switchport hybrid allowed vlan add 1001-1002,3001 egress-tagged enable	Allow 1001-1002 and 3001 configured VLANs on po1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe3	Enter Interface mode for xe3
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#interface xe4	Enter Interface mode for xe4
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#exit	Exit from configuration mode
(config)#interface xe9	Enter Interface mode for xe9
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#interface xe10	Enter Interface mode for xe10
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#exit	Exit from configuration mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit from configuration mode

Spine 1

Spine node where all VTEPs are connected

Generic Configuration

`#configure terminal`	Enter Configure mode.
`(config)#qos enable`	Enabling QoS
`(config)#commit`	Commit the candidate configuration to the running configuration

Interface and Loopback Configuration

`#configure terminal`	Enter Configure mode.
`(config)#qos enable`	Enabling QoS
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 11.11.11.11/32 secondary`	Configure loopback IP address as 11.11.11.11 for Spine1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe8`	Enter Interface mode for xe8
`(config-if)#ip address 10.10.10.2/24`	Configure IP address as 10.10.10.2 on network side of VTEP1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe1`	Enter Interface mode for xe1
`(config-if)#ip address 30.30.30.2/24`	Configure IP address as 30.30.30.2 on network side of VTEP2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe18`	Enter Interface mode for xe18
`(config-if)#ip address 50.50.50.2/24`	Configure IP address as 50.50.50.2 on network side of VTEP3
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

OSPF configuration

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 11.11.11.11`	Configure router-ID as 11.11.11.11 (lo IP address)
`(config-router)#network 11.11.11.11/32 area 0.0.0.0`	Add 11.11.11.11 (lo IP address) network into area 0
`(config-router)#network 10.10.10.0/24 area 0.0.0.0`	Add 10.10.10.0 (VTEP1) network into area 0
`(config-router)#network 30.30.30.0/24 area 0.0.0.0`	Add 30.30.30.0 (VTEP2) network into area 0
`(config-router)#network 50.50.50.0/24 area 0.0.0.0`	Add 50.50.50.0 (VTEP3) network into area 0
`(config-router)#bfd all-interfaces`	Enabling BFD on all OSPF interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

Spine 2

Spine node where all VTEPs are connected

Generic configuration

`#configure terminal`	Enter Configure mode.
`(config)#qos enable`	Enabling QoS
`(config)#commit`	Commit the candidate configuration to the running configuration

Interface and loopback configuration

`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 22.22.22.22/32 secondary`	Configure loopback IP address as 22.22.22.22 for Spine2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe5`	Enter Interface mode for xe5
`(config-if)#ip address 20.20.20.2/24`	Configure IP address as 20.20.20.2 on network side of VTEP1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface ce0`	Enter Interface mode for ce0
`(config-if)#ip address 40.40.40.2/24`	Configure IP address as 40.40.40.2 on network side of VTEP2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe19`	Enter Interface mode for xe19
`(config-if)#ip address 60.60.60.2/24`	Configure IP address as 60.60.60.2 on network side of VTEP3
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

OSPF configuration

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 22.22.22.22`	Configure router-id as 11.11.11.11 (lo IP address)
`(config-router)#network 22.22.22.22/32 area 0.0.0.0`	Add 22.22.22.22 (lo IP address) network into area 0
`(config-router)#network 20.20.20.0/24 area 0.0.0.0`	Add 20.20.20.0 (VTEP1) network into area 0
`(config-router)#network 40.40.40.0/24 area 0.0.0.0`	Add 40.40.40.0 (VTEP2) network into area 0
`(config-router)#network 60.60.60.0/24 area 0.0.0.0`	Add 60.60.60.0 (VTEP3) network into area 0
`(config-router)#bfd all-interfaces`	Enabling BFD on all OSPF interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

Validation

VTEP1

```
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________
10       VNI-RED      L2       NW   ----      ----                           ---- ----
1.1.1.1          3.3.3.3
10       VNI-RED      L2       NW   ----      ----                           ---- ----
1.1.1.1          2.2.2.2
10       VNI-RED      --       AC   po1       00:00:00:00:00:11:11:00:00:00 1001 NON-DF
----             ----
10       VNI-RED      --       AC   po1       00:00:00:00:00:11:11:00:00:00 1002 DF
----             ----
20       VNI-BLUE     L2       NW   ----      ----                           ---- ----
1.1.1.1          3.3.3.3
20       VNI-BLUE     L2       NW   ----      ----                           ---- ----
1.1.1.1          2.2.2.2
20       VNI-BLUE     --       AC   po1       00:00:00:00:00:11:11:00:00:00 3001 NON-DF
----             ----

Total number of entries are 7

VTEP1#show nvo vxlan access-if brief

                 Inner                    Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
po1        1002  ---   500001   10        up      up
po1        1001  ---   500000   10        up      up
po1        3001  ---   500002   20        up      up

Total number of entries are 3

VTEP1#show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 500
BGP table version is 6
1 BGP AS-PATH entries
0 BGP community entries
```

```
Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
2.2.2.2                 4   500  161       163       5     0     0  01:05:15
6      3      0      2      1      0
3.3.3.3                 4   500  157       161       5     0     0  01:05:07
2      0      0      2      0      0

Total number of neighbors 2

Total number of Established sessions 2

VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
1.1.1.1          3.3.3.3          Installed        00:31:11       00:31:11
1.1.1.1          2.2.2.2          Installed        01:05:25       00:31:11

Total number of entries are 2

VTEP1#show bgp l2vpn evpn multihoming es-route

RD[1.1.1.1:1] VRF[evpn-gvrf-1]:
ESI                             PE IP-Address   Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   1.1.1.1         1.1.1.1         VXLAN
00:00:00:00:00:11:11:00:00:00   2.2.2.2         2.2.2.2         VXLAN

RD[2.2.2.2:1]
ESI                             PE IP-Address   Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   2.2.2.2         2.2.2.2         VXLAN

VTEP1#show bgp l2vpn evpn multihoming ethernet-ad-per-es

RD[1.1.1.1:1] VRF[evpn-gvrf-1]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              1.1.1.1         VXLAN

RD[1.1.1.1:11] VRF[VRF1]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              2.2.2.2         VXLAN

RD[1.1.1.1:21] VRF[VRF2]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              2.2.2.2         VXLAN

RD[2.2.2.2:1]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              2.2.2.2         VXLAN
VTEP1#show bgp l2vpn evpn multihoming ethernet-ad-per-evi

RD[1.1.1.1:11] VRF[VRF1]:
```

```
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   10          10             2.2.2.2         VXLAN
00:00:00:00:00:11:11:00:00:00   10          10             1.1.1.1         VXLAN

RD[1.1.1.1:21] VRF[VRF2]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   20          20             2.2.2.2         VXLAN
00:00:00:00:00:11:11:00:00:00   20          20             1.1.1.1         VXLAN

RD[2.2.2.2:11]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   10          10             2.2.2.2         VXLAN

RD[2.2.2.2:21]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   20          20             2.2.2.2         VXLAN

VTEP1#show bgp l2vpn evpn
BGP table version is 6, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1] VRF[evpn-gvrf-1]:
*>   [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                       1.1.1.1              0         100           32768    i  -------
---      VXLAN
*>   [4]:[00:00:00:00:00:11:11:00:00:00]:[32,1.1.1.1]
                       1.1.1.1              0         100           32768    i  -------
---      VXLAN
* i  [4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
                     2.2.2.2            0        100          0       i  2.2.2.2
VXLAN

RD[1.1.1.1:11] VRF[VRF1]:
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                     2.2.2.2            0        100          0       i  2.2.2.2
VXLAN
*>                     1.1.1.1              0         100           32768    i  -------
---      VXLAN
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
```

```
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>   [3]:[10]:[32,1.1.1.1]
                      1.1.1.1               0         100           32768    i  -------
---      VXLAN
* i  [3]:[10]:[32,2.2.2.2]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
* i  [3]:[10]:[32,3.3.3.3]
                    3.3.3.3             0        100          0       i  3.3.3.3
VXLAN

RD[1.1.1.1:21] VRF[VRF2]:
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[20]:[20]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>                     1.1.1.1               0         100           32768    i  -------
---      VXLAN
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>   [3]:[20]:[32,1.1.1.1]
                      1.1.1.1               0         100           32768    i  -------
---      VXLAN
* i  [3]:[20]:[32,2.2.2.2]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
* i  [3]:[20]:[32,3.3.3.3]
                    3.3.3.3             0        100          0       i  3.3.3.3
VXLAN

RD[2.2.2.2:1]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>i  [4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN

RD[2.2.2.2:11]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>i  [3]:[10]:[32,2.2.2.2]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN

RD[2.2.2.2:21]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[20]:[20]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>i  [3]:[20]:[32,2.2.2.2]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
```

```
RD[3.3.3.3:11]
*>i  [3]:[10]:[32,3.3.3.3]
                    3.3.3.3            0        100          0       i  3.3.3.3
VXLAN

RD[3.3.3.3:21]
*>i  [3]:[20]:[32,3.3.3.3]
                    3.3.3.3            0        100          0       i  3.3.3.3
VXLAN

Total number of prefixes 21
```

VTEP2

```
VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID    VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________
10      VNI-RED      L2       NW   ----      ----                           ---- ----
2.2.2.2          1.1.1.1
10      VNI-RED      L2       NW   ----      ----                           ---- ----
2.2.2.2          3.3.3.3
10      VNI-RED      --       AC   xe32      --- Single Homed Port ---      ---- ----
----             ----
10      VNI-RED      --       AC   po1       00:00:00:00:00:11:11:00:00:00 1001 DF
----             ----
10      VNI-RED      --       AC   po1       00:00:00:00:00:11:11:00:00:00 1002 NON-DF
----             ----
20      VNI-BLUE     L2       NW   ----      ----                           ---- ----
2.2.2.2          1.1.1.1
20      VNI-BLUE     L2       NW   ----      ----                           ---- ----
2.2.2.2          3.3.3.3
20      VNI-BLUE     --       AC   po1       00:00:00:00:00:11:11:00:00:00 3001 DF
----             ----

Total number of entries are 8

VTEP2#show nvo vxlan access-if
% Incomplete command.

VTEP2#show nvo vxlan access-if brief

                Inner                    Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
```

```
xe32        ---   ---   500004   10         up       up
po1         1002  ---   500001   10         up       up
po1         1001  ---   500000   10         up       up
po1         3001  ---   500002   20         up       up

Total number of entries are 4

VTEP2#show bgp l2vpn evpn summary
BGP router identifier 2.2.2.2, local AS number 500
BGP table version is 4
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4   500  172       171      4      0      0  01:09:28
6      3      0      2      1      0
3.3.3.3                 4   500  165       173      4      0      0  01:09:29
2      0      0      2      0      0

Total number of neighbors 2

Total number of Established sessions 2

VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
2.2.2.2          1.1.1.1          Installed        01:09:38       00:35:24
2.2.2.2          3.3.3.3          Installed        01:09:39       01:09:39

Total number of entries are 2

VTEP2#show bgp l2vpn evpn multihoming es-route

RD[1.1.1.1:1]
ESI                             PE IP-Address   Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   1.1.1.1         1.1.1.1         VXLAN

RD[2.2.2.2:1] VRF[evpn-gvrf-1]:
ESI                             PE IP-Address   Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   1.1.1.1         1.1.1.1         VXLAN
00:00:00:00:00:11:11:00:00:00   2.2.2.2         2.2.2.2         VXLAN
VTEP2#show bgp l2vpn evpn multihoming ethernet-ad-per-es

RD[1.1.1.1:1]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              1.1.1.1         VXLAN

RD[2.2.2.2:1] VRF[evpn-gvrf-1]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
```

```
00:00:00:00:00:11:11:00:00:00  4294967295  0              2.2.2.2         VXLAN

RD[2.2.2.2:11] VRF[VRF1]:
ESI                            Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00  4294967295  0              1.1.1.1         VXLAN

RD[2.2.2.2:21] VRF[VRF2]:
ESI                            Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00  4294967295  0              1.1.1.1         VXLAN

VTEP2#show bgp l2vpn evpn multihoming ethernet-ad-per-evi

RD[1.1.1.1:11]
ESI                            Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00  10          10             1.1.1.1         VXLAN

RD[1.1.1.1:21]
ESI                            Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00  20          20             1.1.1.1         VXLAN

RD[2.2.2.2:11] VRF[VRF1]:
ESI                            Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00  10          10             2.2.2.2         VXLAN
00:00:00:00:00:11:11:00:00:00  10          10             1.1.1.1         VXLAN

RD[2.2.2.2:21] VRF[VRF2]:
ESI                            Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00  20          20             2.2.2.2         VXLAN
00:00:00:00:00:11:11:00:00:00  20          20             1.1.1.1         VXLAN

VTEP2# show bgp l2vpn evpn
BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
*>i  [4]:[00:00:00:00:00:11:11:00:00:00]:[32,1.1.1.1]
```

```
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN

RD[1.1.1.1:11]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
*>i  [3]:[10]:[32,1.1.1.1]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN

RD[1.1.1.1:21]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[20]:[20]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
*>i  [3]:[20]:[32,1.1.1.1]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN

RD[2.2.2.2:1] VRF[evpn-gvrf-1]:
*>   [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                      2.2.2.2             0         100           32768    i  -------
---      VXLAN
* i  [4]:[00:00:00:00:00:11:11:00:00:00]:[32,1.1.1.1]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
*>   [4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
                      2.2.2.2             0         100           32768    i  -------
---      VXLAN

RD[2.2.2.2:11] VRF[VRF1]:
*>   [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                      2.2.2.2             0         100           32768    i  -------
---      VXLAN
* i                  1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
* i  [3]:[10]:[32,1.1.1.1]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
*>   [3]:[10]:[32,2.2.2.2]
                      2.2.2.2             0         100           32768    i  -------
---      VXLAN
* i  [3]:[10]:[32,3.3.3.3]
                    3.3.3.3             0        100          0       i  3.3.3.3
VXLAN

RD[2.2.2.2:21] VRF[VRF2]:
*>   [1]:[00:00:00:00:00:11:11:00:00:00]:[20]:[20]
                      2.2.2.2             0         100           32768    i  -------
---      VXLAN
```

```
* i                  1.1.1.1            0        100          0       i  1.1.1.1
VXLAN
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                     1.1.1.1            0        100          0       i  1.1.1.1
VXLAN
* i  [3]:[20]:[32,1.1.1.1]
                     1.1.1.1            0        100          0       i  1.1.1.1
VXLAN
*>   [3]:[20]:[32,2.2.2.2]
                       2.2.2.2              0         100           32768    i  -------
---      VXLAN
* i  [3]:[20]:[32,3.3.3.3]
                     3.3.3.3            0        100          0       i  3.3.3.3
VXLAN

RD[3.3.3.3:11]
*>i  [3]:[10]:[32,3.3.3.3]
                     3.3.3.3            0        100          0       i  3.3.3.3
VXLAN

RD[3.3.3.3:21]
*>i  [3]:[20]:[32,3.3.3.3]
                     3.3.3.3            0        100          0       i  3.3.3.3
VXLAN

Total number of prefixes 21
```

VTEP3

```
VTEP3#  show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
10       VNI-RED      L2       NW   ----      ----                         ---- ----
3.3.3.3          2.2.2.2
10       VNI-RED      L2       NW   ----      ----                         ---- ----
3.3.3.3          1.1.1.1
10       VNI-RED      --       AC   xe48      --- Single Homed Port ---    1001 ----
----             ----
10       VNI-RED      --       AC   xe48      --- Single Homed Port ---    1002 ----
----             ----
20       VNI-BLUE     L2       NW   ----      ----                         ---- ----
3.3.3.3          2.2.2.2
20       VNI-BLUE     L2       NW   ----      ----                         ---- ----
3.3.3.3          1.1.1.1
20       VNI-BLUE     --       AC   xe48      --- Single Homed Port ---    3001 ----
----             ----
```

```
Total number of entries are 7

VTEP3#show nvo vxlan access-if brief

                 Inner                    Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
xe48       1002  ---   500001   10        up      up
xe48       1001  ---   500000   10        up      up
xe48       3001  ---   500002   20        up      up

Total number of entries are 3

VTEP3#show bgp l2vpn evpn summary
BGP router identifier 3.3.3.3, local AS number 500
BGP table version is 4
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4   500  177       173      3      0      0  01:11:49
6      3      0      2      1      0
2.2.2.2                 4   500  177       171      2      0      0  01:11:59
6      3      0      2      1      0

Total number of neighbors 2

Total number of Established sessions 2

VTEP3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
=============================================================================
3.3.3.3          2.2.2.2          Installed        01:12:10       01:12:10
3.3.3.3          1.1.1.1          Installed        01:12:00       01:12:00

Total number of entries are 2

VTEP3#show bgp l2vpn evpn multihoming es-route

RD[1.1.1.1:1]
ESI                           PE IP-Address   Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00 1.1.1.1         1.1.1.1         VXLAN

RD[2.2.2.2:1]
ESI                           PE IP-Address   Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00 2.2.2.2         2.2.2.2         VXLAN

VTEP3#show bgp l2vpn evpn multihoming ethernet-ad-per-es
```

```
RD[1.1.1.1:1]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              1.1.1.1         VXLAN

RD[2.2.2.2:1]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              2.2.2.2         VXLAN

RD[3.3.3.3:11] VRF[VRF1]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              2.2.2.2         VXLAN
00:00:00:00:00:11:11:00:00:00   4294967295  0              1.1.1.1         VXLAN

RD[3.3.3.3:21] VRF[VRF2]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   4294967295  0              2.2.2.2         VXLAN
00:00:00:00:00:11:11:00:00:00   4294967295  0              1.1.1.1         VXLAN

VTEP3#show bgp l2vpn evpn multihoming ethernet-ad-per-evi

RD[1.1.1.1:11]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   10          10             1.1.1.1         VXLAN

RD[1.1.1.1:21]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   20          20             1.1.1.1         VXLAN

RD[2.2.2.2:11]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   10          10             2.2.2.2         VXLAN

RD[2.2.2.2:21]
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   20          20             2.2.2.2         VXLAN

RD[3.3.3.3:11] VRF[VRF1]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   10          10             2.2.2.2         VXLAN
00:00:00:00:00:11:11:00:00:00   10          10             1.1.1.1         VXLAN

RD[3.3.3.3:21] VRF[VRF2]:
ESI                             Eth-Tag     VNID/LABEL     Nexthop IP      Encap
00:00:00:00:00:11:11:00:00:00   20          20             2.2.2.2         VXLAN
00:00:00:00:00:11:11:00:00:00   20          20             1.1.1.1         VXLAN

VTEP3#show bgp l2vpn evpn
BGP table version is 4, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network          Next Hop            Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                    1.1.1.1             0         100          0        i  1.1.1.1
VXLAN
*>i  [4]:[00:00:00:00:00:11:11:00:00:00]:[32,1.1.1.1]
                    1.1.1.1             0         100          0        i  1.1.1.1
VXLAN

RD[1.1.1.1:11]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                    1.1.1.1             0         100          0        i  1.1.1.1
VXLAN
*>i  [3]:[10]:[32,1.1.1.1]
                    1.1.1.1             0         100          0        i  1.1.1.1
VXLAN

RD[1.1.1.1:21]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[20]:[20]
                    1.1.1.1             0         100          0        i  1.1.1.1
VXLAN
*>i  [3]:[20]:[32,1.1.1.1]
                    1.1.1.1             0         100          0        i  1.1.1.1
VXLAN

RD[2.2.2.2:1]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                    2.2.2.2             0         100          0        i  2.2.2.2
VXLAN
*>i  [4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
                    2.2.2.2             0         100          0        i  2.2.2.2
VXLAN

RD[2.2.2.2:11]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                    2.2.2.2             0         100          0        i  2.2.2.2
VXLAN
*>i  [3]:[10]:[32,2.2.2.2]
                    2.2.2.2             0         100          0        i  2.2.2.2
VXLAN
```

```
RD[2.2.2.2:21]
*>i  [1]:[00:00:00:00:00:11:11:00:00:00]:[20]:[20]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>i  [3]:[20]:[32,2.2.2.2]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN

RD[3.3.3.3:11] VRF[VRF1]:
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
* i                  1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
* i                  1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
* i  [3]:[10]:[32,1.1.1.1]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
* i  [3]:[10]:[32,2.2.2.2]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>   [3]:[10]:[32,3.3.3.3]
                      3.3.3.3               0         100           32768    i  -------
---      VXLAN

RD[3.3.3.3:21] VRF[VRF2]:
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[20]:[20]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
* i                  1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
* i  [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
* i                  1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
* i  [3]:[20]:[32,1.1.1.1]
                    1.1.1.1             0        100          0       i  1.1.1.1
VXLAN
* i  [3]:[20]:[32,2.2.2.2]
                    2.2.2.2             0        100          0       i  2.2.2.2
VXLAN
*>   [3]:[20]:[32,3.3.3.3]
                      3.3.3.3               0         100           32768    i  -------
---      VXLAN

Total number of prefixes 22
```

Static MAC-IP advertise through Single Home and Multihomed VTEPs

Advertise static MAC IPv4 from MH1 and SH3.

MH1-VTEPs: VTEP1 & VTEP2- same MAC should be configured on both VTEPs under po access-port, configs should be symmetric between MH VTEPs

SH3-VTEP: VTEP3

VTEP1(MH1)

`#configure terminal`	Enter Configure mode.
`(config)# nvo vxlan access-if port-vlan po1 1001`	Enter into VxLAN MH po1 access-port with VLAN 1001
`(config-nvo-acc-if)#mac 0000.1111.1001 ip 11.11.10.1`	Configure static MAC IP
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-port config mode
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit from configuration mode

VTEP2(MH1)

`#configure terminal`	Enter Configure mode.
`(config)#nvo vxlan access-if port-vlan po1 1001`	Enter into VxLAN MH po1 access-port with vlan 1001
`(config-nvo-acc-if)# mac 0000.1111.1001 ip 11.11.10.1`	Configure static MAC IP
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-port config mode
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit from configuration mode

VTEP3(SH)

`#configure terminal`	Enter Configure mode.
`(config)# nvo vxlan access-if port-vlan xe48 1001`	Enter into single-homed access-port - xe48 with VLAN 1001
`(config-nvo-acc-if)#mac 0000.3333.1001 ip 11.11.10.2`	Configure static MAC IP
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-port config mode
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit from configuration mode

Validation

Verify MAC-table in MH VTEPs and Single Home VTEP, MAC will be advertised through ESI value which is advertised from VTEP1 and VTEP2 and VTEP IP from SH VTEP VTEP3.

Verify ARP-cache table in all VTEPs, VTEP1 and VTEP2 will learn VTEP3 IP.

Any ARP request comes for 11.11.10.2, VTEP1/VTEP2 will do proxy-ARP.

VTEP1

```
VTEP1#show nvo vxlan mac-table
=========================================================================================
=====================================
                                                     VXLAN MAC Entries
=========================================================================================
=====================================
VNID       Interface VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type            Status     AccessPortDesc
_________________________________________________________________________________________
___________________________________________________


10         po1       1001   ----         0000.1111.1001 00:00:00:00:00:11:11:00:00:00
Static Local    -------    -------
10         ----      ----   ----         0000.3333.1001 3.3.3.3
Static Remote   -------    -------

Total number of entries are : 2

VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________
10       11.11.10.1       0000.1111.1001 Static Local   ----
10       11.11.10.2       0000.3333.1001 Static Remote  ----
Total number of entries are 2
```

VTEP2

```
VTEP2#show nvo vxlan mac-table
=========================================================================================
=====================================
                                                     VXLAN MAC Entries
=========================================================================================
=====================================
VNID       Interface VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type            Status     AccessPortDesc
_________________________________________________________________________________________
_____________________________________


10         po1       1001   ----         0000.1111.1001 00:00:00:00:00:11:11:00:00:00
Static Local    -------    -------
```

```
10          ----      ----   ----         0000.3333.1001 3.3.3.3
Static Remote   -------    -------

Total number of entries are : 2

VTEP2#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
10       11.11.10.1       0000.1111.1001 Static Local   ----
10       11.11.10.2       0000.3333.1001 Static Remote  ----
Total number of entries are 2
```

VTEP3

```
VTEP3#show nvo vxlan mac-table
================================================================================================
====================================
                                                          VXLAN MAC Entries
================================================================================================
====================================
VNID       Interface VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type            Status     AccessPortDesc
_____________________________________________________________________________________________
_____________________________________

10          ----      1001   ----         0000.1111.1001 00:00:00:00:00:11:11:00:00:00
Static Remote   -------    -------
10          xe48      1001   ----         0000.3333.1001 3.3.3.3
Static Local    -------    -------

Total number of entries are : 2

VTEP3#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
10       11.11.10.1       0000.1111.1001 Static Remote  ----
10       11.11.10.2       0000.3333.1001 Static Local   ----
Total number of entries are 2
```

Dynamic MAC advertise through Single Home and Multihomed VTEPs

Advertise 2 MAC's through CE1 connected IXIA, dynamic MAC entries and verify MAC-table in all VTEPs.

One MAC will be dynamic local in VTEP1 and same will be remote in VTEP2 and other be dynamic local in VTEP2 and same will be remote in VTEP1.

Both MAC's will be in remote in VTEP3.

VTEP1

```
VTEP1#show nvo vxlan mac-table
================================================================================================
====================================
                                                                VXLAN MAC Entries
================================================================================================
====================================
VNID       Interface VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type            Status     AccessPortDesc

________________________________________________________________________________________________
_____________________________________

10         po1       1001   ----         0000.1111.1002 00:00:00:00:00:11:11:00:00:00
Dynamic Local   -------    -------
10         ----      1002   ----         0000.1111.1003 00:00:00:00:00:11:11:00:00:00
Dynamic Remote  -------    -------

Total number of entries are : 2

VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left

_____________________________________________________________________________
10       21.21.21.1       0000.1111.1002 Dynamic Local  ----
10       31.1.31.1        0000.1111.1003 Dynamic Remote ----
Total number of entries are 2
```

VTEP2

```
VTEP2#show nvo vxlan mac-table
================================================================================================
====================================
                                                                VXLAN MAC Entries
================================================================================================
====================================
VNID       Interface VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type            Status     AccessPortDesc

________________________________________________________________________________________________
_____________________________________

10         ----      1001   ----         0000.1111.1002 00:00:00:00:00:11:11:00:00:00
Dynamic Remote  -------    -------
10         po1       1002   ----         0000.1111.1003 00:00:00:00:00:11:11:00:00:00
Dynamic Local   -------    -------

Total number of entries are : 2

VTEP2#show nvo vxla arp-cache
VXLAN ARP-CACHE Information
============================
```

```
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
10       21.21.21.1       0000.1111.1002 Dynamic Remote ----
10       31.1.31.1        0000.1111.1003 Dynamic Local  ----
Total number of entries are 2
```

VTEP3

```
VTEP3#show nvo vxlan mac-table
================================================================================================
================================================
                                                          VXLAN MAC Entries
================================================================================================
================================================
VNID       Interface VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type            Status     AccessPortDesc
________________________________________________________________________________________________
_____________________________________________________

10         ----      1001   ----         0000.1111.1002 00:00:00:00:00:11:11:00:00:00
Dynamic Remote  -------    -------
10         ----      1002   ----         0000.1111.1003 00:00:00:00:00:11:11:00:00:00
Dynamic Remote  -------    -------

Total number of entries are : 2

VTEP3#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
10       21.21.21.1       0000.1111.1002 Dynamic Remote ----
10       31.1.31.1        0000.1111.1003 Dynamic Remote ----
Total number of entries are 2
```

Note: When VTEP1 tunnel goes down, then traffic from VTEP3 will use VTEP2 for forwarding. But Traffic from Switch to VTEP1 will be lost in VTEP1 itself.

When DUT is rebooted, access-if will be in hold down state until ESI hold timer value expiry. After ESI hold timer expiry, access-if port will be up and started learning.

MAC Hold timer will not be applicable on ESI interface, because of mass-withdraw requirement.

A CE can connect to maximum two nodes for multihoming, more than two nodes in a multihoming group is not supported.

All configuration (shutdown, disable learning, disable arp/nd cache, disable arp/nd flood, map vnid, qos profiles, encapsulation) on a multihomed access port should be same on both VTEPs sharing the ESI for multihoming functionalities to work properly.

Multiple ESI values are supported on same VTEP.

CHAPTER 3 VXLAN Hybrid Access Port Configuration

This chapter shows how to configure a hybrid access port which is a Layer 2 Port (configured switchport) that is part of both a VXLAN domain and a Layer 2 bridge with different VLANs.

Overview

A hybrid VXLAN access port is a Layer 2 port that is part of regular Layer 2 bridge (RSTP/MSTP/STP) and a VXLAN bridge. The mapping between Layer 2 bridging and VxLAN untagged and tagged access interface is supported on the same Layer 2 switch port interface. The same VLAN cannot be a part of both a VXLAN domain and a Layer 2 bridge.

If a port is created with “all” VLANS, then the port should not allow VXLAN access-port configurations. If a VXLAN with a specific VLAN is mapped, then configuring VLAN “all” on the same port does not allow that specified VLAN in a Layer 2 bridge. If the VXLAN access-port configurations are removed, then the specified VLAN is added immediately in a Layer 2 bridge.

Ingress traffic with a VXLAN VLAN does not receive Layer 2 traffic or vice-versa. STP states on the Port P1 do not affect VXLAN traffic.

Topology

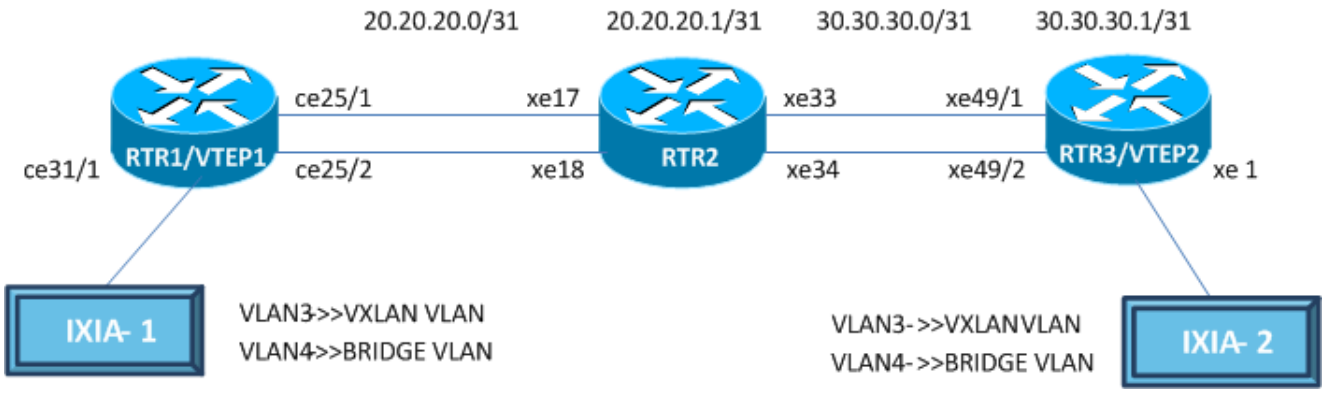

RTR1/VTEP1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode for loopback.
`(config-if)#ip address 10.10.10.10/32 secondary`	Assign secondary IP address.
`(config-if)#exit`	Exit interface mode.
`(config)#mac vrf vrf12345`	Create mac routing/forwarding instance with vrf12345 name and enter VRF mode.
`(config-vrf)#rd 1.1.1.1:1`	Assign Rd value.
`(config-vrf)#route-target both 10.10.10.10:10`	Assign route-target both value.
`(config-vrf)#exit`	Exit VRF configuration mode.
`(config)#bridge 32 protocol ieee vlan-bridge`	Configure the ieee vlan-bridge with Id 32.

(config)#vlan 2-5 bridge 32	Configure the vlans 2-5 for the configured bridge Id 32.
(config)#interface ce25/1	Enter interface mode for ce25/1.
(config-if)#ip address 20.20.20.0/31	Assign IP address 20.20.20.0 in /31 mask.
(config-if)#exit	Exit interface mode.
(config)#interface ce25/2	Enter interface mode for ce25/2.
(config-if)#switchport	Make it L2 interface.
(config-if)#bridge-group 32	Associate the bridge-group 32 to the interface.
(config-if)#switchport mode hybrid	Configure the Hybrid mode.
(config-if)# switchport hybrid allowed vlan add 4 egress-tagged enable	Configure hybrid allowed vlan add 4 to support the created vlan in the L2 Bridge.
(config-if)#exit	Exit interface mode.
(config)#interface ce31/1	Enter interface mode for ce31/1.
(config-if)#switchport	Make it L2 interface.
(config-if)#bridge-group 32	Associate the bridge-group 32 to the interface.
(config-if)#switchport mode hybrid	Configure the Hybrid mode.
(config-if)# switchport hybrid allowed vlan add 4 egress-tagged enable	Configure hybrid allowed vlan add 4 to support the created vlan in the L2 Bridge.
(config-if)#exit	Exit interface mode.
(config)#router bgp 64512	Enter BGP router mode.
(config-router)# bgp router-id 1.1.1.1	Assign BGP router ID
(config-router)#neighbor 20.20.20.1 remote-as 64513	Specify a neighbor router with peer ip address and remote-as defined.
(config-router)#address-family l2vpn evpn	Enter l2vpn address family mode.
(config-router-af)#neighbor 20.20.20.1 activate	Activate the peer into address family mode.
(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family
(config-router-af)#network 10.10.10.10/32	Advertise loopback network into BGP for VTEP ID reachability
(config-router-af)#exit-address-family	Exit ipv4 unicast address family mode
(config-router-af)#exit-address-family	Exit l2vpn address family mode.
(config-router)#exit	Exit BGP router mode.
(config)#nvo vxlan enable	Enable Vxlan.
(config)#nvo vxlan vtep-ip-global 10.10.10.10	Configure the source Vtep-ip.
(config)#nvo vxlan id 16777215 ingress-replication inner-vid-disabled	Configure Vxlan Network identifier with/without inner-vid-disabled configure and enter Vxlan tenant mode.
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf12345	Assign VRF for evpn-bgp to carry evpn route.
(config-nvo)#exit	Exit Vxlan tenant mode.
(config)#nvo vxlan access-if port-vlan ce31/1 3	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) mapping.
(config-nvo-acc-if)#map vnid 16777215	Map Vxlan Identifier to access-port.
(config-nvo-acc-if)#exit	Exit Vxlan access-interface mode.
(config)#commit	Commit the candidate configuration to the running configuration

RTR2

`#configure terminal`	Enter configure mode.
`(config)#bridge 32 protocol ieee vlan-bridge`	Configure the ieee vlan-bridge with Id 32.
`(config)#vlan 2-5 bridge 32`	Configure the vlans 2-5 for the configured bridge Id 32.
`(config)#interface xe17`	Enter interface mode for xe17.
`(config-if)#ip address 20.20.20.1/31`	Assign IP address 20.20.20.1 in /31 mask.
`(config-if)#exit`	Exit interface mode.
`(config)#interface  xe18`	Enter interface mode for xe18.
`(config-if)#switchport`	Make it L2 interface .
`(config-if)#bridge-group 32`	Associate the bridge-group 32 to the interface.
`(config-if)#switchport mode hybrid`	Configure the Hybrid mode.
`(config-if)#switchport hybrid allowed vlan add 4 egress-tagged enable`	Configure hybrid allowed vlan add 4 to support the created vlan in the L2 Bridge.
`(config-if)#exit`	Exit interface mode.
`(config)#interface  xe33`	Enter interface mode for xe33.
`(config-if)#ip address 30.30.30.0/31`	Assign IP address 30.30.30.0 in /31 mask.
`(config-if)#exit`	Exit interface mode.
`(config)#interface  xe34`	Enter interface mode for xe34.
`(config-if)#switchport`	Make it L2 interface.
`(config-if)#bridge-group 32`	Associate the bridge-group 32 to the interface.
`(config-if)#switchport mode hybrid`	Configure the Hybrid mode.
`(config-if)#switchport hybrid allowed vlan add 4 egress-tagged enable`	Configure hybrid allowed vlan add 4 to support the created vlan in the L2 Bridge.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 64513`	Enter BGP router mode.
`(config-router)#bgp router-id 2.2.2.2`	Assign BGP router ID
`(config-router)#neighbor 20.20.20.0 remote-as 64512`	Specify a neighbor router with peer ip address and remote-as defined.
`(config-router)#neighbor 30.30.30.1 remote-as 64514`	Specify a neighbor router with peer ip address and remote-as defined.
`(config-router)#address-family l2vpn evpn`	Enter l2vpn address family mode.
`(config-router-af)#neighbor 20.20.20.0 activate`	Activate the peer into address family mode.
`(config-router-af)#neighbor 30.30.30.1 activate`	Activate the peer into address family mode.
`(config-router-af)#exit-address-family`	Exit l2vpn address family mode.
`(config-router)#exit`	Exit BGP router mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

RTR3/VTEP2

#configure terminal	Enter configure mode.
(config)#interface lo	Enter interface mode for loopback.
(config-if)#ip address 40.40.40.40/32 secondary	Assign secondary IP address.
(config-if)#exit	Exit interface mode.
(config)#mac vrf vrf12345	Create mac routing/forwarding instance with vrf12345 name and enter VRF mode.
(config-vrf)#rd 2.2.2.2:2	Assign Rd value.
(config-vrf)#route-target both 10.10.10.10:10	Assign route-target both value.
(config-vrf)#exit	Exit VRF configuration mode.
(config)#bridge 32 protocol ieee vlan-bridge	Configure the ieee vlan-bridge with Id 32.
(config)#vlan 2-5 bridge 32	Configure the vlans 2-5 for the configured bridge Id 32.
(config)#interface xe 49/1	Enter interface mode for xe49/1.
(config-if)#ip address 30.30.30.1/31	Assign Ip address 30.30.30.1 in /31 mask.
(config-if)#exit	Exit interface mode.
(config)#interface xe49/2	Enter interface mode for xe49/2.
(config-if)#switchport	Make it L2 interface .
(config-if)#bridge-group 32	Associate the bridge-group 32 to the interface.
(config-if)#switchport mode hybrid	Configure the Hybrid mode.
(config-if)#switchport hybrid allowed vlan add 4 egress-tagged enable	Configure hybrid allowed vlan add 4 to support the created vlan in the L2 Bridge.
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#switchport	Make it L2 interface .
(config-if)#bridge-group 32	Associate the bridge-group 32 to the interface.
(config-if)#switchport mode hybrid	Configure the Hybrid mode.
(config-if)#switchport hybrid allowed vlan add 4 egress-tagged enable	Configure hybrid allowed vlan add 4 to support the created vlan in the L2 Bridge.
(config-if)#exit	Exit interface mode.
(config)#router bgp 64514	Enter BGP router mode.
(config-router)# bgp router-id 3.3.3.3	Assign BGP router ID
(config-router)#neighbor 30.30.30.0 remote-as 64513	Specify a neighbor router with peer IP address and remote-as defined.
(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family
(config-router-af)#network 40.40.40.40/32	Advertise loopback network into BGP for VTEP ID reachability
(config-router-af)#exit-address-family	Exit ipv4 unicast address family mode
(config-router)#address-family l2vpn evpn	Enter l2vpn address family mode.
(config-router-af)#neighbor 30.30.30.0 activate	Activate the peer into address family mode.
(config-router-af)#exit-address-family	Exit l2vpn address family mode.
(config-router)#exit	Exit BGP router mode.

(config)#nvo vxlan enable	Enable Vxlan.
(config)#nvo vxlan vtep-ip-global 40.40.40.40	Configure the source Vtep-ip.
(config)#nvo vxlan id 16777215 ingress-replication inner-vid-disabled	Configure Vxlan Network identifier with/without inner-vid-disabled configure and enter Vxlan tenant mode.
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf12345	Assign VRF for evpn-bgp to carry evpn route.
(config-nvo)#exit	Exit Vxlan tenant mode.
(config)#nvo vxlan access-if port-vlan xe1 3	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) mapping.
(config-nvo-acc-if)#map vnid 16777215	Map Vxlan Identifier to access-port.
(config-nvo-acc-if)#exit	Exit Vxlan access-interface mode.
(config)#commit	Commit the candidate configuration to the running configuration

Validation

VTEP1

```
#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 10.10.10.10
!
nvo vxlan id 16777215 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrf12345
!
nvo vxlan access-if port-vlan ce31/1 3
 map vnid 16777215
!

VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
           AC - Access Port
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface     ESI                             VLAN DF-
Status Src-Addr         Dst-Addr

_____________________________________________________________________________________
_____________________________________
16777215 ----         L2       NW   ----          ------                          ---- --
--      10.10.10.10      40.40.40.40
16777215 ----         --       AC   ce31/1         --- Single Homed Port ---     3    -
---      ----            ----
 Total number of entries are 2
```

```
VTEP1#show nvo vxlan vnid 16777215
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface     ESI                             VLAN DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
16777215 ----         L2       NW   ----          ------                          ---- --
--      10.10.10.10      40.40.40.40
16777215 ----         --       AC   ce31/1         --- Single Homed Port ---     3    -
---      ----             ----
 Total number of entries are 2!

VTEP1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     10.10.10.10/32 is directly connected, lo, 01:15:55
C     20.20.20.0/31 is directly connected, xe10/1, 01:07:53
B     40.40.40.40/32 [20/0] via 20.20.20.1, xe10/1, 00:42:54
C     127.0.0.0/8 is directly connected, lo, 1d05h02m

Gateway of last resort is not set


VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
=============================================================================
33.33.33.0       34.34.34.0       Installed     00:26:27     00:26:27
Total number of entries are 1


VTEP1#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.2, local AS number 64512
BGP table version is 10
2 BGP AS-PATH entries
0 BGP community entries
```

```
Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
20.20.20.1              4 64513  108       109     10      0      0  00:48:14
3      0      2      1      0      0

Total number of neighbors 1

Total number of Established sessions 1

VTEP1#show bgp l2vpn evpn
BGP table version is 4, local router ID is 10.10.10.10
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1] VRF[vrf12345]:
*    [2]:[0]:[16777215]:[48,0000.0744.4433]:[0]:[16777215]
                       40.40.40.40          0         100           0       64513 64514
i  20.20.20.1      VXLAN
*>   [2]:[0]:[16777215]:[48,0000.2222.2222]:[0]:[16777215]
                       10.10.10.10          0         100           32768    i  --------
--      VXLAN
*>   [3]:[16777215]:[32,10.10.10.10]
                       10.10.10.10          0         100           32768    i  --------
--      VXLAN
*    [3]:[16777215]:[32,40.40.40.40]
                       40.40.40.40          0         100           0       64513 64514
i  20.20.20.1      VXLAN

RD[2.2.2.2:2]
*>   [2]:[0]:[16777215]:[48,0000.0744.4433]:[0]:[16777215]
                       40.40.40.40          0         100           0       64513 64514
i  20.20.20.1      VXLAN
*>   [3]:[16777215]:[32,40.40.40.40]
                       40.40.40.40          0         100           0       64513 64514
i  20.20.20.1      VXLAN

Total number of prefixes 6

VTEP1#show nvo vxlan mac-table
================================================================================
============================================================
                              VXLAN MAC Entries
```

```
================================================================================================
============================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc

________________________________________________________________________________________________
________________________________________________________________

16777215  ce31/1      3      ----         0000.2222.2222 10.10.10.10
Dynamic Local        -------              -------
16777215  ----       ----   ----         0000.0744.4433 40.40.40.40
Dynamic Remote       -------              -------

Total number of entries are : 2

VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left

________________________________________________________________________________
Total number of entries are 0

VTEP1#show vlan brief
Bridge  VLAN ID     Name         State   H/W Status      Member ports
                                                         (u)-Untagged, (t)-Tagged
======= ======= ================ ======= ========== ==========================
32      1       default          ACTIVE  Success    ce25/2(u) ce31/1(u)
32      2       VLAN0002         ACTIVE  Success
32      3       VLAN0003         ACTIVE  Success
32      4       VLAN0004         ACTIVE  Success    ce25/2(t) ce31/1(t)
32      5       VLAN0005         ACTIVE  Success
```

RTR2

```
RTR2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B     10.10.10.10/32 [20/0] via 20.20.20.0,  xe17, 00:29:48
C     20.20.20.0/31 is directly connected,  xe17, 02:33:29
C     30.30.30.0/31 is directly connected, xe33, 02:31:56
B     40.40.40.40/32 [20/0] via 30.30.30.1, xe33, 02:23:26
C     127.0.0.0/8 is directly connected, lo, 21:17:41

Gateway of last resort is not set
```

```
RTR2#show bgp l2vpn evpn summary
BGP router identifier 11.11.11.1, local AS number 64513
BGP table version is 10
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
20.20.20.0              4 64512  111       112     10      0      0  00:49:36
3      0      2      1      0      0
30.30.30.1              4 64514  101       103     10      0      0  00:45:10
3      0      2      1      0      0

Total number of neighbors 2

Total number of Established sessions 2


#show nvo vxlan vnid 16777215
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface     ESI                             VLAN DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________________
16777215 ----         L2       NW   ----          ------                           ---- --
--      40.40.40.40      10.10.10.10
16777215 ----         --       AC   xe1        --- Single Homed Port ---     3    ----
----             ----
 Total number of entries are 2


RTR2#show vlan brief
Bridge  VLAN ID     Name         State   H/W Status      Member ports
                                                        (u)-Untagged, (t)-Tagged
======= ======= ================ ======= ========== ==========================
32      1       default          ACTIVE  Success    ce10/2(u) ce11/2(u)
32      2       VLAN0002         ACTIVE  Success
32      3       VLAN0003         ACTIVE  Success
32      4       VLAN0004         ACTIVE  Success    ce10/2(t) ce11/2(t)
32      5       VLAN0005         ACTIVE  Success
```

VTEP2

```
#show running-config nvo vxlan
!
```

```
nvo vxlan enable
!
nvo vxlan vtep-ip-global 40.40.40.40
!
nvo vxlan id 16777215 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrf12345
!
nvo vxlan access-if port-vlan xe1 3
 no shutdown
 map vnid 16777215
!

VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name   VNI-Type Type Interface    ESI  VLAN DF-Status Src-Addr   Dst-Addr
_____________________________________________
16777215 ----         L2       NW   ----          ------                         ---- --
--      40.40.40.40      10.10.10.10
16777215 ----        --      AC   xe1        --- Single Homed Port ---     3    ----
----             ----
 Total number of entries are 2

VTEP2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B     10.10.10.10/32 [20/0] via 30.30.30.0, xe11/1, 00:27:32
C     30.30.30.0/31 is directly connected, xe11/1, 00:30:00
C     40.40.40.40/32 is directly connected, lo, 00:31:00
C     127.0.0.0/8 is directly connected, lo, 06:25:00

Gateway of last resort is not set


VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
```

```
40.40.40.40      10.10.10.10      Installed     00:14:40     00:14:40
Total number of entries are 1


VTEP2#show bgp l2vpn evpn summary
BGP router identifier 33.33.33.33, local AS number 64514
BGP table version is 8
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
30.30.30.0              4 64513  267       270      8     0     0  02:09:07
2      0      1      1      0      0

Total number of neighbors 1

Total number of Established sessions 1

VTEP2#show nvo vxlan access-if-config
nvo vxlan access-if port-vlan xe1/1 3
 no shutdown
 map vnid 16777215
!

VTEP2#show bgp l2vpn evpn
BGP table version is 8, local router ID is 40.40.40.40
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1]
*>   [2]:[0]:[16777215]:[48,0000.2222.2222]:[0]:[16777215]
                       10.10.10.10          0         100           0       64513 64512
i  30.30.30.0      VXLAN
*>   [3]:[16777215]:[32,10.10.10.10]
                       10.10.10.10          0         100           0       64513 64512
i  30.30.30.0      VXLAN

RD[2.2.2.2:2] VRF[vrf12345]:
*>   [2]:[0]:[16777215]:[48,0000.0744.4433]:[0]:[16777215]
```

```
                        40.40.40.40          0          100            32768    i  --------
--      VXLAN
*    [2]:[0]:[16777215]:[48,0000.2222.2222]:[0]:[16777215]
                        10.10.10.10          0          100            0        64513 64512
i  30.30.30.0      VXLAN
*    [3]:[16777215]:[32,10.10.10.10]
                        10.10.10.10          0          100            0        64513 64512
i  30.30.30.0      VXLAN
*>   [3]:[16777215]:[32,40.40.40.40]
                        40.40.40.40          0          100            32768    i  --------
--      VXLAN

Total number of prefixes 6


VTEP2#show nvo vxlan mac-table
================================================================================================
============================================================
                                    VXLAN MAC Entries
================================================================================================
============================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc
------------------------------------------------------------------------------------------------
------------------------------------------------------------

16777215  ----       ----   ----         0000.2222.2222 10.10.10.10
Dynamic Remote       -------             -------
16777215  xe1      3      ----         0000.0744.4433 40.40.40.40
Dynamic Local        -------             -------

Total number of entries are : 2

VTEP2#show vlan brief
Bridge  VLAN ID     Name         State   H/W Status      Member ports
                                                         (u)-Untagged, (t)-Tagged
======= ======= ================ ======= ========== ==========================
32      1       default          ACTIVE  Success    xe1(u) xe49/2(u)
32      2       VLAN0002         ACTIVE  Success
32      3       VLAN0003         ACTIVE  Success
32      4       VLAN0004         ACTIVE  Success    xe1(t) xe49/2(t)
32      5       VLAN0005         ACTIVE  Success

VTEP2#
```

CHAPTER 4 VXLAN Trunk Access Port

In VxLAN, most of the use cases demand to carry the complete traffic received on the access interface to another VTEP access-port. Hence, this support of accepting all tagged and untagged traffic received on the mapped physical port.

Topology

The configurations used in this section use the topology in Figure 4-6.

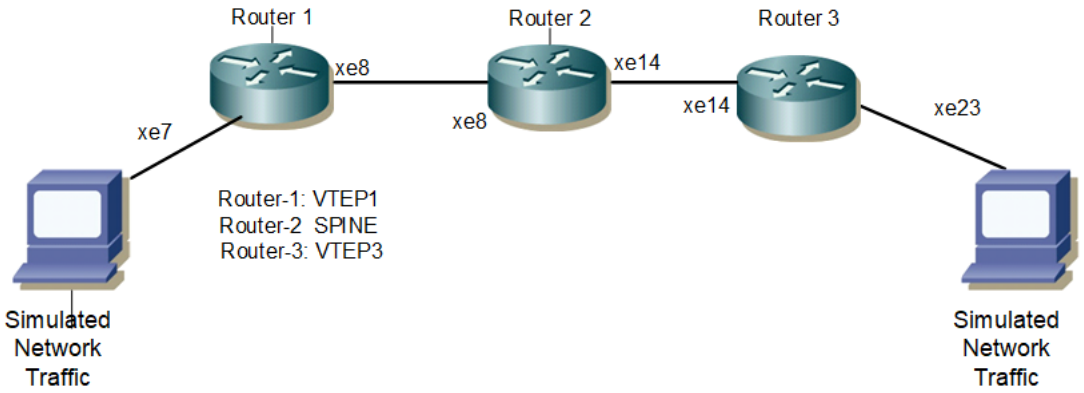

Figure 4-6: VXLAN Trunk Access Port

Base Configuration - L2 VXLAN

VTEP1

Hardware profile and generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#qos enable`	Enable qos
`(config)#commit`	Committing the configurations

Interface and loopback configuration:

`(config)#interface xe7`	Enter Interface mode for xe7
`(config-if)#switchport`	Make it L2 interface
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe8`	Enter Interface mode for xe8
`(config-if)#ip add 10.10.10.1/24`	Configuring the ip address in the network side
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure loopback ip address as 1.1.1.1 for VTEP1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Committing the interface configurations

OSPF configuration:

`(config)#router ospf 1`	Enter into router OSPF mode
`(config-router)#ospf router-id 1.1.1.1`	Configure router-id as 1.1.1.1 (lo ip address)
`(config-router)#network 1.1.1.1/32 area 0.0.0.0`	Add 1.1.1.1 (lo ip address) network into area 0
`(config-router)#network 10.10.10.0/24 area 0.0.0.0`	Add 10.10.10.0(Spine) network into area 0
`(config-router)#bfd all-interfaces`	Enabling bfd on all ospf interface for fast convergence
`(config-router)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Committing the ospf configurations

BGP configuration:

`(config)#router bgp 1`	Enter into Router BGP mode
`(config-router)#bgp router-id 1.1.1.1`	Configure router-id as 1.1.1.1 (lo ip address)
`(config-router)#neighbor 2.2.2.2 remote-as 1`	Specify a VTEP2 loopback ip address and remote-as defined
`(config-router)#neighbor 2.2.2.2 update-source lo`	Configure update as loopback for VTEP2
`(config-router)#neighbor 2.2.2.2 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP2
`(config-router)#address-family l2vpn evpn`	Enter into l2vpn EVPN address family mode
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate 2.2.2.2(VTEP2) into l2vpn evpn address family mode
`(config-router-af)#exit-address-family`	Exit from l2vpn address family mode
`(config-router)#exit`	Exit from Router BGP mode and enter into config mode
`(config)#commit`	Committing the bgp configurations

L2 MAC VRF Configuration:

`(config)#mac vrf L2VRF1`	Create mac routing/forwarding instance with L2VRF1 name and enter into vrf mode
`(config-vrf)#rd 1.1.1.1:1`	Assign RD value
`(config-vrf)#route-target both 1:1`	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
`(config-vrf)#exit`	Exit from vrf mode
`(config)#commit`	Committing the vrf configurations

L2 VXLAN configuration:

`(config)#nvo vxlan enable`	Enable VXLAN
`(config)#nvo vxlan vtep-ip-global 1.1.1.1`	Configure Source vtep-ip-global configuration - Use loopback ip address
`(config)#nvo vxlan id 100 ingress-replication`	Configure VXLAN Network identifier without inner-vid-disabled configured for vxlan trunk access port and enter into VXLAN tenant mode

(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#commit	Committing the vxlan configurations

VTEP2

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#qos enable	Enable qos
(config)#commit	Committing the configurations

Interface and loopback configuration:

(config)#interface xe23	Enter Interface mode for xe23
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe14	Enter Interface mode for xe14
(config-if)#ip add 20.20.20.1/24	Configuring the ip address in the network side
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 2.2.2.2/32 secondary	Configure loopback ip address as 2.2.2.2 for VTEP2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Committing the interface configurations

OSPF configuration:

(config)#router ospf 1	Enter into router OSPF mode
(config-router)#ospf router-id 2.2.2.2	Configure router-id as 2.2.2.2 (lo ip address)
(config-router)#network 2.2.2.2/32 area 0.0.0.0	Add 2.2.2.2 (lo ip address) network into area 0
(config-router)#network 20.20.20.0/24 area 0.0.0.0	Add 20.20.20.0(Spine) network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-router)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Committing the ospf configurations

BGP configuration:

(config)#router bgp 1	Enter into Router BGP mode
(config-router)#bgp router-id 2.2.2.2	Configure router-id as 2.2.2.2 (lo ip address)
(config-router)#neighbor 1.1.1.1 remote-as 1	Specify a VTEP1 loopback ip address and remote-as defined

`(config-router)#neighbor 1.1.1.1 update-source lo`	Configure update as loopback for VTEP2
`(config-router)#neighbor 1.1.1.1 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP1
`(config-router)#address-family l2vpn evpn`	Enter into l2vpn EVPN address family mode
`(config-router-af)#neighbor 1.1.1.1 activate`	Activate 1.1.1.1(VTEP1) into l2vpn evpn address family mode
`(config-router-af)#exit-address-family`	Exit from l2vpn address family mode
`(config-router)#exit`	Exit from Router BGP mode and enter into config mode
`(config)#commit`	Committing the bgp configurations

L2 MAC VRF Configuration:

`(config)#mac vrf L2VRF1`	Create mac routing/forwarding instance with L2VRF1 name and enter into vrf mode
`(config-vrf)#rd 2.2.2.2:1`	Assign RD value
`(config-vrf)#route-target both 1:1`	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
`(config-vrf)#exit`	Exit from vrf mode
`(config)#commit`	Committing the vrf configurations

L2 VXLAN configuration:

`(config)#nvo vxlan enable`	Enable VXLAN
`(config)#nvo vxlan vtep-ip-global 2.2.2.2`	Configure Source vtep-ip-global configuration - Use loopback ip address
`(config)#nvo vxlan id 100 ingress-replication`	Configure VXLAN Network identifier without inner-vid-disabled configured for vxlan trunk access port and enter into VXLAN tenant mode
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1`	Assign vrf for evpn-bgp to carry EVPN route
`(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.
`(config)#commit`	Committing the vxlan configurations

SPINE

Spine node where all VTEPs are connected.

Generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#qos enable`	Enabling qos
`(config)#commit`	Committing the configuration

Interface configuration:

`(config)#interface xe8`	Enter Interface mode for xe8
`(config-if)#ip address 10.10.10.2/24`	Configure ip address as 10.10.10.2 on network side of VTEP1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) interface xe14`	Enter into ce5/1 interface mode
`(config-if)#ip address 20.20.20.2/24`	Configure ip address as 20.20.20.2 on network side of VTEP2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) commit`	Committing the configuration

OSPF configuration:

`(config)#router ospf 1`	Enter into router OSPF mode
`(config-router)#ospf router-id 3.3.3.3`	Configure router-id as 3.3.3.3
`(config-router)#network 10.10.10.0/24 area 0.0.0.0`	Add 10.10.10.0 (VTEP1) network into area 0
`(config-router)#network 20.20.20.0/24 area 0.0.0.0`	Add 20.20.20.0 (VTEP2) network into area 0
`(config-router)#bfd all-interfaces`	Enabling bfd on all ospf interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)commit`	Committing the ospf configuration

VxLAN Trunk Access port as default

In VxLAN, most of the use cases demand to carry the complete traffic received on the access interface to another VTEP access-port. Hence this support of accepting all tagged and untagged traffic received on the mapped physical port.

VTEP1

`(config)#nvo vxlan access-if port xe7 default`	Configuring the vxlan access port as default to receive untagged, single and double tagged traffic
`(config-nvo-acc-if)#map vnid 100`	Map VXLAN Identified to access-port for VXLAN
`(config-nvo-acc-if)#arp-cache disable`	Disable arp-cache - mandatory
`(config-nvo-acc-if)#nd-cache disable`	Disable nd-cache - mandatory
`(config-nvo-acc-if)# mac 0000.1111.0001`	Configure static mac-only
`(config-nvo-acc-if)#exit`	Exit from VXLAN access-interface mode and enter into configuration mode
`(config)#commit`	Committing the vxlan configuration

VTEP2

(config)#nvo vxlan access-if port xe23 default	Configuring the vxlan access port as default to receive untagged, single and double tagged traffic
(config-nvo-acc-if)#map vnid 100	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)#arp-cache disable	Disable arp-cache - mandatory
(config-nvo-acc-if)#nd-cache disable	Disable nd-cache - mandatory
(config-nvo-acc-if)# mac 0000.2222.0001	Configure static mac-only
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#commit	Committing the vxlan configuration

Validation

VTEP1

```
VTEP1#sh run nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan id 100 ingress-replication
 vxlan host-reachability-protocol evpn-bgp L2VRF1
!
nvo vxlan access-if port xe7 default
 map vnid 100
 arp-cache disable
 nd-cache disable
 mac 0000.1111.0001
!
!
VTEP1#sh nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status           Up/Down         Update
================================================================================
1.1.1.1           2.2.2.2          Installed        00:02:49        00:02:49

Total number of entries are 1
VTEP1#sh nvo vxlan mac-table
================================================================================================
================================================
                                                          VXLAN MAC Entries
================================================================================================
================================================
VNID     Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr     VTEP-Ip/ESI
Type             Status     AccessPortDesc
```

```
_____________________________________________________________________________________________
_____________________________________________

100       xe7      ----   ----        ----        0000.1111.0001 1.1.1.1
Static Local     -------     -------
100         ----       ----    ----          0000.2222.0001 2.2.2.2
Static Remote    -------     -------
100       xe7      ----   ----        ----        b0da.1d10.6496 1.1.1.1
Dynamic Local    -------     -------

Total number of entries are : 3

VTEP1#sh nvo vxlan mac-table hardware
=============================================================================================
================================================
                                                              VXLAN MAC Entries
=============================================================================================
================================================
VNID     Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr     VTEP-Ip/ESI
Type             Status     Time-out AccessPortDesc
_____________________________________________________________________________________________
_____________________________________________

100         xe7        ----    ----          0000.1111.0001 1.1.1.1
Static Local     -------     ---       -------
100         ---        ----    ----          0000.2222.0001 2.2.2.2
Remote           -------     ---       -------
100         xe7        ----    ----          b0da.1d10.6496 1.1.1.1
Dynamic Local    -------     300       -------

Total number of entries are 3
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________________
100      ----         L2       NW   ----      ----                         ---- ----
1.1.1.1          2.2.2.2
100      ----         --       AC   xe7       --- Single Homed Port ---     ---- ----
----             ----

Total number of entries are 2
VTEP1#sh nvo vxlan route-count
VXLAN Active route count information
====================================
Max route count   : 32768
Active route count: 3
```

```
-------------------------------------------
 VNID     Total    MACONLY  MACIPv4  MACIPv6
-------------------------------------------
 100      3        3        0        0

Total number of entries are 1
VTEP1#sh nvo vxlan access-if-config
nvo vxlan access-if port xe7 default
 map vnid 100
 arp-cache disable
 nd-cache disable
 mac 0000.1111.0001
!
VTEP1#sh nvo vxlan access-if brief

                 Inner                    Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
xe7        ---   ---   500000   100       up      up

Total number of entries are 1
VTEP1#sh bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 2
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
2.2.2.2                 4     1   15        17      2     0     0  00:04:52
2      0      1      1      0      0

Total number of neighbors 1

Total number of Established sessions 1
VTEP1#sh bgp l2vpn evpn
BGP table version is 2, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap
```

```
RD[1.1.1.1:1] VRF[L2VRF1]:
*>   [2]:[0]:[100]:[48,0000:1111:0001]:[0]:[100]
                     1.1.1.1             0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[100]:[48,0000:2222:0001]:[0]:[100]
                    2.2.2.2             0        100       0    i  2.2.2.2        VXLAN
*>   [2]:[0]:[100]:[48,b0da:1d10:6496]:[0]:[100]
                     1.1.1.1             0        100       32768  i  ----------
VXLAN
*>   [3]:[100]:[32,1.1.1.1]
                     1.1.1.1             0        100       32768  i  ----------
VXLAN
* i  [3]:[100]:[32,2.2.2.2]
                    2.2.2.2             0        100       0    i  2.2.2.2        VXLAN

RD[2.2.2.2:1]
*>i  [2]:[0]:[100]:[48,0000:2222:0001]:[0]:[100]
                    2.2.2.2             0        100       0    i  2.2.2.2        VXLAN
*>i  [3]:[100]:[32,2.2.2.2]
                    2.2.2.2             0        100       0    i  2.2.2.2        VXLAN

Total number of prefixes 7
VTEP1#sh nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
Total number of entries are 0
VTEP1#sh nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________
_____________
Total number of entries are 0
VTEP1#
```

VTEP2

```
VTEP2#sh run nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 100 ingress-replication
 vxlan host-reachability-protocol evpn-bgp L2VRF1
!
```

```
nvo vxlan access-if port xe23 default
 map vnid 100
 arp-cache disable
 nd-cache disable
 mac 0000.2222.0001
!
!
VTEP2#sh nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
2.2.2.2          1.1.1.1          Installed        00:05:47       00:05:47

Total number of entries are 1
VTEP2#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________________
100      ----         L2       NW   ----      ----                           ---- ----
2.2.2.2          1.1.1.1
100      ----         --       AC   xe23      --- Single Homed Port ---      ---- ----
----             ----

Total number of entries are 2
VTEP2#sh nvo vxlan mac-table
================================================================================================
================================================
                                                          VXLAN MAC Entries
================================================================================================
================================================
VNID     Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr     VTEP-Ip/ESI
Type             Status     AccessPortDesc
_____________________________________________________________________________________________
_____________________________________________

100        ----      ----   ----         0000.1111.0001 1.1.1.1
Static Remote   -------    -------
100      xe23     ----  ----       ----       0000.2222.0001 2.2.2.2
Static Local    -------    -------
100        ----      ----   ----         b0da.1d10.6496 1.1.1.1
Dynamic Remote  -------    -------

Total number of entries are : 3

VTEP2#sh nvo vxlan mac-table hardware
```

```
================================================================================================
================================================
                                                                VXLAN MAC Entries
================================================================================================
================================================
VNID     Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr     VTEP-Ip/ESI
Type             Status     Time-out AccessPortDesc
________________________________________________________________________________________________
_____________________________________________

100        ---       ----   ----         0000.1111.0001 1.1.1.1
Remote          -------    ---      -------
100        xe23      ----   ----         0000.2222.0001 2.2.2.2
Static Local    -------    ---      -------
100        ---       ----   ----         b0da.1d10.6496 1.1.1.1
Remote          -------    ---      -------

Total number of entries are 3
VTEP2#sh nvo vxlan route-count
VXLAN Active route count information
==================================
Max route count   : 32768
Active route count: 3

-------------------------------------------
 VNID     Total    MACONLY  MACIPv4  MACIPv6
-------------------------------------------
 100      3        3        0        0

Total number of entries are 1
VTEP2#sh nvo vxlan access-if br

                 Inner                     Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
xe23       ---   ---   500000   100       up      up

Total number of entries are 1
VTEP2#sh nvo vxlan access-if-config
nvo vxlan access-if port xe23 default
 map vnid 100
 arp-cache disable
 nd-cache disable
 mac 0000.2222.0001
!
VTEP2#sh nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
Total number of entries are 0
VTEP2#sh nvo vxlan nd-cache
```

```
VXLAN ND-CACHE Information
==========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________
_____________
Total number of entries are 0

VTEP2#
VTEP2#sh bgp l2vpn evpn summary
BGP router identifier 2.2.2.2, local AS number 1
BGP table version is 2
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4     1   27        27      1     0     0  00:09:54
3      0      2      1      0      0

Total number of neighbors 1

Total number of Established sessions 1
VTEP2#sh bgp l2vpn evpn
BGP table version is 2, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network          Next Hop            Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1]
*>i  [2]:[0]:[100]:[48,0000:1111:0001]:[0]:[100]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN
*>i  [2]:[0]:[100]:[48,b0da:1d10:6496]:[0]:[100]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN
*>i  [3]:[100]:[32,1.1.1.1]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN

RD[2.2.2.2:1] VRF[L2VRF1]:
* i  [2]:[0]:[100]:[48,0000:1111:0001]:[0]:[100]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN
*>   [2]:[0]:[100]:[48,0000:2222:0001]:[0]:[100]
```

```
                      2.2.2.2              0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[100]:[48,b0da:1d10:6496]:[0]:[100]
                     1.1.1.1              0        100       0    i  1.1.1.1        VXLAN
* i  [3]:[100]:[32,1.1.1.1]
                     1.1.1.1              0        100       0    i  1.1.1.1        VXLAN
*>   [3]:[100]:[32,2.2.2.2]
                      2.2.2.2              0        100       32768  i  ----------
VXLAN

Total number of prefixes 8
VTEP2#
```

VxLAN Trunk access port with vlan range

When access port with a specific vlan range configured, all the traffic in that specific range are accepted and forwaded.

VTEP1

`(config)#nvo vxlan access-if port-vlan xe7 2-100`	Configuring the vxlan access port with vlan range 2-100 where traffic in the vlan range 2-100 are accepted
`(config-nvo-acc-if)#map vnid 100`	Map VXLAN Identified to access-port for VXLAN
`(config-nvo-acc-if)#arp-cache disable`	Disable arp-cache - mandatory
`(config-nvo-acc-if)#nd-cache disable`	Disable nd-cache - mandatory
`(config-nvo-acc-if)# mac 0000.1111.0001`	Configure static mac-only
`(config-nvo-acc-if)#exit`	Exit from VXLAN access-interface mode and enter into configuration mode
`(config)#commit`	Committing the vxlan configuration

VTEP2

`(config)#nvo vxlan access-if port-vlan xe23 2-100`	Configuring the vxlan access port with vlan range 2-100 where traffic in the vlan range 2-100 are accepted
`(config-nvo-acc-if)#map vnid 100`	Map VXLAN Identified to access-port for VXLAN
`(config-nvo-acc-if)#arp-cache disable`	Disable arp-cache - mandatory
`(config-nvo-acc-if)#nd-cache disable`	Disable nd-cache - mandatory
`(config-nvo-acc-if)# mac 0000.2222.0001`	Configure static mac-only
`(config-nvo-acc-if)#exit`	Exit from VXLAN access-interface mode and enter into configuration mode
`(config)#commit`	Committing the vxlan configuration

Validations

VTEP1

```
VTEP1#sh run nvo vx
```

```
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan id 100 ingress-replication
 vxlan host-reachability-protocol evpn-bgp L2VRF1
!
nvo vxlan access-if port-vlan xe7 2-100
 map vnid 100
 arp-cache disable
 nd-cache disable
 mac 0000.1111.0001
!
!
VTEP1#
VTEP1#sh nvo vxlan tunnel summary

Total number of entries: 1 [Installed: 1, Resolved: 0, Unresolved: 0]

Total number of entries are 1
VTEP1#sh nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
1.1.1.1          2.2.2.2          Installed        00:20:10       00:20:10

Total number of entries are 1
VTEP1#sh nvo vx
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________________
100      ----         L2       NW   ----      ----                         ---- ----
1.1.1.1          2.2.2.2
100      ----         --       AC   xe7       --- Single Homed Port ---     2    ----
----             ----

Total number of entries are 2
VTEP1#sh nvo vxlan mac-table
================================================================================
================================================
                                                        VXLAN MAC Entries
================================================================================
================================================
```

```
VNID     Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr     VTEP-Ip/ESI
Type             Status     AccessPortDesc
_____________________________________________________________________________________________
_____________________________________________

100      xe7      -----       2-100 ----       0000.1111.0001 1.1.1.1
Static Local    -------    -------
100        ----       ----   ----         0000.2222.0001 2.2.2.2
Static Remote   -------    -------
100      xe7      -----       2-100 ----       b0da.1d10.6496 1.1.1.1
Dynamic Local   -------    -------

Total number of entries are : 3

VTEP1#sh nvo vxlan mac-table hardware
=============================================================================================
=================================================
                                                                VXLAN MAC Entries
=============================================================================================
=================================================
VNID     Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr     VTEP-Ip/ESI
Type             Status     Time-out AccessPortDesc
_____________________________________________________________________________________________
_____________________________________________

100      xe7      -----       2-100 ----       0000.1111.0001 1.1.1.1
Static Local    -------    ---      -------
100        ---       ----   ----         0000.2222.0001 2.2.2.2
Remote          -------    ---      -------
100      xe7      -----       2-100 ----       b0da.1d10.6496 1.1.1.1
Dynamic Local   -------    300      -------

Total number of entries are 3
VTEP1#sh nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
Total number of entries are 0
VTEP1#sh nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________________
_____________
Total number of entries are 0
VTEP1#sh nvo vxlan access-if-config
nvo vxlan access-if port-vlan xe7 2-100
 map vnid 100
 arp-cache disable
 nd-cache disable
 mac 0000.1111.0001
!
```

```
VTEP1#sh nvo vxlan access-if brief

                 Inner                      Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid       status  status
-----------------------------------------------------------
xe7        2     ---   500000   100        up      up

Total number of entries are 1
VTEP1#
VTEP1#sh nvo vxlan route-count
VXLAN Active route count information
====================================
Max route count   : 32768
Active route count: 3

-------------------------------------------
 VNID     Total    MACONLY  MACIPv4  MACIPv6
-------------------------------------------
 100      3        3        0        0

Total number of entries are 1
VTEP1#sh bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
2.2.2.2                 4     1   58        61       5     0     0  00:22:05
2      0      1      1      0      0

Total number of neighbors 1

Total number of Established sessions 1
VTEP1#sh bgp l2vpn evpn
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap
```

```
RD[1.1.1.1:1] VRF[L2VRF1]:
*>   [2]:[0]:[100]:[48,0000:1111:0001]:[0]:[100]
                     1.1.1.1             0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[100]:[48,0000:2222:0001]:[0]:[100]
                    2.2.2.2             0        100       0    i  2.2.2.2        VXLAN
*>   [2]:[0]:[100]:[48,b0da:1d10:6496]:[0]:[100]
                     1.1.1.1             0        100       32768  i  ----------
VXLAN
*>   [3]:[100]:[32,1.1.1.1]
                     1.1.1.1             0        100       32768  i  ----------
VXLAN
* i  [3]:[100]:[32,2.2.2.2]
                    2.2.2.2             0        100       0    i  2.2.2.2        VXLAN

RD[2.2.2.2:1]
*>i  [2]:[0]:[100]:[48,0000:2222:0001]:[0]:[100]
                    2.2.2.2             0        100       0    i  2.2.2.2        VXLAN
*>i  [3]:[100]:[32,2.2.2.2]
                    2.2.2.2             0        100       0    i  2.2.2.2        VXLAN

Total number of prefixes 7
```

VTEP2

```
VTEP2#sh run nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 100 ingress-replication
 vxlan host-reachability-protocol evpn-bgp L2VRF1
!
nvo vxlan access-if port-vlan xe23 2-100
 map vnid 100
 arp-cache disable
 nd-cache disable
 mac 0000.2222.0001
!
!
VTEP2#sh nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
2.2.2.2          1.1.1.1          Installed        00:19:28       00:19:28

Total number of entries are 1
VTEP2#sh nvo vxlan tunnel
```

```
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
2.2.2.2          1.1.1.1          Installed        00:21:06       00:21:06

Total number of entries are 1
VTEP2#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________________
100      ----         L2       NW   ----      ----                         ---- ----
2.2.2.2          1.1.1.1
100      ----         --       AC   xe23      --- Single Homed Port ---     2    ----
----             ----

Total number of entries are 2
VTEP2#sh nvo vxlan mac-table
=============================================================================================
=================================================
                                                          VXLAN MAC Entries
=============================================================================================
=================================================
VNID     Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr     VTEP-Ip/ESI
Type             Status     AccessPortDesc
_____________________________________________________________________________________________
_____________________________________________________

100        ----      ----   ----         0000.1111.0001 1.1.1.1
Static Remote   -------    -------
100      xe23     -----      2-100 ----       0000.2222.0001 2.2.2.2
Static Local    -------    -------
100        ----      ----   ----         b0da.1d10.6496 1.1.1.1
Dynamic Remote  -------    -------

Total number of entries are : 3

VTEP2#sh nvo vxlan mac-table hardware
=============================================================================================
=================================================
                                                          VXLAN MAC Entries
=============================================================================================
=================================================
VNID     Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr     VTEP-Ip/ESI
Type             Status     Time-out AccessPortDesc
_____________________________________________________________________________________________
_____________________________________________________
```

```
100        ---       ----   ----         0000.1111.0001 1.1.1.1
Remote          -------    ---      -------
100       xe23     -----       2-100 ----       0000.2222.0001 2.2.2.2
Static Local    -------    ---      -------
100        ---       ----   ----         b0da.1d10.6496 1.1.1.1
Remote          -------    ---      -------

Total number of entries are 3
VTEP2#sh nvo vxlan route-count
VXLAN Active route count information
====================================
Max route count   : 32768
Active route count: 3

-------------------------------------------
 VNID     Total    MACONLY  MACIPv4  MACIPv6
-------------------------------------------
 100      3        3        0        0

Total number of entries are 1
VTEP2#sh nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
Total number of entries are 0
VTEP2#sh nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________
_____________
Total number of entries are 0
VTEP2#
VTEP2#
VTEP2#sh nvo vxlan access-if brief

                Inner                    Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
xe23       2     ---   500000   100       up      up

Total number of entries are 1
VTEP2#sh nvo vxlan access-if-config
nvo vxlan access-if port-vlan xe23 2-100
 map vnid 100
 arp-cache disable
 nd-cache disable
 mac 0000.2222.0001
!
```

```
VTEP2#sh bgp l2vpn evpn summary
BGP router identifier 2.2.2.2, local AS number 1
BGP table version is 4
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4    1   59        57      3     0     0  00:21:48
3      0      2      1      0      0

Total number of neighbors 1

Total number of Established sessions 1
VTEP2#sh bgp l2vpn evpn
BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1]
*>i  [2]:[0]:[100]:[48,0000:1111:0001]:[0]:[100]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN
*>i  [2]:[0]:[100]:[48,b0da:1d10:6496]:[0]:[100]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN
*>i  [3]:[100]:[32,1.1.1.1]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN

RD[2.2.2.2:1] VRF[L2VRF1]:
* i  [2]:[0]:[100]:[48,0000:1111:0001]:[0]:[100]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN
*>   [2]:[0]:[100]:[48,0000:2222:0001]:[0]:[100]
                      2.2.2.2             0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[100]:[48,b0da:1d10:6496]:[0]:[100]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN
* i  [3]:[100]:[32,1.1.1.1]
                     1.1.1.1             0       100       0    i  1.1.1.1        VXLAN
*>   [3]:[100]:[32,2.2.2.2]
                      2.2.2.2             0        100       32768  i  ----------
VXLAN
```

```
Total number of prefixes 8
VTEP2#
```

CHAPTER 5 EVPN-VXLAN Hybrid Port Support

This chapter contains the configurations for EVPN-VXLAN Hybrid Port Support.

Overview

This feature is to map VxLAN untagged and tagged (port+vlan) access interfaces on the same parent interface (Physical, Dynamic and Static LAG). It also supports the Bridge vlan and vxlan untagged access on the same parent interface.

Topology

The procedures in this section use the topology in Figure 5-7

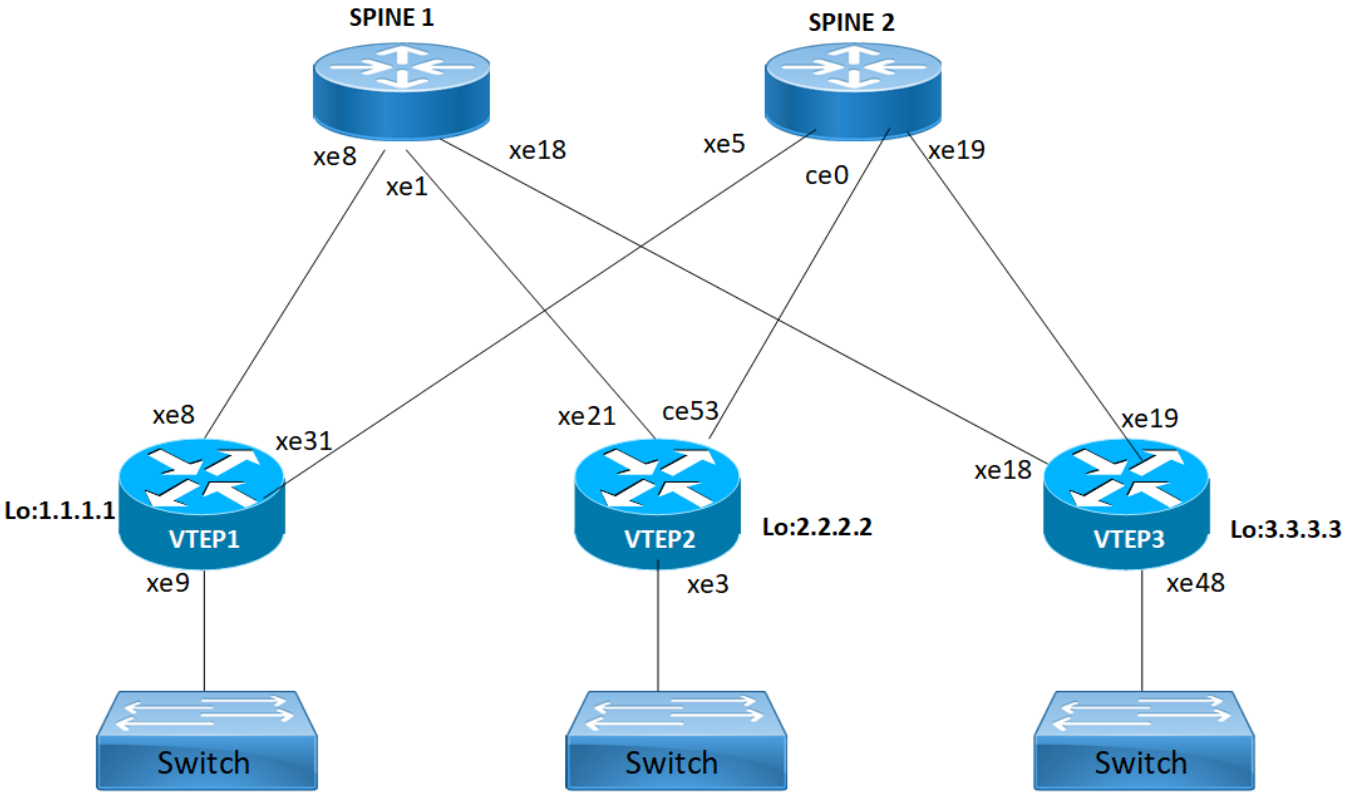

Figure 5-7: EVPN-VxLAN hybrid port

EVPN-VxLAN Hybrid Port Configuration

VTEP1

Interface and Loopback Configuration

(config)#interface xe9	Enter Interface mode for xe9
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 1.1.1.1/32 secondary	Configure loopback ip address as 1.1.1.1 for VTEP1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe8	Enter Interface mode for xe8
(config-if)#ip address 10.10.10.1/24	Configure IP address as 10.10.10.1 on network side of Spine1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe31	Enter Interface mode for xe31
(config-if)#ip address 20.20.20.1/24	Configure IP address as 20.20.20.1 on network side of Spine2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

OSPF Configuration

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 1.1.1.1	Configure router-id as 1.1.1.1 (lo ip address)
(config-router)#network 1.1.1.1/32 area 0.0.0.0	Add 1.1.1.1 (lo IP address) network into area 0
(config-router)#network 10.10.10.0/24 area 0.0.0.0	Add 10.10.10.0 (Spine1) network into area 0
(config-router)#network 20.20.20.0/24 area 0.0.0.0	Add 20.20.20.0 (Spine2) network into area 0
(config-router)#bfd all-interfaces	Enabling BFD on all OSPF interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

BGP Configuration

`(config)#router bgp 500`	Enter into Router BGP mode
`(config-router)#bgp router-id 1.1.1.1`	Configure router-id as 1.1.1.1 (lo IP address)
`(config-router)#neighbor 2.2.2.2 remote-as 500`	Specify a VTEP2 loopback IP address and remote-as defined
`(config-router)#neighbor 2.2.2.2 update-source lo`	Configure update as loopback for VTEP2
`(config-router)#neighbor 2.2.2.2 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP2
`(config-router)#neighbor 3.3.3.3 remote-as 500`	Specify a VTEP3 loopback IP address and remote-as defined
`(config-router)#neighbor 3.3.3.3 update-source lo`	Configure update as loopback for VTEP3
`(config-router)#neighbor 3.3.3.3 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP3
`(config-router)#address-family ipv4 unicast`	Enter into ipv4 unicast address family mode
`(config-router-af)#network 1.1.1.1/32`	Advertise loopback network into BGP for VTEP ID reachability
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate 2.2.2.2 (VTEP2) into ipv4 unicast address family mode
`(config-router-af)#neighbor 3.3.3.3 activate`	Activate 3.3.3.3 (VTEP2) into ipv4 unicast address family mode
`(config-router-af)#exit-address-family`	Exit from ipv4 unicast address family mode
`(config-router)#address-family l2vpn evpn`	Enter into L2VPN EVPN address family mode
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate 2.2.2.2 (VTEP2) into L2VPN evpn address family mode
`(config-router-af)#neighbor 3.3.3.3 activate`	Activate 3.3.3.3 (VTEP3) into L2VPN evpn address family mode
`(config-router-af)#exit-address-family`	Exit from L2VPN address family mode
`(config-router)#exit`	Exit from Router BGP mode and enter into config mode
`(config)#commit`	Commit the candidate configuration to the running configuration

VRF Configuration

`(config)#mac vrf VRF1`	Create mac routing/forwarding instance with VRF1 name and enter into VRF mode
`(config-vrf)#rd 1.1.1.1:11`	Assign RD value
`(config-vrf)#route-target both 9.9.9.9:100`	Assign route-target value for same for import and export. Should be same on all node for VRF1
`(config-vrf)#exit`	Exit from VRF mode
`(config)#commit`	Commit the candidate configuration to the running configuration

VxLAN Configuration

`(config)#nvo vxlan enable`	Enable VxLAN
`(config)#nvo vxlan vtep-ip-global 1.1.1.1`	Configure Source VTEP-IP-global configuration
`(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled`	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1`	Assign VRF for EVPN-BGP to carry EVPN route
`(config-nvo)#vni-name VNI-RED`	Configure VNI-name as VNI-RED
`(config-nvo)#exit`	Exit from VxLAN tenant mode and enter into configuration mode.
`(config)#nvo vxlan id 20 ingress-replication inner-vid-disabled`	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1`	Assign VRF for EVPN-BGP to carry EVPN route
`(config-nvo)#vni-name VNI-BLUE`	Configure VNI-name as VNI-BLUE
`(config-nvo)#exit`	Exit from VxLAN tenant mode and enter into configuration mode.
`(config)#nvo vxlan access-if port xe9`	Enable port mapping i.e. access port to the physical interface
`(config-nvo-acc-if)#map vni-name VNI-RED`	Map VxLAN Identified to access-port for VxLAN
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-interface mode and enter into configuration mode
`(config)#nvo vxlan access-if port-vlan xe9 1001`	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) – physical interface
`(config-nvo-acc-if)#map vni-name VNI-RED`	Map VxLAN Identified to access-port for VxLAN
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-interface mode and enter into configuration mode
`(config)#nvo vxlan access-if port-vlan xe9 2001`	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) – physical interface
`(config-nvo-acc-if)#map vni-name VNI-BLUE`	Map VxLAN Identified to access-port for VxLAN
`(config-nvo-acc-if)#exit`	Exit from VxLAN access-interface mode and enter into configuration mode

(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit from configuration mode

VTEP2

Interface and Loopback Configuration

(config)#interface xe3	Enter Interface mode for xe3
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 2.2.2.2/32 secondary	Configure loopback IP address as 2.2.2.2 for VTEP2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe21	Enter Interface mode for xe21
(config-if)#ip address 30.30.30.1/24	Configure IP address as 30.30.30.1 on network side of Spine1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface ce53	Enter Interface mode for ce53
(config-if)#ip address 40.40.40.1/24	Configure IP address as 40.40.40.1 on network side of Spine2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

OSPF Configuration

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 2.2.2.2	Configure router-id as 2.2.2.2 (lo IP address)
(config-router)#network 2.2.2.2/32 area 0.0.0.0	Add 2.2.2.2 (lo IP address) network into area 0
(config-router)#network 30.30.30.0/24 area 0.0.0.0	Add 30.30.30.0 (Spine1) network into area 0
(config-router)#network 40.40.40.0/24 area 0.0.0.0	Add 40.40.40.0 (Spine2) network into area 0
(config-router)#bfd all-interfaces	Enabling BFD on all OSPF interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

BGP Configuration

`(config)#router bgp 500`	Enter into Router BGP mode
`(config-router)#bgp router-id 2.2.2.2`	Configure router-id as 2.2.2.2 (lo IP address)
`(config-router)#neighbor 1.1.1.1 remote-as 500`	Specify a VTEP1 loopback IP address and remote-as defined
`(config-router)#neighbor 1.1.1.1 update-source lo`	Configure update as loopback for VTEP1
`(config-router)#neighbor 1.1.1.1 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP1
`(config-router)#neighbor 3.3.3.3 remote-as 500`	Specify a VTEP3 loopback IP address and remote-as defined
`(config-router)#neighbor 3.3.3.3 update-source lo`	Configure update as loopback for VTEP3
`(config-router)#neighbor 3.3.3.3 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP3
`(config-router)#address-family ipv4 unicast`	Enter into ipv4 unicast address family mode
`(config-router-af)#network 2.2.2.2/32`	Advertise loopback network into BGP for VTEP ID reachability
`(config-router-af)#neighbor 1.1.1.1 activate`	Activate 1.1.1.1 (VTEP2) into ipv4 unicast address family mode
`(config-router-af)#neighbor 3.3.3.3 activate`	Activate 3.3.3.3 (VTEP2) into ipv4 unicast address family mode
`(config-router-af)#exit-address-family`	Exit from ipv4 unicast address family mode
`(config-router)#address-family l2vpn evpn`	Enter into L2VPN EVPN address family mode
`(config-router-af)#neighbor 1.1.1.1 activate`	Activate 1.1.1.1(VTEP1) into L2VPN evpn address family mode
`(config-router-af)#neighbor 3.3.3.3 activate`	Activate 3.3.3.3(VTEP3) into L2VPN evpn address family mode
`(config-router-af)#exit-address-family`	Exit from L2VPN address family mode
`(config-router)#exit`	Exit from Router BGP mode and enter into config mode
`(config)#commit`	Commit the candidate configuration to the running configuration

VRF Configuration

`(config)# mac vrf VRF1`	Create mac routing/forwarding instance with VRF1 name and enter into VRF mode
`(config-vrf)#rd 2.2.2.2:11`	Assign RD value
`(config-vrf)#route-target both 9.9.9.9:100`	Assign route-target value for same for import and export. Should be same on all node for VRF1
`(config-vrf)#exit`	Exit from VRF mode
`(config)#commit`	Commit the candidate configuration to the running configuration

VxLAN Configuration

(config)#nvo vxlan enable	Enable VxLAN
(config)#nvo vxlan vtep-ip-global 2.2.2.2	Configure Source VTEP-IP-global configuration
(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#vni-name VNI-RED	Configure VNI-name as VNI-RED
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 20 ingress-replication inner-vid-disabled	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#vni-name VNI-BLUE	Configure VNI-name as VNI-BLUE
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan access-if port xe3	Enable port mapping i.e. access port to physical interface
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan xe3 1002	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) – physical interface
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan xe3 2001	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) – physical interface
(config-nvo-acc-if)#map vni-name VNI-BLUE	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit from configuration mode

VTEP3

Interface and loopback configuration

`#configure terminal`	Enter Configure mode.
`(config)#interface xe48`	Enter Interface mode for xe48 (SH3)
`(config-if)#switchport`	Make it L2 interface
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 3.3.3.3/32 secondary`	Configure loopback IP address as 3.3.3.3 for VTEP3
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe18`	Enter Interface mode for xe18
`(config-if)#ip address 50.50.50.1/24`	Configure IP address as 50.50.50.1 on network side of Spine1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe19`	Enter Interface mode for xe19
`(config-if)#ip address 60.60.60.1/24`	Configure IP address as 60.60.60.1 on network side of Spine2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

OSPF Configuration

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 3.3.3.3`	Configure router-ID as 3.3.3.3 (lo IP address)
`(config-router)#network 3.3.3.3/32 area 0.0.0.0`	Add 3.3.3.3 (lo IP address) network into area 0
`(config-router)#network 50.50.50.0/24 area 0.0.0.0`	Add 50.50.50.0 (Spine1) network into area 0
`(config-router)#network 60.60.60.0/24 area 0.0.0.0`	Add 60.60.60.0 (Spine2) network into area 0
`(config-router)#bfd all-interfaces`	Enabling BFD on all OSPF interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

BGP Configuration

(config)#router bgp 500	Enter into Router BGP mode
(config-router)#bgp router-id 3.3.3.3	Configure router-ID as 3.3.3.3 (lo ip address)
(config-router)#neighbor 1.1.1.1 remote-as 500	Specify a VTEP1 loopback IP address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Configure update as loopback for VTEP1
(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP1
(config-router)#neighbor 2.2.2.2 remote-as 500	Specify a VTEP2 loopback IP address and remote-as defined
(config-router)#neighbor 2.2.2.2 update-source lo	Configure update as loopback for VTEP2
(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP3
(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family mode
(config-router-af)#network 3.3.3.3/32	Advertise loopback network into BGP for VTEP ID reachability
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1 (VTEP2) into ipv4 unicast address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2 (VTEP2) into ipv4 unicast address family mode
(config-router-af)#exit-address-family	Exit from ipv4 unicast address family mode
(config-router)#address-family l2vpn evpn	Enter into L2VPN EVPN address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1 (VTEP1) into L2VPN evpn address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2 (VTEP2) into L2VPN evpn address family mode
(config-router-af)#exit-address-family	Exit from L2VPN address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to the running configuration

VRF Configuration

(config)# mac vrf VRF1	Create MAC routing/forwarding instance with VRF1 name and enter into VRF mode
(config-vrf)#rd 3.3.3.3:11	Assign RD value
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for VRF1
(config-vrf)#exit	Exit from VRF mode
(config)#commit	Commit the candidate configuration to the running configuration

VxLAN Configuration

(config)#nvo vxlan enable	Enable VxLAN
(config)#nvo vxlan vtep-ip-global 3.3.3.3	Configure Source VTEP-IP-global configuration
(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#vni-name VNI-RED	Configure VNI-name as VNI-RED
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 20 ingress-replication inner-vid-disabled	Configure VxLAN Network identifier with/without inner-VID-disabled configure and enter into VxLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#vni-name VNI-BLUE	Configure VNI-name as VNI-BLUE
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan access-if port xe48	Enable port mapping i.e. access port to physical interface
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan xe48 1001	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) – physical interface
(config-nvo-acc-if)#map vni-name VNI-RED	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan xe48 2001	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) – physical interface
(config-nvo-acc-if)#map vni-name VNI-BLUE	Map VxLAN Identified to access-port for VxLAN
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode

(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit from configuration mode

Spine 1

Spine node where all VTEPs are connected

Interface and Loopback Configuration

#configure terminal	Enter Configure mode.
(config)#qos enable	Enabling QoS
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 11.11.11.11/32 secondary	Configure loopback IP address as 11.11.11.11 for Spine1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe8	Enter Interface mode for xe8
(config-if)#ip address 10.10.10.2/24	Configure IP address as 10.10.10.2 on network side of VTEP1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe1	Enter Interface mode for xe1
(config-if)#ip address 30.30.30.2/24	Configure IP address as 30.30.30.2 on network side of VTEP2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe18	Enter Interface mode for xe18
(config-if)#ip address 50.50.50.2/24	Configure IP address as 50.50.50.2 on network side of VTEP3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

OSPF configuration

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 11.11.11.11	Configure router-ID as 11.11.11.11 (lo IP address)
(config-router)#network 11.11.11.11/32 area 0.0.0.0	Add 11.11.11.11 (lo IP address) network into area 0
(config-router)#network 10.10.10.0/24 area 0.0.0.0	Add 10.10.10.0 (VTEP1) network into area 0
(config-router)#network 30.30.30.0/24 area 0.0.0.0	Add 30.30.30.0 (VTEP2) network into area 0
(config-router)#network 50.50.50.0/24 area 0.0.0.0	Add 50.50.50.0 (VTEP3) network into area 0

`(config-router)#bfd all-interfaces`	Enabling BFD on all OSPF interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

Spine 2

Spine node where all VTEPs are connected

Interface and loopback configuration

`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 22.22.22.22/32 secondary`	Configure loopback IP address as 22.22.22.22 for Spine2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe5`	Enter Interface mode for xe5
`(config-if)#ip address 20.20.20.2/24`	Configure IP address as 20.20.20.2 on network side of VTEP1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface ce0`	Enter Interface mode for ce0
`(config-if)#ip address 40.40.40.2/24`	Configure IP address as 40.40.40.2 on network side of VTEP2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe19`	Enter Interface mode for xe19
`(config-if)#ip address 60.60.60.2/24`	Configure IP address as 60.60.60.2 on network side of VTEP3
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

OSPF configuration

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 22.22.22.22	Configure router-id as 11.11.11.11 (lo IP address)
(config-router)#network 22.22.22.22/32 area 0.0.0.0	Add 22.22.22.22 (lo IP address) network into area 0
(config-router)#network 20.20.20.0/24 area 0.0.0.0	Add 20.20.20.0 (VTEP1) network into area 0
(config-router)#network 40.40.40.0/24 area 0.0.0.0	Add 40.40.40.0 (VTEP2) network into area 0
(config-router)#network 60.60.60.0/24 area 0.0.0.0	Add 60.60.60.0 (VTEP3) network into area 0
(config-router)#bfd all-interfaces	Enabling BFD on all OSPF interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

Validation

VTEP1

```
VTEP1#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
10       VNI-RED      L2       NW   ----      ----                            ----      -
---      1.1.1.1          3.3.3.3
10       VNI-RED      L2       NW   ----      ----                            ----      -
---      1.1.1.1          2.2.2.2
10       VNI-RED      --       AC   xe9       --- Single Homed Port ---     0         --
--       ----             ----
10       VNI-RED      --       AC   xe9       --- Single Homed Port ---     1001      --
--       ----             ----
20       VNI-BLUE     L2       NW   ----      ----                            ----      -
---      1.1.1.1          3.3.3.3
20       VNI-BLUE     L2       NW   ----      ----                            ----      -
---      1.1.1.1          2.2.2.2
20       VNI-BLUE     --       AC   xe9       --- Single Homed Port ---     2001      --
--       ----             ----

Total number of entries are 7
```

```
VTEP1#sh nvo vxlan access-if brief

                 Inner                      Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid       status  status
-----------------------------------------------------------
xe9        ---   ---   500000   10         up      up
xe9        1001  ---   500001   10         up      up
xe9        2001  ---   500002   20         up      up

Total number of entries are 3


VTEP1#show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 500
BGP table version is 4
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
2.2.2.2                 4   500   53        54      3     0     0  00:20:58
2      0      0      2      0      0
3.3.3.3                 4   500   53        53      3     0     0  00:20:58
2      0      0      2      0      0

Total number of neighbors 2

Total number of Established sessions 2


VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=====================================================================
1.1.1.1          3.3.3.3          Installed     00:10:44      00:10:44
1.1.1.1          2.2.2.2          Installed     00:10:58      00:10:58

Total number of entries are 2


VTEP1#show bgp l2vpn evpn
BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
```

```
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network          Next Hop            Metric    LocPrf      Weight    Path  Peer
Encap

RD[1.1.1.1:1] VRF[VRF1]:
*> [3]:[10]:[32,1.1.1.1]
                     1.1.1.1               0        100       32768  i  ----------
VXLAN
* i[3]:[10]:[32,2.2.2.2]
                    2.2.2.2               0        100       0    i  2.2.2.2        VXLAN
* i[3]:[10]:[32,3.3.3.3]
                    3.3.3.3               0        100       0    i  3.3.3.3        VXLAN
*> [3]:[20]:[32,1.1.1.1]
                     1.1.1.1               0        100       32768  i  ----------
VXLAN
* i[3]:[20]:[32,2.2.2.2]
                    2.2.2.2               0        100       0    i  2.2.2.2        VXLAN
* i[3]:[20]:[32,3.3.3.3]
                    3.3.3.3               0        100       0    i  3.3.3.3        VXLAN

RD[2.2.2.2:1]
*>i[3]:[10]:[32,2.2.2.2]
                    2.2.2.2               0        100       0    i  2.2.2.2        VXLAN
*>i[3]:[20]:[32,2.2.2.2]
                    2.2.2.2               0        100       0    i  2.2.2.2        VXLAN

RD[3.3.3.3:1]
*>i[3]:[10]:[32,3.3.3.3]
                    3.3.3.3               0        100       0    i  3.3.3.3        VXLAN
*>i[3]:[20]:[32,3.3.3.3]
                    3.3.3.3               0        100       0    i  3.3.3.3        VXLAN

Total number of prefixes 10
```

VTEP2

```
VTEP2#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________________
```

```
10       VNI-RED      L2       NW   ----      ----                          ----      -
---      2.2.2.2          1.1.1.1
10       VNI-RED      L2       NW   ----      ----                          ----      -
---      2.2.2.2          3.3.3.3
10       VNI-RED      --       AC   xe8       --- Single Homed Port ---     0         -
---      ----             ----
10       VNI-RED      --       AC   xe8       --- Single Homed Port ---     1001      -
---      ----             ----
20       VNI-BLUE     L2       NW   ----      ----                          ----      -
---      2.2.2.2          1.1.1.1
20       VNI-BLUE     L2       NW   ----      ----                          ----      -
---      2.2.2.2          3.3.3.3
20       VNI-BLUE     --       AC   xe8       --- Single Homed Port ---     2001      -
---      ----             ----

Total number of entries are 7


VTEP2#sh nvo vxlan access-if brief

                 Inner                     Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
xe3        ---   ---   500000   10        up      up
xe3        1001  ---   500001   10        up      up
xe3        2001  ---   500002   20        up      up

Total number of entries are 3


VTEP2#show bgp l2vpn evpn summary
BGP router identifier 2.2.2.2, local AS number 500
BGP table version is 4
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4   500   53        53       3     0     0  00:20:56
2      0      0      2      0      0
3.3.3.3                 4   500   51        54       3     0     0  00:20:56
2      0      0      2      0      0

Total number of neighbors 2

Total number of Established sessions 2


VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================
2.2.2.2          1.1.1.1          Installed     00:11:01      00:11:01
```

```
2.2.2.2          3.3.3.3          Installed     00:10:47      00:10:47

Total number of entries are 2


VTEP2#show bgp l2vpn evpn
BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1]
*>i  [3]:[10]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1        VXLAN
*>i  [3]:[20]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1        VXLAN

RD[2.2.2.2:1] VRF[VRF1]:
* i  [3]:[10]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1        VXLAN
*>   [3]:[10]:[32,2.2.2.2]
                      2.2.2.2              0        100       32768  i  ----------
VXLAN
* i  [3]:[10]:[32,3.3.3.3]
                     3.3.3.3             0        100       0    i  3.3.3.3        VXLAN
* i  [3]:[20]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1        VXLAN
*>   [3]:[20]:[32,2.2.2.2]
                      2.2.2.2              0        100       32768  i  ----------
VXLAN
* i  [3]:[20]:[32,3.3.3.3]
                     3.3.3.3             0        100       0    i  3.3.3.3        VXLAN

RD[3.3.3.3:1]
*>i  [3]:[10]:[32,3.3.3.3]
                     3.3.3.3             0        100       0    i  3.3.3.3        VXLAN
*>i  [3]:[20]:[32,3.3.3.3]
                     3.3.3.3             0        100       0    i  3.3.3.3        VXLAN

Total number of prefixes 10
```

VTEP3

```
VTEP3#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
10       VNI-RED      L2       NW   ----      ----                            ----      -
---      3.3.3.3          2.2.2.2
10       VNI-RED      L2       NW   ----      ----                            ----      -
---      3.3.3.3          1.1.1.1
10       VNI-RED      --       AC   ce13/1    --- Single Homed Port ---     0         -
---      ----             ----
10       VNI-RED      --       AC   ce13/1    --- Single Homed Port ---     1001      -
---      ----             ----
20       VNI-BLUE     L2       NW   ----      ----                            ----      -
---      3.3.3.3          2.2.2.2
20       VNI-BLUE     L2       NW   ----      ----                            ----      -
---      3.3.3.3          1.1.1.1
20       VNI-BLUE     --       AC   ce13/1    --- Single Homed Port ---     2001      -
---      ----             ----

Total number of entries are 7

VTEP3#sh nvo vxlan access-if brief
                 Inner                      Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid       status  status
-----------------------------------------------------------
xe48     ---   ---   500000   10         up      up
xe48     1001  ---   500001   10         up      up
xe48     2001  ---   500002   20         up      up

Total number of entries are 3

VTEP3#show bgp l2vpn evpn summary
BGP router identifier 3.3.3.3, local AS number 500
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4   500   52        52      3     0     0  00:20:37
2      0      0      2      0      0
2.2.2.2                 4   500   52        51      3     0     0  00:20:37
2      0      0      2      0      0

Total number of neighbors 2
```

```
Total number of Established sessions 2

VTEP3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================
3.3.3.3          2.2.2.2          Installed     00:10:50      00:10:50
3.3.3.3          1.1.1.1          Installed     00:10:50      00:10:50

Total number of entries are 2

VTEP3#show bgp l2vpn evpn
BGP table version is 3, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[1.1.1.1:1]
*>i  [3]:[10]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1        VXLAN
*>i  [3]:[20]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1        VXLAN

RD[2.2.2.2:1]
*>i  [3]:[10]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2        VXLAN
*>i  [3]:[20]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2        VXLAN

RD[3.3.3.3:1] VRF[VRF1]:
* i  [3]:[10]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1        VXLAN
* i  [3]:[10]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2        VXLAN
*>   [3]:[10]:[32,3.3.3.3]
                      3.3.3.3              0        100       32768  i  ----------
VXLAN
* i  [3]:[20]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1        VXLAN
```

```
* i   [3]:[20]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>    [3]:[20]:[32,3.3.3.3]
                       3.3.3.3               0        100       32768  i  ----------
VXLAN

Total number of prefixes 10
```

CHAPTER 6 VXLAN Quality of Service Configuration

This chapter contains the configurations for VXLAN Quality of Service (QoS) .

Overview

VXLAN enables multiple tenants to operate in a data center. Each tenant is assigned a priority group to prioritize their traffic. Cloud carriers want to use quality of service to differentiate different applications.

Data center networks are being increasingly used by telecommunications operators as well as by enterprises. Currently these networks are organized as one large Layer 2 network in a single building. In some cases such a network is extended geographically using Virtual Local Area Network (VLAN) technologies as an even larger Layer 2 network connecting the virtual machines (VM), each with its own MAC address.

Multiple tenants might want their own isolated network domain. In a data center hosting multiple tenants, each tenant may independently assign MAC addresses and VLAN IDs and this might lead to duplication.

Cloud carriers wish to categorize the traffic based on the application such as voice, video, etc. Based on the type of the application different traffic classes may be identified and different priority levels can be assigned to each. To do so, quality of service marking is needed in VXLAN.

This chapter shows how to mark packet headers with the VXLAN tunnel end point (VTEP) when the frames are introduced by the virtual machines. The (re)marking /setting of QoS field DSCP/TOS in the VXLAN IP header is done with the two modes which are set globally.

Topology

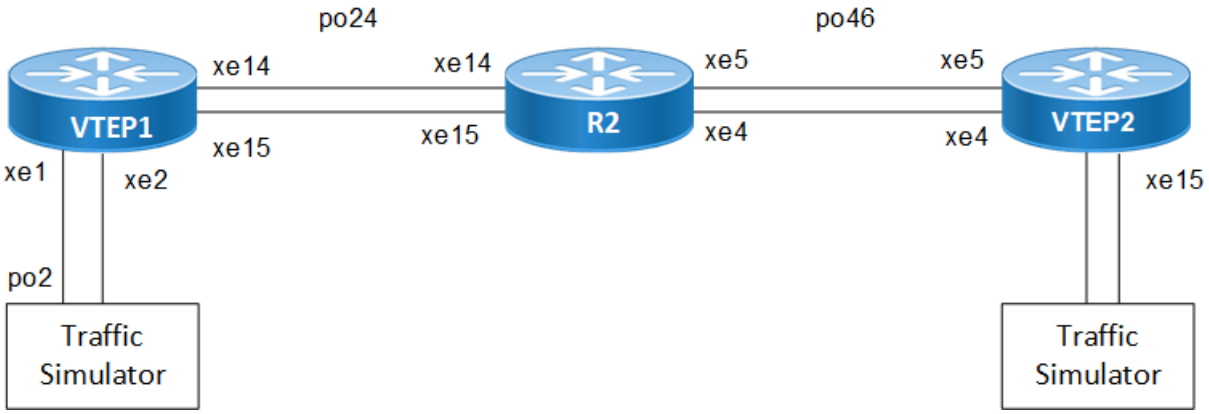

Figure 6-8: VXLAN QoS

COS-DSCP

RTR1/VTEP1

`#configure terminal`	Enter Configure mode.
`(config)#mac vrf vrf1`	Create mac routing/forwarding instance with vrf1 name and enter into vrf mode

(config-vrf)#rd 1.1.1.1:11	Assign RD value
(config-vrf)#route-target both 10.10.10.10:100	Assign route-target value for import/export
(config-vrf)#exit	Exit from vrf mode
(config)#qos enable	Enable qos
(config)#qos profile cos-to-queue COS-QUE	Create qos profile for mapping traffic towards tunnel from access-if.
(config-ingress-cos-map)#cos 2 queue 3	Configure particular COS value to the queue value for configured profile.
(config-ingress-cos-map)#exit	Exit from qos profile config mode
(config)#qos profile queue-color-to-dscp QUE-DSCP	Create qos profile for attaching in vxlan tunnel egress.
(config-egress-dscp-map)#queue 3 dscp 16	Configure particular queue value to the dscp value for configured profile.
(config-egress-dscp-encap-map)#exit	Exit from qos profile config mode
(config)#interface po2	Create a port channel po2
(config-if)#switchport	Configure port as switchport
(config-if)#load-interval 30	Set load-interval
(config-if)#interface po24	Create a port channel po24
(config-if)#load-interval 30	Configure port as switchport
(config-if)#ip address 24.1.1.1/30	Set load-interval
(config-if)#interface lo	Enter in to loopback interface
(config-if)#ip address 1.1.1.1/32 secondary	Configure ip address
(config-if)#interface xe1	Enter in to interface mode
(config-if)#channel-group 2 mode active	Map to channel-group
(config-if)#interface xe2	Enter in to interface mode
(config-if)#channel-group 2 mode active	Map to channel-group
(config-if)#interface xe14	Enter in to interface mode
(config-if)#channel-group 24 mode ac-tive	Map to channel-group
(config-if)#interface xe15	Enter in to interface mode
(config-if)#channel-group 24 mode ac-tive	Map to channel-group
(config-if)#router ospf 1	Create ospf instance
(config-router)#ospf router-id 1.1.1.1	Configure ospf router-id
(config-router)#network 1.1.1.1/32 area 0.0.0.0	Configure loopback network address in to ospf
(config-router)#network 24.1.1.0/30 area 0.0.0.0	Configure network address in to ospf
(config-router)#router bgp 100	Enter into Router BGP mode
(config-router)#neighbor 6.6.6.6 re-mote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 6.6.6.6 up-date-source lo	Specify the neighbor to use loopback address as source
(config-router)#address-family l2vpn evpn	Enter into l2vpn evpn address-family
(config-router-af)#neighbor 6.6.6.6 activate	Activate the neighbor to address-family

(config-router)#nvo vxlan vtep-ip-global 1.1.1.1	Configure Source vtep-ip-global configuration
(config)#nvo vxlan tunnel qos-map-mode cos-dscp egress QUE-DSCP	Configure the mapping qos profile in to vxlan tunnel egress
(config)#nvo vxlan id 1 in-gress-replication inner-vid-disabled	Create vnid 1 and disable inner-vid
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)#nvo vxlan access-if port-vlan po2 1001	Create vxlan access-if with vlan 1001
(config-nvo-acc-if)#no shutdown	No shut the vxlan access-if
(config-nvo-acc-if)#map vnid 1	Map vnid to the vxlan access-if
(config-nvo-acc-if)#map qos-profile cos-to-queue COS-QUE	Map qos profile for vxlan access-if ingress traffic from CE
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode.
(config)#commit	Commit the candidate configuration to the running configuration

R2

#configure terminal	Enter Configure mode.
(config)#interface po24	Create port channel
(config-if)#load-interval 30	Set load-interval
(config-if)#ip address 24.1.1.2/30	Assign ip address
(config-if)#interface po46	Create port channel
(config-if)#load-interval 30	Set load-interval
(config-if)#ip address 46.1.1.1/30	Assign ip address
(config-if)#interface lo	Enter in to loopback interface
(config-if)#ip address 4.4.4.4/32 sec-ondary	Assign secondary ip address
(config-if)#interface xe4	Enter into interface mode
(config-if)#channel-group 46 mode active	Map port channel to the interface
(config-if)#interface xe5	Enter into interface mode
(config-if)#channel-group 46 mode active	Map port channel to the interface
(config-if)#interface xe14	Enter into interface mode
(config-if)#channel-group 24 mode active	Map port channel to the interface
(config-if)#interface xe15	Enter into interface mode
(config-if)#channel-group 24 mode active	Map port channel to the interface
(config-if)#router ospf 1	Create ospf instance
(config-router)#ospf router-id 4.4.4.4	Configure ospf router-id
(config-router)#network 4.4.4.4/32 area 0.0.0.0	Configure ospf network address with respective area
(config-router)#network 24.1.1.0/30 area 0.0.0.0	Configure ospf network address with respective area

(config-router)#network 46.1.1.0/30 area 0.0.0.0	Configure ospf network address with respective area
(config-router)#exit	Exit from router mode.
(config)#commit	Commit the candidate configuration to the running configuration

RTR3/VTEP2

#configure terminal	Enter Configuration mode
(config)#mac vrf vrf1	Create mac routing/forwarding instance with vrf1 name and enter into vrf mode
(config-vrf)#rd 6.6.6.6:11	Assign RD value
(config-vrf)#route-target both 10.10.10.10:100	Assign route-target value for import/export
(config-vrf)#exit	Exit from vrf mode
(config)#qos enable	Enable QOS
(config)#qos statistics	Enable QOS statistics
(config)#qos profile queue-color-to-cos QUE-COS	Create qos profile for mapping incoming traffic from tunnel to access-if.
(config-egress-cos-map)#queue 4 cos 5	Configure particular queue value to the cos value for configured profile.
(config-egress-cos-map)#qos profile dscp-to-queue DSCP-QUE	Create qos profile for attaching in vxlan tunnel ingress.
(config-ingress-dscp-map)#dscp 16 queue 4	Configure particular dscp value to the queue value for configured profile.
(config-egress-dscp-map)#interface po46	Create port channel
(config-if)#load-interval 30	Set load interval
(config-if)#ip address 46.1.1.2/30	Assign ip address
(config-if)#interface lo	Enter into loopback interface
(config-if)#ip address 6.6.6.6/32 secondary	Assign secondary ip address
(config-if)#interface xe4	Enter into interface mode
(config-if)#channel-group 46 mode active	Map channel group into the interface
(config-if)#interface xe5	Enter into interface mode
(config-if)#channel-group 46 mode active	Map channel group into the interface
(config-if)#interface xe15	Enter into interface mode
(config-if)#switchport	Make interface as L2 port
(config-if)#load-interval 30	Set load interval
(config-if)#router ospf 1	Create ospf instance
(config-router)#ospf router-id 6.6.6.6	Configure ospf router-id
(config-router)#network 6.6.6.6/32 area 0.0.0.0	Configure ospf network address with respective area
(config-router)#network 46.1.1.0/30 area 0.0.0.0	Configure ospf network address with respective area
(config-router)#router bgp 100	Enter into Router BGP mode

(config-router)#neighbor 1.1.1.1 re-mote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 1.1.1.1 up-date-source lo	Specify the neighbor to use loopback address as source
(config-router)#address-family l2vpn evpn	Enter into l2vpn evpn address-family
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor to address-family
(config)#nvo vxlan vtep-ip-global 6.6.6.6	Configure Source vtep-ip-global configuration
(config)#nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP-QUE	Configure the mapping qos profile in to vxlan tunnel ingress
(config)#nvo vxlan id 1 in-gress-replication inner-vid-disabled	Create vnid 1 and disable inner-vid
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)#nvo vxlan access-if port-vlan xe15 1000	Create vxlan access-if with vlan 1000
(config-nvo-acc-if)#no shutdown	No shut the vxlan access-if
(config-nvo-acc-if)#map vnid 1	Map vnid to the vxlan access-if
(config-nvo-acc-if)#map qos-profile queue-color-to-cos QUE-COS	Map qos profile for vxlan access-if egress traffic to CE
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#commit	Commit the candidate configuration to the running configuration

Validation

As per the QoS configuration, when L2 traffic with cos value 2 sent to VTEP1 access-if, the packets forwarded to queue 3 and packets in queue 3 are mapped with overlay dscp value 16 while egress out of tunnel. At VTEP2, when packets with overlay dscp value 16 ingresses at tunnel, it is forwarded to queue 4 and packets of queue 4 are remarked with cos value 5.

RTR1/VTEP1

```
VTEP1#sh run nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan tunnel qos-map-mode cos-dscp egress QUE-DSCP
!
nvo vxlan id 1 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrf1
!
nvo vxlan access-if port-vlan po2 1001
 map vnid 1
 map qos-profile cos-to-queue COS-QUE
!
VTEP1#show run qos
qos enable
```

```
!
qos profile cos-to-queue COS-QUE
 cos 2 dei all queue 3
!
qos profile queue-color-to-dscp QUE-DSCP
 queue 3 color all dscp 16
!

VTEP1#show interface xe14 count queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes      |  Dropped pkts   |
Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
q0                 (E) 12517376 0                 0                   0                 0
q1                 (E) 12517376 0                 0                   0                 0
q2                 (E) 12517376 0                 0                   0                 0
q3                 (E) 12517376 205284588         188040683524        0                 0
q4                 (E) 12517376 0                 0                   0                 0
q5                 (E) 12517376 0                 0                   0                 0
q6                 (E) 12517376 0                 0                   0                 0
q7                 (E) 12517376 7518              1007412             0                 0

VTEP1#show interface xe15 count queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes      |  Dropped pkts   |
Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
q0                 (E) 12517376 0                 0                   0                 0
q1                 (E) 12517376 0                 0                   0                 0
q2                 (E) 12517376 0                 0                   0                 0
q3                 (E) 12517376 205624494         188352040168        0                 0
q4                 (E) 12517376 0                 0                   0                 0
q5                 (E) 12517376 0                 0                   0                 0
q6                 (E) 12517376 0                 0                   0                 0
q7                 (E) 12517376 9006              1136741             0                 0

VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
           AC - Access Port
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                           VLAN DF-Status
Src-Addr         Dst-Addr
```

```
_____________________________________________________________________________________________
_____________________________________
1        ----         L2       NW   ----      ----                          ---- ----
1.1.1.1          6.6.6.6
1       ----        AC   po2      --- Single Homed port ---    1001 ----     ----
----

Total number of entries are 3

VTEP1#show nvo vxlan mac-table
=============================================================================================
============================================================
                                                                VXLAN MAC Entries
=============================================================================================
============================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc
_____________________________________________________________________________________________
_____________________________________________________________

1         po2        1001   ----         0000.2000.9991 1.1.1.1
Dynamic Local        -------             -------

Total number of entries are : 1

VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
=============================================================================
1.1.1.1          6.6.6.6          Installed        00:11:29        00:11:29

Total number of entries are 2

VTEP1#show qos-profile COS-QUE
profile name: COS-QUE
profile type: cos-to-queue
profile attached to 1 instances
configured mapping:
 cos 2 dei all queue 3
Detailed mapping:
 ---------------+----------------- | ---------------+-----------------
      INPUT     |     OUTPUT       |      INPUT     |     OUTPUT
 ---------------+----------------- | ---------------+-----------------
   COS  |  DEI  | Queue |  Color   |   COS  |  DEI  | Queue |  Color
 -------+-------+-------+--------- | -------+-------+-------+---------
  0       0       0       green    |  0       1       0       yellow
  1       0       1       green    |  1       1       1       yellow
  2       0       3       green    |  2       1       3       yellow
  3       0       3       green    |  3       1       3       yellow
  4       0       4       green    |  4       1       4       yellow
  5       0       5       green    |  5       1       5       yellow
  6       0       6       green    |  6       1       6       yellow
```

```
  7       0       7       green    |  7       1       7       yellow

VTEP1#show qos-profile QUE-DSCP
profile name: QUE-DSCP
profile type: queue-color-to-dscp
profile attached to 1 instances
configured mapping:
 queue 3 color all dscp 16
Detailed mapping:
 ----------------+-------- | ----------------+-------- | ----------------+--------
      INPUT      | OUTPUT  |      INPUT      | OUTPUT  |      INPUT      | OUTPUT
 ----------------+-------- | ----------------+-------- | ----------------+--------
  Queue | Color  | DSCP    |  Queue | Color  | DSCP    |  Queue | Color  | DSCP
 -------+--------+-------- | -------+--------+-------- | -------+--------+--------
  0       green    0       |  0       yellow   0       |  0       red      0
  1       green    10      |  1       yellow   12      |  1       red      14
  2       green    18      |  2       yellow   20      |  2       red      22
  3       green    16      |  3       yellow   16      |  3       red      16
  4       green    34      |  4       yellow   36      |  4       red      38
  5       green    40      |  5       yellow   40      |  5       red      40
  6       green    48      |  6       yellow   48      |  6       red      48
  7       green    56      |  7       yellow   56      |  7       red      56
```

RTR2/VTEP2

```
VTEP2#show run nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 6.6.6.6
!
nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP-QUE
!
nvo vxlan id 1 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrf1
!
nvo vxlan access-if port-vlan xe15 1000
 map vnid 1
 map qos-profile queue-color-to-cos QUE-COS
!
!
VTEP2#show run qos
qos enable
qos statistics
!
qos profile queue-color-to-cos QUE-COS
 queue 4 color all cos 5
!
qos profile dscp-to-queue DSCP-QUE
```

```
 dscp 16 queue 4
!
VTEP2#show nvo vxlan mac-table

================================================================================================
============================================================
                                                              VXLAN MAC Entries
================================================================================================
============================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc
------------------------------------------------------------------------------------------------
------------------------------------------------------------

1         ----       ----   ----         0000.2000.9991 1.1.1.1
Dynamic Remote       -------             -------

Total number of entries are : 1

VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
6.6.6.6          1.1.1.1          Installed       00:09:39        00:09:39

Total number of entries are 2
VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status
Src-Addr          Dst-Addr
------------------------------------------------------------------------------------------------
----------------------------------------
1        ----         L2       NW   ----      ----                           ---- ----
6.6.6.6          1.1.1.1
1       ----        AC   xe15     --- Single Homed port ---    1000 ----     ----
----

Total number of entries are 3

VTEP2#show interface xe15 count queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+------------------+----------------
+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes     |  Dropped pkts   |
Dropped bytes   |
+--------------------+--------+-----------------+------------------+----------------
+-------------------+
q0                (E) 12517376 0                 0                  0                 0
```

```
q1                (E) 12517376 0                 0                  0                  0
q2                (E) 12517376 0                 0                  0                  0
q3                (E) 12517376 0                 0                  0                  0
q4                (E) 12517376 37895872          36455829826        0                  0
q5                (E) 12517376 0                 0                  0                  0
q6                (E) 12517376 0                 0                  0                  0
q7                (E) 12517376 0                 0                  0                  0

VTEP2#show qos-profile QUE-COS
profile name: QUE-COS
profile type: queue-color-to-cos
profile attached to 1 instances
configured mapping:
 queue 4 color all cos 5
Detailed mapping:
 ----------------+-------- | ----------------+-------- | ----------------+--------
      INPUT      | OUTPUT  |      INPUT      | OUTPUT  |      INPUT      | OUTPUT
 ----------------+-------- | ----------------+-------- | ----------------+--------
  Queue | Color  | COS     |  Queue | Color  | COS     |  Queue | Color  | COS
 -------+--------+-------- | -------+--------+-------- | -------+--------+--------
  0       green    0       |  0       yellow   0       |  0       red      0
  1       green    1       |  1       yellow   1       |  1       red      1
  2       green    2       |  2       yellow   2       |  2       red      2
  3       green    3       |  3       yellow   3       |  3       red      3
  4       green    5       |  4       yellow   5       |  4       red      5
  5       green    5       |  5       yellow   5       |  5       red      5
  6       green    6       |  6       yellow   6       |  6       red      6
  7       green    7       |  7       yellow   7       |  7       red      7

VTEP2#show qos-profile DSCP-QUE
profile name: DSCP-QUE
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
 dscp 16 queue 4
Detailed mapping:
 Ingress map for dscp to queue-color
 -------+----------------- | -------+----------------- | -------+----------------- | --
-----+-----------------
  INPUT |     OUTPUT       |  INPUT |     OUTPUT       |  INPUT |     OUTPUT       |
INPUT |     OUTPUT
 -------+----------------- | -------+----------------- | -------+----------------- | --
-----+-----------------
  DSCP  | Queue | Color    |  DSCP  | Queue | Color    |  DSCP  | Queue | Color    |  DSCP
| Queue | Color
 -------+-------+--------- | -------+-------+--------- | -------+-------+--------- | --
-----+-------+---------
  0       0       green    |  16      4       green    |  32      4       green    |  48
6       green
  1       0       green    |  17      2       green    |  33      4       green    |  49
6       green
```

```
  2       0       green    |  18      2       green    |  34      4       green    |  50
6       green
  3       0       green    |  19      2       green    |  35      4       green    |  51
6       green
  4       0       green    |  20      2       yellow   |  36      4       yellow   |  52
6       green
  5       0       green    |  21      2       green    |  37      4       green    |  53
6       green
  6       0       green    |  22      2       red      |  38      4       red      |  54
6       green
  7       0       green    |  23      2       green    |  39      4       green    |  55
6       green
  8       1       green    |  24      3       green    |  40      5       green    |  56
7       green
  9       1       green    |  25      3       green    |  41      5       green    |  57
7       green
  10      1       green    |  26      3       green    |  42      5       green    |  58
7       green
  11      1       green    |  27      3       green    |  43      5       green    |  59
7       green
  12      1       yellow   |  28      3       yellow   |  44      5       green    |  60
7       green
  13      1       green    |  29      3       green    |  45      5       green    |  61
7       green
  14      1       red      |  30      3       red      |  46      5       green    |  62
7       green
  15      1       green    |  31      3       green    |  47      5       green    |  63
7       green
```

CHAPTER 7 MAC LIMIT ENFORCEMENT AT EVI AND AC FOR EVPN-VxLAN

This chapter includes step-by-step configurations for EVPN MAC Limit Enforcement at EVI and AC for EVPN-VxLAN.

Overview

The MAC Limit feature in EVPN-VxLAN provides a mechanism to control the number of MAC addresses learned at both the Ethernet Virtual Private Network Instance (EVI) level and the Attachment Circuit (AC) level. This feature enhances network efficiency and security by preventing MAC table overflow and mitigating potential denial-of- service attacks in EVPN deployments using VxLAN as the overlay, supporting Leaf and Spine CLOS fabric.

This document outlines the topology, configuration steps, and Command Line Interface (CLI) details for implementing MAC limit enforcement in EVPN-VxLAN environments at the EVI and AC levels.

Feature Characteristics

- MAC Limit Scope: Limits can be enforced at the EVI level (aggregate MACs across all ACs, BGP-learned MACs, and static MACs) or at the individual AC level (MACs learned from SERVER to LEAF).
- Actions on Limit Breach: Supports logging and error-disabling for AC-level limits; only logging is supported for EVI-level limits.
- Threshold Watermarks: High and low watermark thresholds trigger syslog messages when MAC counts exceed or fall below configured percentages (default high watermark: 90%).
- Error Disable Recovery: Configurable recovery timer to automatically re-enable error-disabled ACs after a specified period.
- Overlay Specific: The feature is tailored for VxLAN overlay in EVPN deployments.
- Software-Based: When the MAC limit is reached, further MAC learning is stopped, but flooding continues.

Benefits

- Prevents MAC Table Overflow: Enhances system stability in EVPN-VxLAN environments.
- Granular Control: Allows MAC learning limits at both EVI and AC levels.
- Monitoring Support: Provides logging for tracking and troubleshooting.
- Automatic Recovery: Enables configurable timers for recovery from error-disabled states.
- Enhanced Security: Limits MAC address learning to improve security in multi-tenant EVPN-VxLAN setups.

Prerequisites

- Configure EVPN with VxLAN as the overlay protocol.
- Interfaces should be configured as switchports with VLAN encapsulation.

Topology

The EVPN MAC Limit feature can be deployed in any standard EVPN-VxLAN topology with SERVER-LEAF connectivity.

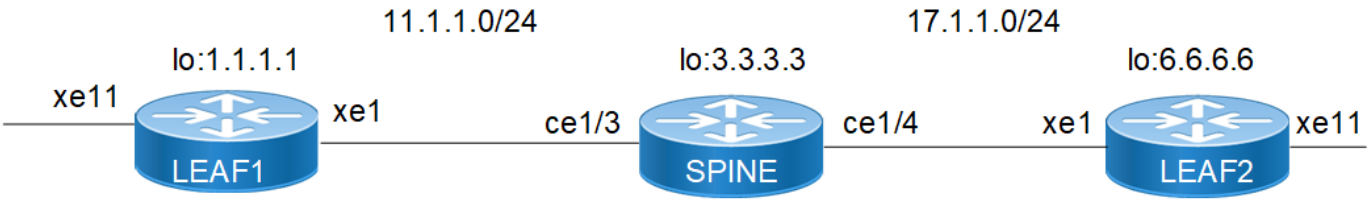

Figure 7-9: VXLAN EVPN MAC limit enforcement

Configuration

The following configuration steps demonstrate how to enable and configure MAC limit enforcement on LEAF1 for EVPN-VxLAN.

EVPN prerequisite configurations:

LEAF1

1. Enable VxLAN to allow configuration of overlay services.

```
nvo vxlan enable
```

2. Configure global VTEP IP address.

```
nvo vxlan vtep-ip-global 1.1.1.1
```

3. Configure MAC VRF.

```
mac vrf vxlan_l2_elan_sh
description vxlan_l2_elan_sh
rd 1.1.1.1:100
route-target both 100:100
```

4. Configure VxLAN instance network identifier.

```
nvo vxlan id 50 ingress-replication inner-vid-disabled
vxlan host-reachability-protocol evpn-bgp vxlan_l2_elan_sh
```

5. Map each port-VLAN sub-interface to the VxLAN network identifier.

```
nvo vxlan access-if port-vlan xe11 50
map vnid 50
nvo vxlan access-if port-vlan xe11 51
map vnid 50
```

```
nvo vxlan access-if port-vlan xe11 52
map vnid 50
```

LEAF2

1. Enable VxLAN to allow configuration of overlay services.

```
nvo vxlan enable
```

2. Configure global VTEP IP address.

```
nvo vxlan vtep-ip-global 6.6.6.6
```

3. Configure MAC VRF.

```
mac vrf vxlan_l2_elan_sh2
description vrf vxlan_l2_elan_sh2
rd 6.6.6.6:100
route-target both 100:100
```

4. Configure VxLAN instance network identifier.

```
nvo vxlan id 50 ingress-replication inner-vid-disabled
vxlan host-reachability-protocol evpn-bgp vxlan_l2_elan_sh2
```

5. Map each port-VLAN sub-interface to the VxLAN network identifier.

```
nvo vxlan access-if port-vlan xe11 50
map vnid 50
```

EVPN-MAC-Limit profile configuration:

LEAF1

Configure mac-limit profile with learning-limit and with default values of high water-mark, low water-mark and action and map it to AC.

```
mac-limit-profile SH1
learning-limit 10
nvo vxlan access-if port-vlan xe11 50
map vnid 50
learning limit SH1
```

Configuration snapshot:

LEAF1

```
hostname LEAF1
!
nvo vxlan enable
!
evpn vxlan multihoming enable
!
```

```
mac vrf vxlan_l2_elan_sh
 description vxlan_l2_elan_sh
 rd 1.1.1.1:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan id 50 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vxlan_l2_elan_sh
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
!
interface xe1
 load-interval 30
 ip address 11.1.1.1/24
!
interface xe11
 switchport
 load-interval 30
!
router ospf 1
 ospf router-id 1.1.1.1
 network 1.1.1.1/32 area 0.0.0.0
 network 11.1.1.0/24 area 0.0.0.0
!
router bgp 1
 bgp router-id 1.1.1.1
 neighbor 6.6.6.6 remote-as 1
 neighbor 6.6.6.6 update-source lo
 !
 address-family l2vpn evpn
 neighbor 6.6.6.6 activate
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan xe11 50
 map vnid 50
 learning limit SH1
!
nvo vxlan access-if port-vlan xe11 51
 map vnid 50
!
nvo vxlan access-if port-vlan xe11 52
 map vnid 50
!
```

SPINE:

```
hostname SPINE
!
interface ce1/3
 load-interval 30
 ip address 11.1.1.2/24
!
interface ce1/4
 load-interval 30
 ip address 17.1.1.2/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 3.3.3.3/24
 ipv6 address ::1/128
!
router ospf 1
 ospf router-id 3.3.3.3
 network 3.3.3.3/32 area 0.0.0.0
 network 11.1.1.0/24 area 0.0.0.0
 network 17.1.1.0/24 area 0.0.0.0
!
```

LEAF2

```
hostname LEAF2
!
nvo vxlan enable
!
mac vrf vxlan_l2_elan_sh2
 description vrf vxlan_l2_elan_sh2
 rd 6.6.6.6:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 6.6.6.6
!
nvo vxlan id 50 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vxlan_l2_elan_sh2
!
interface lo
 ip address 127.0.0.1/8
 ip address 6.6.6.6/32 secondary
 ipv6 address ::1/128
!
interface xe1
 load-interval 30
 ip address 17.1.1.1/24
!
interface xe11
 switchport
```

```
 load-interval 30
!
router ospf 1
 ospf router-id 6.6.6.6
 network 6.6.6.6/32 area 0.0.0.0
 network 17.1.1.0/24 area 0.0.0.0
!
router bgp 1
 bgp router-id 6.6.6.6
 neighbor 1.1.1.1 remote-as 1
 neighbor 1.1.1.1 update-source lo
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan xe11 50
 map vnid 50
!
```

Validation

To verify MAC limit enforcement, send traffic with varying numbers of source MAC addresses and observe the system behavior.

```
LEAF1#show mac-limit profiles
Profile-Name      Limit    Action    High-Watermark(%, v)  Low-Watermark(%, v)
Errdisable-timeout
================================================================================
================================
SH1                10      log-only         90, 9                70, 7
0
LEAF1#
```

- Send Traffic with 10 Source MACs to LEAF1.
- Generate traffic with 10 unique source MACs on xe11.
- When the MAC count reaches the high watermark threshold (default 90% of 10 = 9 MACs), a syslog message is generated.

```
LEAF1#2025 Jun 14 04:45:19.450 : LEAF1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50:
Mac limit for AC xe11, High threshold MAC count 9 with high watermark of 9
```

- Send Traffic with 20 Source MACs.
- Generate traffic with 20 unique source MACs, exceeding the learning limit of 10.

```
LEAF1#2025 Jun 14 04:46:30.610 : LEAF1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50:
Mac limit for AC xe11, Exceeded MAC count 11 with learning limit of 10
```

- Now stop the traffic and check for Low water logs during un learning.

```
2025 Jun 14 04:53:52.280 : LEAF1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac
limit for AC xe11, Low threshold MAC count 6 with low watermark of 7
```

Now modifying the mac-limit profile with non default high water-mark, low water-mark and action as error disable:

```
mac-limit-profile SH1
learning-limit 10
action log-errdisable
high-watermark 80
low-watermark 70
nvo vxlan access-if port-vlan xe11 50
 map vnid 50
 learning limit SH1

LEAF1#show mac-limit profiles
Profile-Name       Limit     Action         High-Watermark(%, v)  Low-Watermark(%, v)
Errdisable-timeout
================================================================================================
================================
SH1                10       log-errdisable         80, 8              70, 7
0
LEAF1#
```

- Send Traffic with 10 Source MACs to LEAF1.
- Generate traffic with 10 unique source MACs on xe11.
- When the MAC count reaches the high watermark threshold (default 90% of 10 = 9 MACs), a syslog message is generated.

```
2025 Jun 14 04:55:22.638 : LEAF1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac
limit for AC xe11, High threshold MAC count 8 with high watermark of 8
```

- Send Traffic with 20 Source MACs.
- Generate traffic with 20 unique source MACs, exceeding the learning limit of 10 and check AC is going to error disable.

```
2025 Jun 14 04:55:22.639 : LEAF1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac
limit for AC xe11, Exceeded MAC count 11 with learning limit of 10
2025 Jun 14 04:55:22.639 : LEAF1 : NSM : CRITI : [IFMGR_ERR_DISABLE_DOWN_2]: Attachment
Circuit with the nvo access-interface xe11 50 on EVPN instance 50 errdisabled
successfully due to EVPN-MAC-LIMIT. Configured error disable timeout 0
2025 Jun 14 04:55:22.652 : LEAF1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac
limit for AC xe11, Low threshold MAC count 6 with low watermark of 7
LEAF1#show nvo vxlan access-if brief

                     Inner                        Admin   Link
Interface  Vlan      vlan  Ifindex    Vnid       status  status
-------------------------------------------------------------
xe11       50        ---   0x7a120    50         up      down(ED)
xe11       51        ---   0x7a122    50         up      up
xe11       52        ---   0x7a123    50         up      up
po100      100       ---   0x7a121    100        up      up

Total number of entries are 4
LEAF1#

LEAF1#show interface brief xe11

Codes: ETH - Ethernet, LB - Loopback, AGG - Aggregate, MLAG - MLAG Aggregate
       FR - Frame Relay, TUN -Tunnel, PBB - PBB Logical Port, VP - Virtual Port
       CVP - Channelised Virtual Port, METH - Management Ethernet, UNK- Unknown
```

```
       ED - ErrDisabled, PD - Protocol Down, AD - Admin Down, IA - InActive
       PD(Min L/B) - Protocol Down Min-Links/Bandwidth
       OTD  - Object Tracking Down
       DV - DDM Violation, NA - Not Applicable
       NOM - No operational members, PVID - Port Vlan-id
       Ctl - Control Port (Br-Breakout/Bu-Bundle)
       HD  - ESI Hold Timer Down
-------------------------------------------------------------------------------------
------------------
Ethernet          Type      PVID  Mode     Status   Reason    Speed   Port    Ctl Br/Bu
Loopbk
Interface     Ch #
-------------------------------------------------------------------------------------
------------------
xe11                ETH      --    --         up     none     10g     --       No
No
LEAF1#
```

Note: Action as log-only, is applicable for both AC (attachment circuit) and vnid.

Note: Action as log-errdisable, is applicable only for AC (attachment circuit).

CLI Commands

This commands section describes the mac limit enforcement at EVI and AC for EVPN-VxLAN.

mac-limit-profile

Use this command to create a MAC limit profile to enforce Layer 2 MAC limits. Use the no form to delete the profile.

Command Syntax

```
mac-limit-profile <PROFILE-NAME>
[no] mac-limit-profile <PROFILE-NAME>
```

Parameters

`<PROFILE NAME>` Name of the MAC limit profile

Default

None

Command Mode

Configure mode

Applicability

This command is introduced in OcNOS/Version/6.6.1.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#learning-limit 100
OcNOS(config-mac-limit-profile)#high-watermark 80
OcNOS(config-mac-limit-profile)#low-watermark 40
OcNOS(config-mac-limit-profile)#action log-errdisable
OcNOS(config-mac-limit-profile)#errdisable-timeout 60
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#no mac-limit-profile profile_1
OcNOS(config)#commit
```

learning-limit

Use this command to set the learning limit. Use the no form to reset to the default limit

Command Syntax

```
learning-limit <1-131071>
[no] learning-limit
```

Parameters

<1-131071>	Maximum number of MAC addresses (1 to 131071).

Default

131071

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS/Version/6.6.1.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#learning-limit 100
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no learning-limit
OcNOS(config-mac-limit-profile)#commit
```

action

Use this command to set the action type after the MAC limit is reached. If set to log-only, only logs are generated. If set to log-errdisable, both logs and error-disable actions take effect. Use the no form to reset to the default action.

Command Syntax

```
action (log-only | log-errdisable)
[no] action
```

Parameters

`log-only`	Maximum number of MAC addresses (1 to 131071)
`log-errdisable`	Generates logs and error-disables the interface

Default

log-only

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS/Version/6.6.1.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#action log-errdisable
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no action
OcNOS(config-mac-limit-profile)#commit
```

high-watermark

Use this command to set the high-watermark percentage before the MAC limit is reached. Once the high watermark is reached, an operator log is generated. Use the no form to reset to the default.

Command Syntax

```
high-watermark <1-100>
[no] high-watermark
```

Parameters

`<1-100>`	Percentage of the MAC limit (1 to 100)

Default

90

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS/Version/6.6.1.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#high-watermark 80
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no high-watermark
OcNOS(config-mac-limit-profile)#commit
```

low-watermark

Use this command to set the low watermark percentage. Once the low watermark is reached, an operator log is generated. Use the no form to reset to the default.

Command Syntax

```
low-watermark <1-100>
[no] low-watermark
```

Parameters

`<1-100>`	Percentage of the MAC limit (1 to 100)

Default

70

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS/Version/6.6.1.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#low-watermark 40
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no low-watermark
OcNOS(config-mac-limit-profile)#commit
```

errdisable-timeout

Use this command to set the error-disable timeout value in seconds. Once error-disable occurs, this is the duration for which the instance will be operationally shut before being restored. Use the no form to reset to the default.

Command Syntax

```
errdisable-timeout <0-86400>
[no] errdisable-timeout
```

Parameters

<0-86400>	Timeout duration in seconds (0 to 86400). A value of 0 disables automatic recovery.

Default

0

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS/Version/6.6.1.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#low-watermark 40
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no low-watermark
OcNOS(config-mac-limit-profile)#commit
```

learning limit

Use this command to associate a MAC limit profile with an EVPN-VxLAN instance or access interface. Use the no form to disassociate the profile.

Command Syntax

```
learning limit <PROFILE-NAME>
[no] learning limit
```

Parameters

<PROFILE-NAME>	Name of the MAC limit profile to apply

Default

None

Command Mode

EVPN-VXLAN-MODE, ACC-IF-EVPN-MODE

Applicability

This command is introduced in OcNOS/Version/6.6.1.

Example

```
Configure an EVPN-VXLAN EVI with a mac-limit-profile
LEAF1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
LEAF1(config)#mac-limit-profile SH1
LEAF1(config-mac-limit-profile)#learning-limit 10
LEAF1(config-mac-limit-profile)#commit
LEAF1(config-mac-limit-profile)#

LEAF1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
LEAF1(config)#
LEAF1(config)#nvo vxlan id 50 ingress-replication
LEAF1(config-nvo)#learning limit SH1
LEAF1(config-nvo)#commit
LEAF1(config-nvo)#end
LEAF1#
LEAF1#
LEAF1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
LEAF1(config)#
LEAF1(config)#nvo vxlan id 50 ingress-replication
LEAF1(config-nvo)#no learning limit
LEAF1(config-nvo)#commit
LEAF1(config-nvo)#end
LEAF1#

Configure an EVPN-VxLAN access interface with a mac-limit-profile
LEAF1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
LEAF1(config)#mac-limit-profile SH1
LEAF1(config-mac-limit-profile)#learning-limit 10
LEAF1(config-mac-limit-profile)#action log-errdisable
LEAF1(config-mac-limit-profile)#high-watermark 80
LEAF1(config-mac-limit-profile)#low-watermark 70
LEAF1(config-mac-limit-profile)#
LEAF1(config-mac-limit-profile)#commit

LEAF1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
LEAF1(config)#nvo vxlan access-if port-vlan xe11 50
LEAF1(config-nvo-acc-if)#learning limit SH1
LEAF1(config-nvo-acc-if)#commit

LEAF1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
LEAF1(config)#nvo vxlan access-if port-vlan xe11 50
LEAF1(config-nvo-acc-if)#no learning limit
LEAF1(config-nvo-acc-if)#commit
```

CHAPTER 8 EVPN L3 Gateway with VXLAN Stitching

Overview

Layer 3 EVPN Gateway Stitching in OcNOS enables seamless IP connectivity between isolated VXLAN EVPN domains or VRFs without merging their control planes or Layer 2 broadcast domains. This approach is designed for multi-site data center interconnect, multi-PoD fabrics, and hybrid cloud deployments, where scalable and secure routing between domains is required.

EVPN stitching is supported on TD3 and TD4 data center platforms using EVPN Type-5 routes at border leaf or spine devices. The solution provides route translation, loop prevention, policy-based route control, and operational visibility, ensuring flexible and modular data center expansion.

Feature Characteristics

- Maintains independent EVPN control planes across domains while enabling IP-level connectivity.
- Uses EVPN Type-5 route stitching at border leaf/spine nodes for prefix exchange.
- Supports multiple domains, Multi-PoD fabrics, and inter-DC connectivity without merging VNIs or broadcast domains.

Benefits

- Provides IP-level communication between isolated VXLAN EVPN domains without merging their control planes.
- Enables multi-site, multi-PoD, and hybrid cloud deployments with independent scaling per domain.
- Route-maps and policy tools allow selective advertisement and filtering of prefixes.

Prerequisites

Below CLIs need to be configured, saved, and the node should be rebooted before proceeding to configure the EVPN L3 gateway VXLAN stitching commands.

```
OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#nvo vxlan enable
OcNOS(config)#nvo vxlan irb
OcNOS(config)#commit
OcNOS(config)#end
OcNOS#write
Building configuration...
[OK]
OcNOS#reload
Are you sure you would like to reset the system? (y/n): y
Device rebooted by ocnos user from cmlsh at time:Tue Oct  7 18:33:51 2025
OcNOS#
```

Limitation

- **Type-5 Only** – Only EVPN Type-5 route stitching is supported. Other route types (e.g., Type-2) must be converted to Type-5 before stitching.
- **Static VXLAN Tunnels** – Not supported; stitching is limited to dynamic EVPN-based VXLAN.
- **Management Support** – SNMP and OpenConfig are not supported in this release; only CLI and Netconf are available.
- **Platform Scope** – Supported only on TD3 and TD4 data center platforms.
- **Scale** – Supported up to four border gateways.
- **ECMP overlay over L3 Gateway** - Not supported
- **peer-type fabric-external option** - Not supported under BGP l2vpn-unnumbered-mode for the l2vpn evpn address family.
- Advanced QoS, ACLs, or telemetry extensions across stitched domains are not included.

Configuration

Topology

The diagram depicts a multi-site data center fabric topology with Spine–Leaf architecture and Super-Spine layers, showing inter-site connectivity between Site A and Site B using VxLAN/EVPN with L3/L2 stitching.

Figure 8-10 displays a sample VxLAN/EVPN with L3/L2 stitching topology.

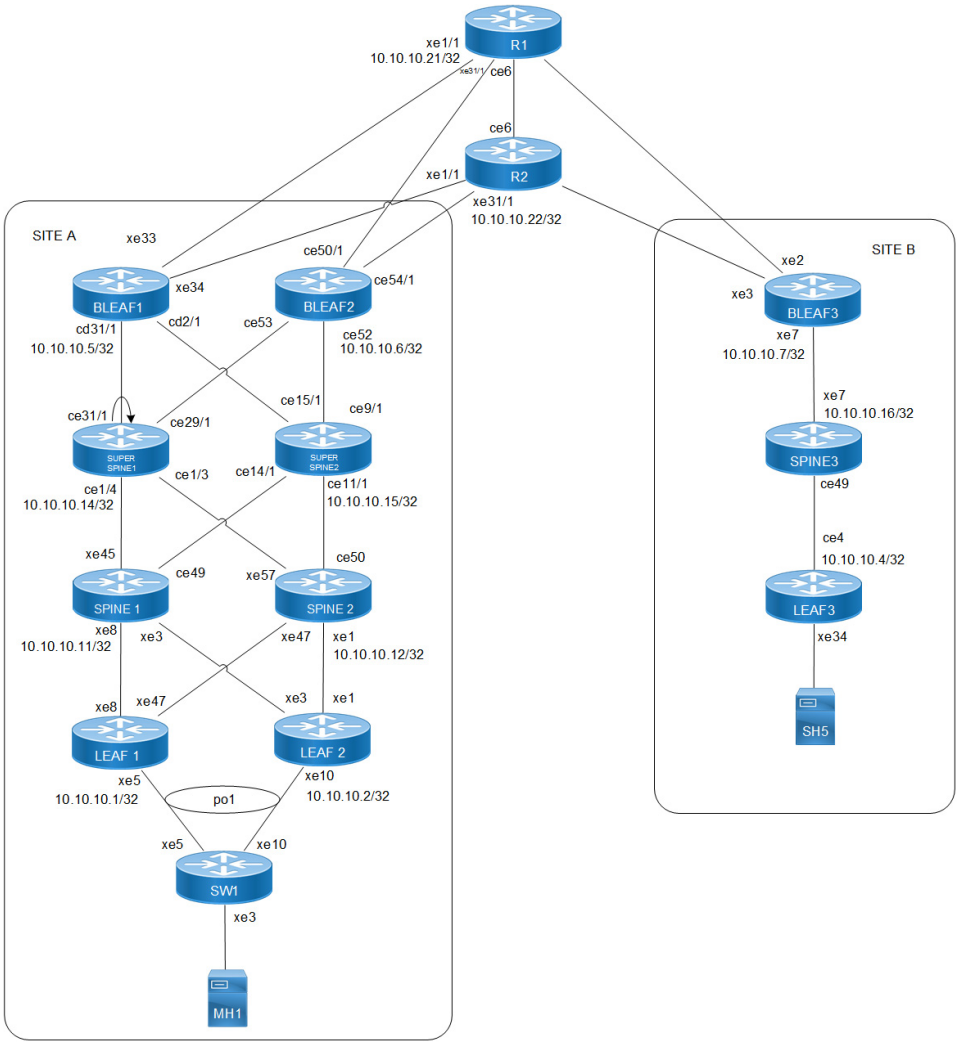

Figure 8-10: EVPN L3 Gateway with VXLAN Stitching topology

BLEAF1 Configuration

1. Enable VXLAN and IRB.

```
#configure terminal
(config)#nvo vxlan enable
(config)#nvo vxlan irb
```

2. Enable EVPN VXLAN multihoming.

```
(config)#evpn vxlan multihoming enable
```

3. Configure L3 VRF with EVPN stitching and enable EVPN stitching.

```
(config)#ip vrf l3_vrf_1
(config-vrf)#rd 7005:1001
(config-vrf)#route-target both 1001:1001
(config-vrf)#l3vni 1000
(config-vrf)#evpn-stitching
(config-vrf-evpn)#rd 7005:7001
(config-vrf-evpn)#route-target both 7001:7001
(config-vrf-evpn)#host-route-extend-l2vni 1001-1002
```

```
(config-vrf-evpn)#translation-l3vni 2000
(config-vrf-evpn)#exit-config-vrf-evpn-mode
(config-vrf)#commit
(config-vrf)#exit
```

4. Configure global VTEP IP.

```
(config)#nvo vxlan vtep-ip-global 10.10.10.5
```

5. Define route-maps for EVPN control.

```
(config)#route-map R1 permit 1
(config-route-map)#set extcommunity soo 1004:1004
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R2 deny 1
(config-route-map)#match extcommunity S1
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R2 permit 2
(config-route-map)#match evpn-route-type type-5
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R2 permit 3
(config-route-map)#match evpn-route-type type-2
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R3 permit 1
(config-route-map)#set extcommunity soo 2008:2008
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R4 deny 1
(config-route-map)#match extcommunity S2
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R4 permit 2
(config-route-map)#match evpn-route-type type-5
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R4 permit 3
(config-route-map)#match evpn-route-type type-2
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R4 permit 4
(config-route-map)#match evpn-route-type type-1
(config-route-map)#commit
(config-route-map)#exit
```

6. Configure interfaces.

```
(config)#interface cd32/1
(config-if)#load-interval 30
```

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

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.5/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface xe33
(config-if)#load-interval 30
(config-if)#ip address 17.20.20.1/24
(config-if)#commit
(config-if)#exit

(config)#interface xe34
(config-if)#load-interval 30
(config-if)#ip address 17.21.21.1/24
(config-if)#commit
(config-if)#exit
```

7. Enable OSPF for underlay.

```
(config)#router ospf 100
(config-router)#bfd all-interfaces
(config-router)#network 10.10.10.5/32 area 0.0.0.0
(config-router)#network 17.20.20.0/24 area 0.0.0.0
(config-router)#network 17.21.21.0/24 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

8. Configure BGP global parameters and BGP EVPN overlay peers.

```
(config)#router bgp 65004
(config-router)#bgp router-id 10.10.10.5
(config-router)#bgp auto-policy-soft-reset enable
(config-router)#no bgp inbound-route-filter
(config-router)#neighbor 10.10.10.21 remote-as 65005
(config-router)#neighbor 10.10.10.21 update-source lo
(config-router)#neighbor 10.10.10.22 remote-as 65006
(config-router)#neighbor 10.10.10.22 update-source lo
```

9. Configure BGP Unnumbered (DCI / External Fabric) and EVPN for Unnumbered Peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor cd31/1 remote-as external
(config-router-unnum)#neighbor cd31/1 as-origination-interval 1
(config-router-unnum)#neighbor cd31/1 advertisement-interval 0
(config-router-unnum)#neighbor cd31/1 fall-over bfd
(config-router-unnum)#neighbor cd2/1 remote-as external
(config-router-unnum)#neighbor cd2/1 as-origination-interval 1
(config-router-unnum)#neighbor cd2/1 advertisement-interval 0
(config-router-unnum)#neighbor cd2/1 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
```

10. Configure BGP IPv4 Unicast Address Family with v4 Unnumbered Neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.5/32
(config-router-af)#max-paths ebgp 4
```

```
(config-router-af)#redistribute connected
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor cd31/1 activate
(config-router-v4-unnum)#neighbor cd2/1 activate
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
```

11. Configure BGP L2VPN EVPN with Numbered and Unnumbered Neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor 10.10.10.21 activate
(config-router-af)#neighbor 10.10.10.21 peer-type fabric-external
(config-router-af)#neighbor 10.10.10.21 allowas-in 1
(config-router-af)#neighbor 10.10.10.21 route-map R2 in
(config-router-af)#neighbor 10.10.10.21 route-map R1 out
(config-router-af)#neighbor 10.10.10.22 activate
(config-router-af)#neighbor 10.10.10.22 peer-type fabric-external
(config-router-af)#neighbor 10.10.10.22 allowas-in 1
(config-router-af)#neighbor 10.10.10.22 route-map R2 in
(config-router-af)#neighbor 10.10.10.22 route-map R1 out

(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-v4-unnum)#neighbor cd31/1 activate
(config-router-v4-unnum)#neighbor cd31/1 allowas-in 1
(config-router-v4-unnum)#neighbor cd31/1 route-map R4 in
(config-router-v4-unnum)#neighbor cd31/1 route-map R3 out
(config-router-v4-unnum)#neighbor cd2/1 activate
(config-router-v4-unnum)#neighbor cd2/1 allowas-in 1
(config-router-v4-unnum)#neighbor cd2/1 route-map R4 in
(config-router-v4-unnum)#neighbor cd2/1 route-map R3 out
(config-router-v4-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

12. Define extended community lists.

```
(config)#ip extcommunity-list standard S1 permit soo 1004:1004
(config)#ip extcommunity-list standard S2 permit soo 2008:2008
(config)#end
```

BLEAF2 Configuration

1. Enable VXLAN and IRB.

```
#configure terminal
(config)#nvo vxlan enable
(config)#nvo vxlan irb
```

2. Enable EVPN VXLAN multihoming.

```
(config)#evpn vxlan multihoming enable
```

3. Configure L3 VRF with EVPN stitching.

```
(config)#ip vrf l3_vrf_1
(config-vrf)#rd 7006:1001
(config-vrf)#route-target both 1001:1001
(config-vrf)#l3vni 1000
(config-vrf)#evpn-stitching
(config-vrf-evpn)#rd 7006:7001
```

```
(config-vrf-evpn)#route-target both 7001:7001
(config-vrf-evpn)#host-route-extend-l2vni 1001-1002
(config-vrf-evpn)#translation-l3vni 2000
(config-vrf-evpn)#exit-config-vrf-evpn-mode
(config-vrf)#commit
(config-vrf)#exit
```

4. Configure the Global VTEP IP address.

```
(config)#nvo vxlan vtep-ip-global 10.10.10.6
```

5. Define route-maps for EVPN control.

```
(config)#route-map R1 permit 1
(config-route-map)#set extcommunity soo 1004:1004
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R2 deny 1
(config-route-map)#match extcommunity S1
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R2 permit 2
(config-route-map)#match evpn-route-type type-5
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R2 permit 3
(config-route-map)#match evpn-route-type type-2
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R3 permit 1
(config-route-map)#set extcommunity soo 2008:2008
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R4 deny 1
(config-route-map)#match extcommunity S2
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R4 permit 2
(config-route-map)#match evpn-route-type type-5
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R4 permit 3
(config-route-map)#match evpn-route-type type-2
(config-route-map)#commit
(config-route-map)#exit

(config)#route-map R4 permit 4
(config-route-map)#match evpn-route-type type-1
(config-route-map)#commit
(config-route-map)#exit
```

6. Configure interfaces.

```
(config)#interface ce50/1
(config-if)#load-interval 30
(config-if)#ip address 17.22.22.1/24
(config-if)#commit
(config-if)#exit

(config)#interface ce52
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce53
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce54/1
(config-if)#load-interval 30
(config-if)#ip address 17.23.23.1/24
(config-if)#commit
(config-if)#exit

(config)#interface eth0
(config-if)#ip vrf forwarding management
(config-if)#ip address dhcp
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.6/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit
```

7. Enable OSPF for underlay.

```
(config)#router ospf 100
(config-router)#bfd all-interfaces
(config-router)#network 10.10.10.6/32 area 0.0.0.0
(config-router)#network 17.22.22.0/24 area 0.0.0.0
(config-router)#network 17.23.23.0/24 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

8. Configure BGP global parameters and BGP EVPN overlay peers.

```
(config)#router bgp 65004
(config-router)#bgp router-id 10.10.10.6
(config-router)#bgp auto-policy-soft-reset enable
(config-router)#no bgp inbound-route-filter
(config-router)#neighbor 10.10.10.21 remote-as 65005
(config-router)#neighbor 10.10.10.21 update-source lo
(config-router)#neighbor 10.10.10.21 as-origination-interval 1
(config-router)#neighbor 10.10.10.21 advertisement-interval 0
(config-router)#neighbor 10.10.10.22 remote-as 65006
(config-router)#neighbor 10.10.10.22 update-source lo
(config-router)#neighbor 10.10.10.22 as-origination-interval 1
(config-router)#neighbor 10.10.10.22 advertisement-interval 0
```

9. Configure BGP unnumbered (DCI / External Fabric) and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor ce52 remote-as external
(config-router-unnum)#neighbor ce52 as-origination-interval 1
(config-router-unnum)#neighbor ce52 advertisement-interval 0
(config-router-unnum)#neighbor ce52 fall-over bfd
(config-router-unnum)#neighbor ce53 remote-as external
(config-router-unnum)#neighbor ce53 as-origination-interval 1
(config-router-unnum)#neighbor ce53 advertisement-interval 0
(config-router-unnum)#neighbor ce53 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
(config-router)#commit
(config-router)#exit
```

10. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.6/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-v4-unnum)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor ce52 activate
(config-router-v4-unnum)#neighbor ce53 activate
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

11. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config)#address-family l2vpn evpn
(config-af)#neighbor 10.10.10.21 activate
(config-af)#neighbor 10.10.10.21 peer-type fabric-external
(config-af)#neighbor 10.10.10.21 allowas-in 1
(config-af)#neighbor 10.10.10.21 route-map R2 in
(config-af)#neighbor 10.10.10.21 route-map R1 out
(config-af)#neighbor 10.10.10.22 activate
(config-af)#neighbor 10.10.10.22 peer-type fabric-external
(config-af)#neighbor 10.10.10.22 allowas-in 1
(config-af)#neighbor 10.10.10.22 route-map R2 in
(config-af)#neighbor 10.10.10.22 route-map R1 out
(config-router-l2vpn-unnum)#bgp l2vpn-unnumbered-mode
(config-router-l2vpn-unnum)#neighbor ce52 activate
(config-router-l2vpn-unnum)#neighbor ce52 allowas-in 1
(config-router-l2vpn-unnum)#neighbor ce52 route-map R4 in
(config-router-l2vpn-unnum)#neighbor ce52 route-map R3 out
(config-router-l2vpn-unnum)#neighbor ce53 activate
(config-router-l2vpn-unnum)#neighbor ce53 allowas-in 1
(config-router-l2vpn-unnum)#neighbor ce53 route-map R4 in
(config-router-l2vpn-unnum)#neighbor ce53 route-map R3 out
(config-router-l2vpn-unnum)#exit-l2vpn-unnumbered-mode
(config-router)#exit-address-family
(config-router)#commit
(config-router)#exit
```

12. Define extended community lists.

```
(config)#ip extcommunity-list standard S1 permit soo 1004:1004
(config)#ip extcommunity-list standard S2 permit soo 2008:2008
(config)#end
```

BLEAF3 Configuration

1. Enable VXLAN and IRB.

```
#configure terminal
(config)#nvo vxlan enable
(config)#nvo vxlan irb
```

2. Enable EVPN VXLAN multihoming.

```
(config)#evpn vxlan multihoming enable
```

3. Configure L3 VRF with EVPN stitching.

```
(config)#ip vrf l3_vrf_1
(config-vrf)#rd 7007:5001
(config-vrf)#route-target both 5001:5001
(config-vrf)#l3vni 1000
(config-vrf)#evpn-stitching
(config-vrf-evpn)#rd 7007:7001
(config-vrf-evpn)#route-target both 7001:7001
(config-vrf-evpn)#host-route-extend-l2vni 1001-1002
(config-vrf-evpn)#translation-l3vni 2000
(config-vrf-evpn)#exit-config-vrf-evpn-mode
```

4. Configure the Global VTEP IP address.

```
(config)#nvo vxlan vtep-ip-global 10.10.10.7
```

5. Define route-maps for EVPN control.

```
(config)#route-map R1 permit 1
(config-route-map)#set extcommunity soo 1005:1005
(config-route-map)#exit
(config-router)#commit
(config-router)#exit

(config)#route-map R2 deny 1
(config-route-map)#match extcommunity S1
(config-route-map)#exit
(config-router)#commit
(config-router)#exit

(config)#route-map R2 permit 2
(config-route-map)#match evpn-route-type type-5
(config-route-map)#exit
(config-router)#commit
(config-router)#exit

(config)#route-map R2 permit 3
(config-route-map)#match evpn-route-type type-2
(config-route-map)#exit
(config-router)#commit
(config-router)#exit
```

6. Configure interfaces.

```
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.7/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe2
(config-if)#load-interval 30
(config-if)#ip address 17.19.19.1/24
(config-if)#commit
(config-if)#exit

(config)#interface xe3
(config-if)#load-interval 30
(config-if)#ip address 17.18.18.1/24
(config-if)#commit
(config-if)#exit

(config)#interface xe4
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit
```

7. Enable OSPF for underlay.

```
(config)#router ospf 100
(config-router)#bfd all-interfaces
(config-router)#network 10.10.10.7/32 area 0.0.0.0
(config-router)#network 17.18.18.0/24 area 0.0.0.0
(config-router)#network 17.19.19.0/24 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

8. Configure BGP global parameters and BGP EVPN overlay peers.

```
(config-router)#router bgp 65004
(config-router)#bgp router-id 10.10.10.7
(config-router)#bgp auto-policy-soft-reset enable
(config-router)#no bgp inbound-route-filter
(config-router)#neighbor 10.10.10.21 remote-as 65005
(config-router)#neighbor 10.10.10.21 update-source lo
(config-router)#neighbor 10.10.10.21 as-origination-interval 1
(config-router)#neighbor 10.10.10.21 advertisement-interval 0
(config-router)#neighbor 10.10.10.22 remote-as 65006
(config-router)#neighbor 10.10.10.22 update-source lo
(config-router)#neighbor 10.10.10.22 as-origination-interval 1
(config-router)#neighbor 10.10.10.22 advertisement-interval 0
```

9. Configure BGP unnumbered (DCI / External Fabric) and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor xe7 remote-as external
(config-router-unnum)#neighbor xe7 as-origination-interval 1
(config-router-unnum)#neighbor xe7 advertisement-interval 0
(config-router-unnum)#neighbor xe7 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
(config-router)#commit
(config-router)#exit
```

10. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config)#address-family ipv4 unicast
(config-af)#network 10.10.10.7/32
(config-af)#max-paths ebgp 4
(config-af)#redistribute connected
(config-af-unnum)#bgp v4-unnumbered-mode
```

```
(config-af-unnum)#neighbor xe7 activate
(config-af-unnum)#exit-v4-unnumbered-mode
(config-af)#commit
(config-af)#exit-address-family
```

11. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config)#address-family l2vpn evpn
(config-af)#neighbor 10.10.10.21 activate
(config-af)#neighbor 10.10.10.21 peer-type fabric-external
(config-af)#neighbor 10.10.10.21 allowas-in 1
(config-af)#neighbor 10.10.10.21 route-map R2 in
(config-af)#neighbor 10.10.10.21 route-map R1 out
(config-af)#neighbor 10.10.10.22 activate
(config-af)#neighbor 10.10.10.22 peer-type fabric-external
(config-af)#neighbor 10.10.10.22 allowas-in 1
(config-af)#neighbor 10.10.10.22 route-map R2 in
(config-af)#neighbor 10.10.10.22 route-map R1 out
(config-af)#bgp l2vpn-unnumbered-mode
(config-af-unnum)#neighbor xe7 activate
(config-af-unnum)#neighbor xe7 allowas-in 1
(config-af-unnum)#exit-l2vpn-unnumbered-mode
(config-af)#commit
(config-af)#exit-address-family
```

12. Define extended community lists.

```
(config)#ip extcommunity-list standard S1 permit soo 1005:1005
(config)#end
```

LEAF1 Configuration

1. Enable VXLAN and IRB.

```
#configure terminal
(config)#nvo vxlan enable
(config)#nvo vxlan irb
```

2. Configure EVPN ESI hold-time.

```
(config)#evpn esi hold-time 90
```

3. Enable EVPN VXLAN multihoming.

```
(config)#evpn vxlan multihoming enable
```

4. Configure L3 VRF with EVPN stitching.

```
(config)#ip vrf l3_vrf_1
(config-vrf)#rd 7001:1001
(config-vrf)#route-target both 1001:1001
(config-vrf)#l3vni 1000
(config-vrf)#commit
(config-vrf)#exit
```

5. Configure a MAC VRF (L2 VRF).

```
(config)#mac vrf l2_vrf_1
(config-vrf)#rd 10.10.10.1:1
(config-vrf)#route-target both 101:101
(config-vrf)#commit
(config-vrf)#exit
!
(config)#mac vrf l2_vrf_2
```

```
(config-vrf)#rd 10.10.10.1:2
(config-vrf)#route-target both 102:102
(config-vrf)#commit
(config-vrf)#exit
!
(config)#mac vrf l2_vrf_3
(config-vrf)#rd 10.10.10.1:3
(config-vrf)#route-target both 103:103
(config-vrf)#commit
(config-vrf)#exit
```

6. Configure the Global VTEP ip address.

```
(config)#nvo vxlan vtep-ip-global 10.10.10.1
```

7. Configure VXLAN L2 VNI with EVPN and IRB.

```
(config)#nvo vxlan id 1001 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_1
(config-nvo)#evpn irb1001
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit

(config)#nvo vxlan id 1002 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_2
(config-nvo)#evpn irb1002
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit

(config)#nvo vxlan id 1003 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_3
(config-nvo)#evpn irb1003
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit
```

8. Configure interfaces.

```
(config)#interface po1
(config-if)#switchport
(config-if)#bridge-group 1 spanning-tree disable
(config-if)#switchport mode hybrid
(config-if)#switchport mode hybrid acceptable-frame-type all
(config-if)#load-interval 30
(config-if)#evpn multi-homed system-mac 0000.aaaa.bbbb
(config-if)#commit
(config-if)#exit

(config)#interface irb1001
(config-if)#ip vrf forwarding l3_vrf_1
(config-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-if)#ip address 101.1.1.1/8 anycast
(config-if)#ipv6 address 1001::1/48 anycast
(config-if)#commit
(config-if)#exit

(config)#interface irb1002
(config-if)#ip vrf forwarding l3_vrf_1
```

```
(config-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-if)#ip address 102.1.1.1/8 anycast
(config-if)#ipv6 address 1002::1/48 anycast
(config-if)#commit
(config-if)#exit

(config)#interface irb1003
(config-if)#ip vrf forwarding l3_vrf_1
(config-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-if)#ip address 103.1.1.1/8 anycast
(config-if)#ipv6 address 1003::1/48 anycast
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.1/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface lo.management
(config-if)#ip vrf forwarding management
(config-if)#ip address 127.0.0.1/8
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface xe5
(config-if)#channel-group 1 mode active
(config-if)#commit
(config-if)#exit

(config)#interface xe8
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface xe47
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit
!
(config)#interface xe57
(config-if)#switchport
(config-if)#commit
(config-if)#exit
```

9. Configure EVPN IRB forwarding with anycast gateway MAC.

```
(config)#evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
```

10. Configure BGP global parameters.

```
(config)#router bgp 65001
(config-router)#bgp router-id 10.10.10.1
(config-router)#no bgp inbound-route-filter
```

11. Configure BGP unnumbered and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor xe8 as-origination-interval 1
(config-router-unnum)#neighbor xe8 advertisement-interval 0
(config-router-unnum)#neighbor xe8 fall-over bfd
(config-router-unnum)#neighbor xe47 as-origination-interval 1
(config-router-unnum)#neighbor xe47 advertisement-interval 0
(config-router-unnum)#neighbor xe47 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
(config-router)#commit
(config-router)#exit
```

12. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config)#address-family ipv4 unicast
(config-af)#network 10.10.10.1/32
(config-af)#max-paths ebgp 4
(config-af)#redistribute connected
(config-af)#redistribute static
(config-router-v4-unnum)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor xe8 activate
(config-router-v4-unnum)#neighbor xe8 allowas-in 1
(config-router-v4-unnum)#neighbor xe47 activate
(config-router-v4-unnum)#neighbor xe47 allowas-in 1
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
```

13. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-l2vpn-unnum)#neighbor xe8 activate
(config-router-l2vpn-unnum)#neighbor xe8 allowas-in 1
(config-router-l2vpn-unnum)#neighbor xe47 activate
(config-router-l2vpn-unnum)#neighbor xe47 allowas-in 1
(config-router-l2vpn-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)#exit-address-family
```

14. Configure BGP IPv4 address family for an L3 VRF.

```
(config-router)#address-family ipv4 vrf l3_vrf_1
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-af)#exit-address-family

(config-router)#address-family ipv6 vrf l3_vrf_1
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

15. Configure a Static Route in an L3 VRF.

```
(config)#ip route vrf l3_vrf_1 101.1.1.240/32 irb1001
```

16. Enable port-vlan mappings and Map VxLAN Identified to access-port for VxLAN.

```
(config)#nvo vxlan access-if port-vlan po1 101
(config-nvo-acc-if)#map vnid 1001
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit
```

```
(config)#nvo vxlan access-if port-vlan po1 102
(config-nvo-acc-if)#map vnid 1002
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit
(config)#nvo vxlan access-if port-vlan po1 103
(config-nvo-acc-if)#map vnid 1003
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit
```

LEAF2 Configuration

1. Enable VXLAN and IRB.

```
#configure terminal
(config)#nvo vxlan enable
(config)#nvo vxlan irb
```

2. Configure EVPN ESI hold-time.

```
(config)#evpn esi hold-time 90
```

3. Enable EVPN VXLAN multihoming.

```
(config)#evpn vxlan multihoming enable
```

4. Configure L3 VRF with EVPN stitching.

```
(config)#ip vrf l3_vrf_1
(config-vrf)#rd 7002:1001
(config-vrf)#route-target both 1001:1001
(config-vrf)#l3vni 1000
(config-vrf)#commit
(config-vrf)#exit
```

5. Configure a MAC VRF (L2 VRF).

```
(config)#mac vrf l2_vrf_1
(config-vrf)#rd 10.10.10.2:1
(config-vrf)#route-target both 101:101
(config-vrf)#commit
(config-vrf)#exit

(config)#mac vrf l2_vrf_2
(config-vrf)#rd 10.10.10.2:2
(config-vrf)#route-target both 102:102
(config-vrf)#commit
(config-vrf)#exit

(config)#mac vrf l2_vrf_3
(config-vrf)#rd 10.10.10.2:3
(config-vrf)#route-target both 103:103
(config-vrf)#commit
(config-vrf)#exit
```

6. Configure the Global VTEP ip address.

```
(config)#nvo vxlan vtep-ip-global 10.10.10.2
```

7. Configure VXLAN L2 VNI with EVPN and IRB.

```
(config)#nvo vxlan id 1001 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_1
(config-nvo)#evpn irb1001
```

```
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit

(config)#nvo vxlan id 1002 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_2
(config-nvo)#evpn irb1002
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit

(config)#nvo vxlan id 1003 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_3
(config-nvo)#evpn irb1003
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit
```

8. Configure interfaces.

```
(config)#interface po1
(config-if)#switchport
(config-if)#bridge-group 1 spanning-tree disable
(config-if)#switchport mode hybrid
(config-if)#switchport mode hybrid acceptable-frame-type all
(config-if)#load-interval 30
(config-if)#evpn multi-homed system-mac 0000.aaaa.bbbb
(config-if)#commit
(config-if)#exit

(config)#interface irb1001
(config-if)#ip vrf forwarding l3_vrf_1
(config-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-if)#ip address 101.1.1.1/8 anycast
(config-if)#ipv6 address 1001::1/48 anycast
(config-if)#commit
(config-if)#exit

(config)#interface irb1002
(config-if)#ip vrf forwarding l3_vrf_1
(config-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-if)#ip address 102.1.1.1/8 anycast
(config-if)#ipv6 address 1002::1/48 anycast
(config-if)#commit
(config-if)#exit

(config)#interface irb1003
(config-if)#ip vrf forwarding l3_vrf_1
(config-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-if)#ip address 103.1.1.1/8 anycast
(config-if)#ipv6 address 1003::1/48 anycast
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.2/32 secondary
(config-if)#ipv6 address ::1/128
```

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

(config)#interface xe1
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface xe3
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface xe10
(config-if)#channel-group 1 mode active
(config-if)#commit
(config-if)#exit
```

9. Configure EVPN IRB forwarding with Anycast gateway MAC.

```
(config)#evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
```

10. Configure BGP global parameters.

```
(config)#router bgp 65001
(config-router)#bgp router-id 10.10.10.2
(config-router)#no bgp inbound-route-filter
```

11. Configure BGP unnumbered and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor xe3 as-origination-interval 1
(config-router-unnum)#neighbor xe3 advertisement-interval 0
(config-router-unnum)#neighbor xe3 fall-over bfd
(config-router-unnum)#neighbor xe1 as-origination-interval 1
(config-router-unnum)#neighbor xe1 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
```

12. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.2/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-v4-unnum)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor xe3 activate
(config-router-v4-unnum)#neighbor xe3 allowas-in 1
(config-router-v4-unnum)#neighbor xe1 activate
(config-router-v4-unnum)#neighbor xe1 allowas-in 1
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)# exit-address-family
```

13. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)# bgp l2vpn-unnumbered-mode
(config-router-l2vpn-unnum)#neighbor xe3 activate
(config-router-l2vpn-unnum)#neighbor xe3 allowas-in 1
(config-router-l2vpn-unnum)#neighbor xe1 activate
(config-router-l2vpn-unnum)#neighbor xe1 allowas-in 1
(config-router-l2vpn-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)#exit-address-family
```

14. Configure BGP IPv4 address family for an L3 VRF.

```
(config-router)#address-family ipv4 vrf l3_vrf_1
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-af)#exit-address-family

(config-router)#address-family ipv6 vrf l3_vrf_1
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

15. Configure a static route in an L3 VRF.

```
(config)#ip route vrf l3_vrf_1 101.1.1.240/32 irb1001
```

16. Enable port-vlan mappings and Map VxLAN Identified to access-port for VxLAN.

```
(config)#nvo vxlan access-if port-vlan po1 101
(config-nvo-acc-if)#map vnid 1001
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit

(config)#nvo vxlan access-if port-vlan po1 102
(config-nvo-acc-if)#map vnid 1002
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit

(config)#nvo vxlan access-if port-vlan po1 103
(config-nvo-acc-if)#map vnid 1003
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit
```

LEAF3 Configuration

1. Enable VXLAN and IRB.

```
#configure terminal
(config)#nvo vxlan enable
(config)#nvo vxlan irb
```

2. Enable EVPN VXLAN multihoming.

```
(config)#evpn vxlan multihoming enable
```

3. Configure L3 VRF with EVPN stitching.

```
(config)#ip vrf l3_vrf_1
(config-vrf)#rd 7004:5001
(config-vrf)#route-target both 5001:5001
(config-vrf)#l3vni 1000
(config-vrf)#commit
(config-vrf)#exit
```

4. Configure a MAC VRF (L2 VRF).

```
(config)#mac vrf l2_vrf_1
(config-vrf)#rd 10.10.10.4:1
(config-vrf)#route-target both 201:201
```

```
(config-vrf)#commit
(config-vrf)#exit

(config)#mac vrf l2_vrf_2
(config-vrf)#rd 10.10.10.4:2
(config-vrf)#route-target both 202:202
(config-vrf)#commit
(config-vrf)#exit

(config)#mac vrf l2_vrf_3
(config-vrf)#rd 10.10.10.4:3
(config-vrf)#route-target both 203:203
(config-vrf)#commit
(config-vrf)#exit
```

5. Configure the Global VTEP ip address.

```
(config)#nvo vxlan vtep-ip-global 10.10.10.4
```

6. Configure VXLAN L2 VNI with EVPN and IRB.

```
(config)#nvo vxlan id 1001 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_1
(config-nvo)#evpn irb1001
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit

(config)#nvo vxlan id 1002 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_2
(config-nvo)#evpn irb1002
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit

(config)#nvo vxlan id 1003 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp l2_vrf_3
(config-nvo)#evpn irb1003
(config-nvo)#evpn irb-advertise-host-route
(config-nvo)#commit
(config-nvo)#exit
```

7. Configure interfaces.

```
(config)#interface ce4
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce31
(config-if)#commit
(config-if)#exit

(config)#interface irb1001
(config-if)#ip vrf forwarding l3_vrf_1
(config-if)#ip address 201.1.1.1/8
(config-if)#ipv6 address 2001::1/48
(config-if)#commit
(config-if)#exit
```

```
(config)#interface irb1002
(config-if)#ip vrf forwarding l3_vrf_1
(config-if)#ip address 202.1.1.1/8
(config-if)#ipv6 address 2002::1/48
(config-if)#commit
(config-if)#exit

(config)#interface irb1003
(config-if)#ip vrf forwarding l3_vrf_1
(config-if)#ip address 203.1.1.1/8
(config-if)#ipv6 address 2003::1/48
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.4/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface xe1/2
(config-if)#switchport
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface xe34
(config-if)#switchport
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit
```

8. Configure BGP Global Parameters.

```
(config)#router bgp 65001
(config-router)#bgp router-id 10.10.10.4
(config-router)#no bgp inbound-route-filter
```

9. Configure BGP unnumbered and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor ce4 as-origination-interval 1
(config-router-unnum)#neighbor ce4 advertisement-interval 0
(config-router-unnum)#neighbor ce4 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
```

10. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.4/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor ce4 activate
(config-router-v4-unnum)#neighbor ce4 allowas-in 1
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
```

11. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-v4-unnum)#neighbor ce4 activate
(config-router-v4-unnum)#neighbor ce4 allowas-in 1
(config-router-v4-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)#exit-address-family
```

12. Configure BGP IPv4 address family for an L3 VRF.

```
(config-router)#address-family ipv4 vrf l3_vrf_1
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-af)#exit-address-family

(config-router)#address-family ipv6 vrf l3_vrf_1
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

13. Configure a static route in an L3 VRF.

```
(config)#ip route vrf l3_vrf_1 201.1.1.240/32 irb1001
```

14. Enable port-vlan mappings and Map VxLAN Identified to access-port for VxLAN.

```
(config)#nvo vxlan access-if port-vlan xe34 2001
(config-nvo-acc-if)#map vnid 1001
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit

(config)#nvo vxlan access-if port-vlan xe34 2002
(config-nvo-acc-if)#map vnid 1002
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit

(config)#nvo vxlan access-if port-vlan xe34 2003
(config-nvo-acc-if)#map vnid 1003
(config-nvo-acc-if)#commit
(config-nvo-acc-if)#exit
```

R1 Configuration

1. Configure the interfaces.

```
#configure terminal
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.21/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface lo.management
(config-if)#ip vrf forwarding management
(config-if)#ip address 127.0.0.1/8
(config-if)#ipv6 address ::1/128
```

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

(config)#interface xe1/1
(config-if)#load-interval 30
(config-if)#ip address 17.20.20.2/24
(config-if)#commit
(config-if)#exit

(config)#interface xe1/2
(config-if)#load-interval 30
(config-if)#ip address 17.19.19.2/24
(config-if)#commit
(config-if)#exit

(config)#interface xe6/1
(config-if)#ip address 17.24.24.2/24
(config-if)#commit
(config-if)#exit

(config)#interface xe31/1
(config-if)#load-interval 30
(config-if)#ip address 17.22.22.2/24
(config-if)#commit
(config-if)#exit
```

2. Enable OSPF for underlay.

```
(config)#router ospf 100
(config-router)#bfd all-interfaces
(config-router)#network 10.10.10.21/32 area 0.0.0.0
(config-router)#network 17.19.19.0/24 area 0.0.0.0
(config-router)#network 17.20.20.0/24 area 0.0.0.0
(config-router)#network 17.22.22.0/24 area 0.0.0.0
(config-router)#network 17.24.24.0/24 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

3. Configure BGP global parameters.

```
(config)#router bgp 65005
(config-router)#bgp router-id 10.10.10.21
(config-router)#no bgp inbound-route-filter
(config-router)#neighbor 10.10.10.5 remote-as 65004
(config-router)#neighbor 10.10.10.5 update-source lo
(config-router)#neighbor 10.10.10.5 as-origination-interval 1
(config-router)#neighbor 10.10.10.5 advertisement-interval 0
(config-router)#neighbor 10.10.10.6 remote-as 65004
(config-router)#neighbor 10.10.10.6 update-source lo
(config-router)#neighbor 10.10.10.6 as-origination-interval 1
(config-router)#neighbor 10.10.10.6 advertisement-interval 0
(config-router)#neighbor 10.10.10.7 remote-as 65004
(config-router)#neighbor 10.10.10.7 update-source lo
(config-router)#neighbor 10.10.10.7 as-origination-interval 1
(config-router)#neighbor 10.10.10.7 advertisement-interval 0
(config-router)#neighbor 10.10.10.22 remote-as 65006
(config-router)#neighbor 10.10.10.22 update-source lo
(config-router)#neighbor 10.10.10.22 as-origination-interval 1
```

```
(config-router)#neighbor 10.10.10.22 advertisement-interval 0
```

4. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.21/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#exit-address-family
```

5. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor 10.10.-10.5 activate
(config-router-af)#neighbor 10.10.10.5 allowas-in 1
(config-router-af)#neighbor 10.10.10.6 activate
(config-router-af)#neighbor 10.10.10.6 allowas-in 1
(config-router-af)#neighbor 10.10.10.7 activate
(config-router-af)#neighbor 10.10.10.7 allowas-in 1
(config-router-af)#neighbor 10.10.10.22 activate
(config-router-af)#neighbor 10.10.10.22 allowas-in 1
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

R2 Configuration

1. Configure the interfaces.

```
#configure terminal
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.22/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface xe1/1
(config-if)#load-interval 30
(config-if)#ip address 17.21.21.2/24
(config-if)#commit
(config-if)#exit

(config)#interface xe1/2
(config-if)#load-interval 30
(config-if)#ip address 17.18.18.2/24
(config-if)#commit
(config-if)#exit

(config)#interface xe6/1
(config-if)#ip address 17.24.24.1/24
(config-if)#commit
(config-if)#exit

(config)#interface xe31/1
(config-if)#load-interval 30
(config-if)#ip address 17.23.23.2/24
(config-if)#commit
(config-if)#exit
```

2. Enable OSPF for underlay.

```
(config)#router ospf 100
(config-router)#bfd all-interfaces
(config-router)#network 10.10.10.22/32 area 0.0.0.0
(config-router)#network 17.18.18.0/24 area 0.0.0.0
(config-router)#network 17.21.21.0/24 area 0.0.0.0
(config-router)#network 17.23.23.0/24 area 0.0.0.0
(config-router)#network 17.24.24.0/24 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

3. Configure BGP global parameters.

```
(config)#router bgp 65006
(config-router)#bgp router-id 10.10.10.22
(config-router)#no bgp inbound-route-filter
(config-router)#neighbor 10.10.10.5 remote-as 65004
(config-router)#neighbor 10.10.10.5 update-source lo
(config-router)#neighbor 10.10.10.5 as-origination-interval 1
(config-router)#neighbor 10.10.10.5 advertisement-interval 0
(config-router)#neighbor 10.10.10.6 remote-as 65004
(config-router)#neighbor 10.10.10.6 update-source lo
(config-router)#neighbor 10.10.10.6 as-origination-interval 1
(config-router)#neighbor 10.10.10.6 advertisement-interval 0
(config-router)#neighbor 10.10.10.7 remote-as 65004
(config-router)#neighbor 10.10.10.7 update-source lo
(config-router)#neighbor 10.10.10.7 as-origination-interval 1
(config-router)#neighbor 10.10.10.7 advertisement-interval 0
(config-router)#neighbor 10.10.10.21 remote-as 65005
(config-router)#neighbor 10.10.10.21 update-source lo
(config-router)#neighbor 10.10.10.21 as-origination-interval 1
(config-router)#neighbor 10.10.10.21 advertisement-interval 0
```

4. Configure BGP IPv4 Unicast address family.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.22/32
(config-router-af)#redistribute connected
(config-router-af)#exit-address-family
```

5. Configure BGP L2VPN EVPN.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor 10.10.10.5 activate
(config-router-af)#neighbor 10.10.10.5 allowas-in 1
(config-router-af)#neighbor 10.10.10.6 activate
(config-router-af)#neighbor 10.10.10.6 allowas-in 1
(config-router-af)#neighbor 10.10.10.7 activate
(config-router-af)#neighbor 10.10.10.7 allowas-in 1
(config-router-af)#neighbor 10.10.10.21 activate
(config-router-af)#neighbor 10.10.10.21 allowas-in 1
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

SPINE1 Configuration

1. Configure the interfaces.

```
#configure terminal
```

```
(config)#interface ce49
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.11/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface xe8
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface xe45
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface xe47
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit
```

2. Configure BGP global parameters.

```
(config)#router bgp 65002
(config-router)#bgp router-id 10.10.10.11
(config-router)#no bgp inbound-route-filter
```

3. Configure BGP unnumbered and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor xe3 as-origination-interval 1
(config-router-unnum)#neighbor xe3 advertisement-interval 0
(config-router-unnum)#neighbor xe3 fall-over bfd
(config-router-unnum)#neighbor ce49 as-origination-interval 1
(config-router-unnum)#neighbor ce49 advertisement-interval 0
(config-router-unnum)#neighbor ce49 fall-over bfd
(config-router-unnum)#neighbor xe8 as-origination-interval 1
(config-router-unnum)#neighbor xe8 advertisement-interval 0
(config-router-unnum)#neighbor xe8 fall-over bfd
(config-router-unnum)#neighbor xe45 as-origination-interval 1
(config-router-unnum)#neighbor xe45 advertisement-interval 0
(config-router-unnum)#neighbor xe45 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
```

4. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.11/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-v4-unnum)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor xe3 activate
(config-router-v4-unnum)#neighbor xe3 allowas-in 1
```

```
(config-router-v4-unnum)#neighbor ce49 activate
(config-router-v4-unnum)#neighbor ce49 allowas-in 1
(config-router-v4-unnum)#neighbor xe8 activate
(config-router-v4-unnum)#neighbor xe8 allowas-in 1
(config-router-v4-unnum)#neighbor xe45 activate
(config-router-v4-unnum)#neighbor xe45 allowas-in 1
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
```

5. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-v4-unnum)#neighbor xe3 activate
(config-router-v4-unnum)#neighbor xe3 allowas-in 1
(config-router-v4-unnum)#neighbor ce49 activate
(config-router-v4-unnum)#neighbor ce49 allowas-in 1
(config-router-v4-unnum)#neighbor xe8 activate
(config-router-v4-unnum)#neighbor xe8 allowas-in 1
(config-router-v4-unnum)#neighbor xe45 activate
(config-router-v4-unnum)#neighbor xe45 allowas-in 1
(config-router-v4-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

SPINE2 Configuration

1. Configure the interfaces.

```
#configure terminal
(config)#interface ce50
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.12/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface xe1
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface xe47
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface xe57
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit
```

2. Configure BGP global parameters.

```
(config)#router bgp 65002
(config-router)#bgp router-id 10.10.10.12
(config-router)#no bgp inbound-route-filter
```

3. Configure BGP unnumbered and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor xe57 as-origination-interval 1
(config-router-unnum)#neighbor xe57 advertisement-interval 0
(config-router-unnum)#neighbor xe57 fall-over bfd
(config-router-unnum)#neighbor xe47 as-origination-interval 1
(config-router-unnum)#neighbor xe47 advertisement-interval 0
(config-router-unnum)#neighbor xe47 fall-over bfd
(config-router-unnum)#neighbor ce50 as-origination-interval 1
(config-router-unnum)#neighbor ce50 advertisement-interval 0
(config-router-unnum)#neighbor xe1 as-origination-interval 1
(config-router-unnum)#neighbor xe1 advertisement-interval 0
(config-router-unnum)#neighbor xe1 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
```

4. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.12/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor xe57 activate
(config-router-v4-unnum)#neighbor xe57 allowas-in 1
(config-router-v4-unnum)#neighbor xe47 activate
(config-router-v4-unnum)#neighbor xe47 allowas-in 1
(config-router-v4-unnum)#neighbor ce50 activate
(config-router-v4-unnum)#neighbor ce50 allowas-in 1
(config-router-v4-unnum)#neighbor xe1 activate
(config-router-v4-unnum)#neighbor xe1 allowas-in 1
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
```

5. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-v4-unnum)#neighbor xe57 activate
(config-router-v4-unnum)#neighbor xe57 allowas-in 1
(config-router-v4-unnum)#neighbor xe47 activate
(config-router-v4-unnum)#neighbor xe47 allowas-in 1
(config-router-v4-unnum)#neighbor ce50 activate
(config-router-v4-unnum)#neighbor ce50 allowas-in 1
(config-router-v4-unnum)#neighbor xe1 activate
(config-router-v4-unnum)#neighbor xe1 allowas-in 1
(config-router-v4-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

SPINE3 Configuration

1. Configure the interfaces.

```
#configure terminal
(config)#interface ce49
(config-if)#load-interval 30
(config-if)#commit
(config-if)#ecit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.16/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#ecit

(config)#interface xe4
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit
```

2. Configure BGP global parameters.

```
(config)#router bgp 65002
(config-router)#bgp router-id 10.10.10.16
(config-router)#no bgp inbound-route-filter
```

3. Configure BGP unnumbered and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor xe7 as-origination-interval 1
(config-router-unnum)#neighbor xe7 advertisement-interval 0
(config-router-unnum)#neighbor xe7 fall-over bfd
(config-router-unnum)#neighbor ce49 as-origination-interval 1
(config-router-unnum)#neighbor ce49 advertisement-interval 0
(config-router-unnum)#neighbor ce49 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
```

4. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.12/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static

(config-router-v4-unnum)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor xe57 activate
(config-router-v4-unnum)#neighbor xe57 allowas-in 1
(config-router-v4-unnum)#neighbor xe47 activate
(config-router-v4-unnum)#neighbor xe47 allowas-in 1
(config-router-v4-unnum)#neighbor ce50 activate
(config-router-v4-unnum)#neighbor ce50 allowas-in 1
(config-router-v4-unnum)#neighbor xe1 activate
(config-router-v4-unnum)#neighbor xe1 allowas-in 1
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
```

5. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-v4-unnum)#neighbor xe7 activate
(config-router-v4-unnum)#neighbor xe7 allowas-in 1
(config-router-v4-unnum)#neighbor ce49 activate
(config-router-v4-unnum)#neighbor ce49 allowas-in 1
(config-router-v4-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

SuperSpine1 Configuration

1. Configure the interfaces.

```
#configure terminal
(config)#interface ce1/3
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce1/4
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce2/1
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce29/1
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce31/1
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.14/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit
```

2. Configure BGP global parameters.

```
(config)#router bgp 65003
(config-router)#bgp router-id 10.10.10.14
(config-router)#no bgp inbound-route-filter
```

3. Configure BGP unnumbered and EVPN for unnumbered peers.

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor ce1/4 as-origination-interval 1
(config-router-unnum)#neighbor ce1/4 advertisement-interval 0
```

```
(config-router-unnum)#neighbor ce1/4 fall-over bfd
(config-router-unnum)#neighbor ce31/1 as-origination-interval 1
(config-router-unnum)#neighbor ce31/1 advertisement-interval 0
(config-router-unnum)#neighbor ce31/1 fall-over bfd
(config-router-unnum)#neighbor ce29/1 as-origination-interval 1
(config-router-unnum)#neighbor ce29/1 advertisement-interval 0
(config-router-unnum)#neighbor ce29/1 fall-over bfd
(config-router-unnum)#neighbor ce1/3 as-origination-interval 1
(config-router-unnum)#neighbor ce1/3 advertisement-interval 0
(config-router-unnum)#neighbor ce1/3 fall-over bfd
(config-router-unnum)#exit-unnumbered-mode
```

4. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.14/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-v4-unnum)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor ce1/4 activate
(config-router-v4-unnum)#neighbor ce1/4 allowas-in 1
(config-router-v4-unnum)#neighbor ce31/1 activate
(config-router-v4-unnum)#neighbor ce31/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce29/1 activate
(config-router-v4-unnum)#neighbor ce29/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce1/3 activate
(config-router-v4-unnum)#neighbor ce1/3 allowas-in 1
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
```

5. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-v4-unnum)#neighbor ce1/4 activate
(config-router-v4-unnum)#neighbor ce1/4 allowas-in 1
(config-router-v4-unnum)#neighbor ce31/1 activate
(config-router-v4-unnum)#neighbor ce31/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce29/1 activate
(config-router-v4-unnum)#neighbor ce29/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce1/3 activate
(config-router-v4-unnum)#neighbor ce1/3 allowas-in 1
(config-router-v4-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)# exit-address-family
(config-router)#commit
(config-router)#exit
```

SuperSpine2 Configuration

1. Configure the interfaces.

```
#configure terminal
(config)#interface ce9/1
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce11/1
```

```
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce14/1
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce15/1
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 10.10.10.15/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit
```

2. Configure BGP global parameters, unnumbered, and EVPN for unnumbered peers.

```
(config)#router bgp 65003
(config-router)#bgp router-id 10.10.10.15
(config-router)#no bgp inbound-route-filter
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor ce14/1 fall-over bfd
(config-router-unnum)#neighbor ce15/1 fall-over bfd
(config-router-unnum)#neighbor ce9/1 fall-over bfd
(config-router-unnum)#neighbor ce11/1 fall-over bfd
(config-router-unnum)#neighbor ce14/1 as-origination-interval 1
(config-router-unnum)#neighbor ce14/1 advertisement-interval 0
(config-router-unnum)#neighbor ce15/1 as-origination-interval 1
(config-router-unnum)#neighbor ce15/1 advertisement-interval 0
(config-router-unnum)#neighbor ce9/1 as-origination-interval 1
(config-router-unnum)#neighbor ce9/1 advertisement-interval 0
(config-router-unnum)#neighbor ce11/1 as-origination-interval 1
(config-router-unnum)#neighbor ce11/1 advertisement-interval 0
(config-router-unnum)#exit-unnumbered-mode
```

3. Configure BGP IPv4 Unicast address family with v4 unnumbered neighbors.

```
(config-router)#address-family ipv4 unicast
(config-router-af)#network 10.10.10.15/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#redistribute connected
(config-router-af)#redistribute static
(config-router-v4-unnum)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor ce14/1 activate
(config-router-v4-unnum)#neighbor ce14/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce15/1 activate
(config-router-v4-unnum)#neighbor ce15/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce9/1 activate
(config-router-v4-unnum)#neighbor ce9/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce11/1 activate
(config-router-v4-unnum)#neighbor ce11/1 allowas-in 1
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
```

4. Configure BGP L2VPN EVPN with numbered and unnumbered neighbors.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-v4-unnum)#neighbor ce14/1 activate
(config-router-v4-unnum)#neighbor ce14/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce15/1 activate
(config-router-v4-unnum)#neighbor ce15/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce9/1 activate
(config-router-v4-unnum)#neighbor ce9/1 allowas-in 1
(config-router-v4-unnum)#neighbor ce11/1 activate
(config-router-v4-unnum)#neighbor ce11/1 allowas-in 1
(config-router-v4-unnum)#exit-l2vpn-unnumbered-mode
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

SW1 Configuration

1. Configure the interfaces.

```
#configure terminal
(config)#interface po1
(config-if)#switchport
(config-if)#bridge-group 1 spanning-tree disable
(config-if)#switchport mode hybrid
(config-if)#switchport mode hybrid acceptable-frame-type all
(config-if)#switchport hybrid allowed vlan all
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit

(config)#interface ce53
(config-if)#commit
(config-if)#exit

(config)#interface ce54
(config-if)#commit
(config-if)#exit

(config)#interface xe5
(config-if)#channel-group 1 mode active
(config-if)#commit
(config-if)#exit

(config)#interface xe10
(config-if)#channel-group 1 mode active
(config-if)#commit
(config-if)#exit

(config)#interface xe57
(config-if)#switchport
(config-if)#bridge-group 1 spanning-tree disable
(config-if)#switchport mode hybrid
(config-if)#switchport mode hybrid acceptable-frame-type all
(config-if)#switchport hybrid allowed vlan all
```

```
(config-if)#load-interval 30
(config-if)#commit
(config-if)#exit
```

Running Configuration

BLEAF1

```
(config)#nvo vxlan enable
!
(config)#nvo vxlan irb
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_1
 rd 7005:1001
 route-target both 1001:1001
 l3vni 1000
 evpn-stitching
  rd 7005:7001
  route-target both 7001:7001
  host-route-extend-l2vni 1001-1002
  translation-l3vni 2000
  exit-config-vrf-evpn-mode
!
nvo vxlan vtep-ip-global 10.10.10.5
!
route-map R1 permit 1
 set extcommunity soo 1004:1004
!
route-map R2 deny 1
 match extcommunity S1
!
route-map R2 permit 2
 match evpn-route-type type-5
!
route-map R2 permit 3
 match evpn-route-type type-2
!
route-map R3 permit 1
 set extcommunity soo 2008:2008
!
route-map R4 deny 1
 match extcommunity S2
!
route-map R4 permit 2
 match evpn-route-type type-5
!
route-map R4 permit 3
 match evpn-route-type type-2
!
route-map R4 permit 4
 match evpn-route-type type-1
!
interface cd32/1
```

```
 load-interval 30
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.5/32 secondary
 ipv6 address ::1/128
!
interface xe33
 load-interval 30
 ip address 17.20.20.1/24
!
interface xe34
 load-interval 30
 ip address 17.21.21.1/24
!
router ospf 100
 bfd all-interfaces
 network 10.10.10.5/32 area 0.0.0.0
 network 17.20.20.0/24 area 0.0.0.0
 network 17.21.21.0/24 area 0.0.0.0
!
router bgp 65004
 bgp router-id 10.10.10.5
 bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor 10.10.10.21 remote-as 65005
 neighbor 10.10.10.21 update-source lo
 neighbor 10.10.10.22 remote-as 65006
 neighbor 10.10.10.22 update-source lo
 !
 bgp unnumbered-mode
  neighbor cd31/1 remote-as external
  neighbor cd31/1 as-origination-interval 1
  neighbor cd31/1 advertisement-interval 0
  neighbor cd31/1 fall-over bfd
  neighbor cd2/1 remote-as external
  neighbor cd2/1 as-origination-interval 1
  neighbor cd2/1 advertisement-interval 0
  neighbor cd2/1 fall-over bfd
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
  network 10.10.10.5/32
  max-paths ebgp 4
  redistribute connected
  !
  bgp v4-unnumbered-mode
  neighbor cd31/1 activate
  neighbor cd2/1 activate
  exit-v4-unnumbered-mode
  !
 exit-address-family
 !
 address-family l2vpn evpn
  neighbor 10.10.10.21 activate
  neighbor 10.10.10.21 peer-type fabric-external
  neighbor 10.10.10.21 allowas-in 1
```

```
  neighbor 10.10.10.21 route-map R2 in
  neighbor 10.10.10.21 route-map R1 out
  neighbor 10.10.10.22 activate
  neighbor 10.10.10.22 peer-type fabric-external
  neighbor 10.10.10.22 allowas-in 1
  neighbor 10.10.10.22 route-map R2 in
  neighbor 10.10.10.22 route-map R1 out
 !
  bgp l2vpn-unnumbered-mode
  neighbor cd31/1 activate
  neighbor cd31/1 allowas-in 1
  neighbor cd31/1 route-map R4 in
  neighbor cd31/1 route-map R3 out
  neighbor cd2/1 activate
  neighbor cd2/1 allowas-in 1
  neighbor cd2/1 route-map R4 in
  neighbor cd2/1 route-map R3 out
  exit-l2vpn-unnumbered-mode
  !
 exit-address-family
 !
!
ip extcommunity-list standard S1 permit soo 1004:1004
ip extcommunity-list standard S2 permit soo 2008:2008
!
```

BLEAF2

```
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_1
 rd 7006:1001
 route-target both 1001:1001
 l3vni 1000
 evpn-stitching
 rd 7006:7001
 route-target both 7001:7001
 host-route-extend-l2vni 1001-1002
  translation-l3vni 2000
  exit-config-vrf-evpn-mode
!
nvo vxlan vtep-ip-global 10.10.10.6
!
route-map R1 permit 1
 set extcommunity soo 1004:1004
!
route-map R2 deny 1
 match extcommunity S1
!
route-map R2 permit 2
 match evpn-route-type type-5
```

```
!
route-map R2 permit 3
 match evpn-route-type type-2
!
route-map R3 permit 1
 set extcommunity soo 2008:2008
!
route-map R4 deny 1
 match extcommunity S2
!
route-map R4 permit 2
 match evpn-route-type type-5
!
route-map R4 permit 3
 match evpn-route-type type-2
!
route-map R4 permit 4
 match evpn-route-type type-1
!
interface ce50/1
 load-interval 30
 ip address 17.22.22.1/24
!
interface ce52
 load-interval 30
!
interface ce53
 load-interval 30
!
interface ce54/1
 load-interval 30
 ip address 17.23.23.1/24
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.6/32 secondary
 ipv6 address ::1/128
!
router ospf 100
 bfd all-interfaces
 network 10.10.10.6/32 area 0.0.0.0
 network 17.22.22.0/24 area 0.0.0.0
 network 17.23.23.0/24 area 0.0.0.0
!
router bgp 65004
 bgp router-id 10.10.10.6
 bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor 10.10.10.21 remote-as 65005
 neighbor 10.10.10.21 update-source lo
 neighbor 10.10.10.21 as-origination-interval 1
 neighbor 10.10.10.21 advertisement-interval 0
 neighbor 10.10.10.22 remote-as 65006
```

```
 neighbor 10.10.10.22 update-source lo
 neighbor 10.10.10.22 as-origination-interval 1
 neighbor 10.10.10.22 advertisement-interval 0
 !
 bgp unnumbered-mode
  neighbor ce52 remote-as external
  neighbor ce52 as-origination-interval 1
  neighbor ce52 advertisement-interval 0
  neighbor ce52 fall-over bfd
  neighbor ce53 remote-as external
  neighbor ce53 as-origination-interval 1
  neighbor ce53 advertisement-interval 0
  neighbor ce53 fall-over bfd
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
  network 10.10.10.6/32
  max-paths ebgp 4
  redistribute connected
!
 bgp v4-unnumbered-mode
 neighbor ce52 activate
 neighbor ce53 activate
 exit-v4-unnumbered-mode
!
 exit-address-family
!
 address-family l2vpn evpn
  neighbor 10.10.10.21 activate
  neighbor 10.10.10.21 peer-type fabric-external
  neighbor 10.10.10.21 allowas-in 1
  neighbor 10.10.10.21 route-map R2 in
  neighbor 10.10.10.21 route-map R1 out
  neighbor 10.10.10.22 activate
  neighbor 10.10.10.22 peer-type fabric-external
  neighbor 10.10.10.22 allowas-in 1
  neighbor 10.10.10.22 route-map R2 in
  neighbor 10.10.10.22 route-map R1 out
   bgp l2vpn-unnumbered-mode
   neighbor ce52 activate
   neighbor ce52 allowas-in 1
  neighbor ce52 route-map R4 in
  neighbor ce52 route-map R3 out
  neighbor ce53 activate
  neighbor ce53 allowas-in 1
  neighbor ce53 route-map R4 in
  neighbor ce53 route-map R3 out
  exit-l2vpn-unnumbered-mode
  !
 exit-address-family
 !
ip extcommunity-list standard S1 permit soo 1004:1004
ip extcommunity-list standard S2 permit soo 2008:2008
!
end
```

BLEAF3

```
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_1
 rd 7007:5001
 route-target both 5001:5001
 l3vni 1000
 evpn-stitching
  rd 7007:7001
  route-target both 7001:7001
  host-route-extend-l2vni 1001-1002
  translation-l3vni 2000
  exit-config-vrf-evpn-mode
!
nvo vxlan vtep-ip-global 10.10.10.7
!
route-map R1 permit 1
 set extcommunity soo 1005:1005
!
route-map R2 deny 1
 match extcommunity S1
!
route-map R2 permit 2
 match evpn-route-type type-5
!
route-map R2 permit 3
 match evpn-route-type type-2
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.7/32 secondary
 ipv6 address ::1/128
!
interface xe2
 load-interval 30
 ip address 17.19.19.1/24
!
interface xe3
 load-interval 30
 ip address 17.18.18.1/24
!
interface xe4
 load-interval 30
!
router ospf 100
 bfd all-interfaces
 network 10.10.10.7/32 area 0.0.0.0
 network 17.18.18.0/24 area 0.0.0.0
 network 17.19.19.0/24 area 0.0.0.0
!
```

```
router bgp 65004
 bgp router-id 10.10.10.7
 bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor 10.10.10.21 remote-as 65005
 neighbor 10.10.10.21 update-source lo
 neighbor 10.10.10.21 as-origination-interval 1
 neighbor 10.10.10.21 advertisement-interval 0
 neighbor 10.10.10.22 remote-as 65006
 neighbor 10.10.10.22 update-source lo
 neighbor 10.10.10.22 as-origination-interval 1
 neighbor 10.10.10.22 advertisement-interval 0
!
 bgp unnumbered-mode
  neighbor xe7 remote-as external
  neighbor xe7 as-origination-interval 1
  neighbor xe7 advertisement-interval 0
  neighbor xe7 fall-over bfd
  exit-unnumbered-mode
  !
 address-family ipv4 unicast
  network 10.10.10.7/32
  max-paths ebgp 4
  redistribute connected
  bgp v4-unnumbered-mode
   neighbor xe7 activate
   exit-v4-unnumbered-mode
  !
 exit-address-family
 !
 address-family l2vpn evpn
  neighbor 10.10.10.21 activate
  neighbor 10.10.10.21 peer-type fabric-external
  neighbor 10.10.10.21 allowas-in 1
  neighbor 10.10.10.21 route-map R2 in
  neighbor 10.10.10.21 route-map R1 out
  neighbor 10.10.10.22 activate
  neighbor 10.10.10.22 peer-type fabric-external
  neighbor 10.10.10.22 allowas-in 1
  neighbor 10.10.10.22 route-map R2 in
  neighbor 10.10.10.22 route-map R1 out
  bgp l2vpn-unnumbered-mode
   neighbor xe7 activate
   neighbor xe7 allowas-in 1
   exit-l2vpn-unnumbered-mode
   !
  exit-address-family
!
ip extcommunity-list standard S1 permit soo 1005:1005
!
```

LEAF1 Configuration

```
!
nvo vxlan enable
!
```

```
nvo vxlan irb
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_1
 rd 7001:1001
 route-target both 1001:1001
 l3vni 1000
!
mac vrf l2_vrf_1
 rd 10.10.10.1:1
 route-target both 101:101
!
mac vrf l2_vrf_2
 rd 10.10.10.1:2
 route-target both 102:102
!
mac vrf l2_vrf_3
 rd 10.10.10.1:3
 route-target both 103:103
!
nvo vxlan vtep-ip-global 10.10.10.1
!
nvo vxlan id 1001 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_1
 evpn irb1001
 evpn irb-advertise-host-route
!
nvo vxlan id 1002 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_2
 evpn irb1002
 evpn irb-advertise-host-route
!
nvo vxlan id 1003 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_3
 evpn irb1003
 evpn irb-advertise-host-route
!
interface po1
 switchport
 bridge-group 1 spanning-tree disable
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 load-interval 30
 evpn multi-homed system-mac 0000.aaaa.bbbb
!
interface irb1001
 ip vrf forwarding l3_vrf_1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 101.1.1.1/8 anycast
 ipv6 address 1001::1/48 anycast
!
interface irb1002
 ip vrf forwarding l3_vrf_1
 evpn irb-if-forwarding anycast-gateway-mac
```

```
 ip address 102.1.1.1/8 anycast
 ipv6 address 1002::1/48 anycast
!
interface irb1003
 ip vrf forwarding l3_vrf_1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 103.1.1.1/8 anycast
 ipv6 address 1003::1/48 anycast
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.1/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe5
 channel-group 1 mode active
!
interface xe8
 load-interval 30
!
interface xe47
 load-interval 30
!
interface xe57
 switchport
!
evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
!
router bgp 65001
 bgp router-id 10.10.10.1
 no bgp inbound-route-filter
 !
 bgp unnumbered-mode
  neighbor xe8 as-origination-interval 1
  neighbor xe8 advertisement-interval 0
  neighbor xe8 fall-over bfd
  neighbor xe47 as-origination-interval 1
  neighbor xe47 advertisement-interval 0
  neighbor xe47 fall-over bfd
  exit-unnumbered-mode
  !
 address-family ipv4 unicast
  network 10.10.10.1/32
  max-paths ebgp 4
  redistribute connected
  redistribute static
  bgp v4-unnumbered-mode
   neighbor xe8 activate
   neighbor xe8 allowas-in 1
   neighbor xe47 activate
   neighbor xe47 allowas-in 1
   exit-v4-unnumbered-mode
```

```
   !
  exit-address-family
  !
 address-family l2vpn evpn
  bgp l2vpn-unnumbered-mode
   neighbor xe8 activate
   neighbor xe8 allowas-in 1
   neighbor xe47 activate
   neighbor xe47 allowas-in 1
   exit-l2vpn-unnumbered-mode
   !
  exit-address-family
  !
 address-family ipv4 vrf l3_vrf_1
  max-paths ebgp 4
  redistribute connected
  redistribute static
  exit-address-family
  !
 address-family ipv6 vrf l3_vrf_1
  max-paths ebgp 4
  redistribute connected
  redistribute static
  exit-address-family
  !
!
ip route vrf l3_vrf_1 101.1.1.240/32 irb1001
!
nvo vxlan access-if port-vlan po1 101
 map vnid 1001
!
nvo vxlan access-if port-vlan po1 102
 map vnid 1002
!
nvo vxlan access-if port-vlan po1 103
 map vnid 1003
!
```

LEAF2

```
!
nvo vxlan enable
nvo vxlan irb
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_1
 rd 7002:1001
 route-target both 1001:1001
 l3vni 1000
!
mac vrf l2_vrf_1
 rd 10.10.10.2:1
 route-target both 101:101
```

```
!
mac vrf l2_vrf_2
 rd 10.10.10.2:2
 route-target both 102:102
!
mac vrf l2_vrf_3
 rd 10.10.10.2:3
 route-target both 103:103
!
nvo vxlan vtep-ip-global 10.10.10.2
!
nvo vxlan id 1001 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_1
 evpn irb1001
 evpn irb-advertise-host-route
!
nvo vxlan id 1002 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_2
 evpn irb1002
 evpn irb-advertise-host-route
!
nvo vxlan id 1003 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_3
 evpn irb1003
 evpn irb-advertise-host-route
!
interface po1
 switchport
 bridge-group 1 spanning-tree disable
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 load-interval 30
 evpn multi-homed system-mac 0000.aaaa.bbbb
!
interface irb1001
 ip vrf forwarding l3_vrf_1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 101.1.1.1/8 anycast
 ipv6 address 1001::1/48 anycast
!
interface irb1002
 ip vrf forwarding l3_vrf_1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 102.1.1.1/8 anycast
 ipv6 address 1002::1/48 anycast
!
interface irb1003
 ip vrf forwarding l3_vrf_1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 103.1.1.1/8 anycast
 ipv6 address 1003::1/48 anycast
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.2/32 secondary
 ipv6 address ::1/128
!
```

```
interface xe1
 load-interval 30
!
interface xe3
 load-interval 30
!
interface xe10
 channel-group 1 mode active
!
evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
!
router bgp 65001
 bgp router-id 10.10.10.2
 no bgp inbound-route-filter
 bgp unnumbered-mode
  neighbor xe3 as-origination-interval 1
  neighbor xe3 advertisement-interval 0
  neighbor xe3 fall-over bfd
  neighbor xe1 as-origination-interval 1
  neighbor xe1 fall-over bfd
  exit-unnumbered-mode
  !
 address-family ipv4 unicast
  network 10.10.10.2/32
  max-paths ebgp 4
  redistribute connected
  redistribute static
  bgp v4-unnumbered-mode
   neighbor xe3 activate
   neighbor xe3 allowas-in 1
   neighbor xe1 activate
   neighbor xe1 allowas-in 1
   exit-v4-unnumbered-mode
   !
  exit-address-family
  !
 address-family l2vpn evpn
  bgp l2vpn-unnumbered-mode
   neighbor xe3 activate
   neighbor xe3 allowas-in 1
   neighbor xe1 activate
   neighbor xe1 allowas-in 1
   exit-l2vpn-unnumbered-mode
   !
  exit-address-family
  !
 address-family ipv4 vrf l3_vrf_1
  max-paths ebgp 4
  redistribute connected
  redistribute static
  exit-address-family
  !
 address-family ipv6 vrf l3_vrf_1
  max-paths ebgp 4
  redistribute connected
  redistribute static
  exit-address-family
```

```
  !
!
ip route vrf l3_vrf_1 101.1.1.240/32 irb1001
!
nvo vxlan access-if port-vlan po1 101
 map vnid 1001
!
nvo vxlan access-if port-vlan po1 102
 map vnid 1002
!
nvo vxlan access-if port-vlan po1 103
 map vnid 1003
!
```

LEAF3

```
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_1
 rd 7004:5001
 route-target both 5001:5001
 l3vni 1000
!
mac vrf l2_vrf_1
 rd 10.10.10.4:1
 route-target both 201:201
!
mac vrf l2_vrf_2
 rd 10.10.10.4:2
 route-target both 202:202
!
mac vrf l2_vrf_3
 rd 10.10.10.4:3
 route-target both 203:203
!
nvo vxlan vtep-ip-global 10.10.10.4
!
nvo vxlan id 1001 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_1
 evpn irb1001
 evpn irb-advertise-host-route
!
nvo vxlan id 1002 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_2
 evpn irb1002
 evpn irb-advertise-host-route
!
nvo vxlan id 1003 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_3
 evpn irb1003
 evpn irb-advertise-host-route
```

```
!
interface ce4
 load-interval 30
!
interface ce31
!
interface irb1001
 ip vrf forwarding l3_vrf_1
 ip address 201.1.1.1/8
 ipv6 address 2001::1/48
!
interface irb1002
 ip vrf forwarding l3_vrf_1
 ip address 202.1.1.1/8
 ipv6 address 2002::1/48
!
interface irb1003
 ip vrf forwarding l3_vrf_1
 ip address 203.1.1.1/8
 ipv6 address 2003::1/48
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.4/32 secondary
 ipv6 address ::1/128
!
interface xe1/2
 switchport
 load-interval 30
!
interface xe34
 switchport
 load-interval 30
!
router bgp 65001
 bgp router-id 10.10.10.4
 no bgp inbound-route-filter
 bgp unnumbered-mode
  neighbor ce4 as-origination-interval 1
  neighbor ce4 advertisement-interval 0
  neighbor ce4 fall-over bfd
  exit-unnumbered-mode
  !
 address-family ipv4 unicast
  network 10.10.10.4/32
  max-paths ebgp 4
  redistribute connected
  redistribute static
  bgp v4-unnumbered-mode
   neighbor ce4 activate
   neighbor ce4 allowas-in 1
   exit-v4-unnumbered-mode
   !
  exit-address-family
  !
 address-family l2vpn evpn
  bgp l2vpn-unnumbered-mode
```

```
   neighbor ce4 activate
   neighbor ce4 allowas-in 1
   exit-l2vpn-unnumbered-mode
   !
  exit-address-family
  !
 address-family ipv4 vrf l3_vrf_1
  max-paths ebgp 4
  redistribute connected
  redistribute static
  exit-address-family
  !
 address-family ipv6 vrf l3_vrf_1
  max-paths ebgp 4
  redistribute connected
  redistribute static
  exit-address-family
  !
!
ip route vrf l3_vrf_1 201.1.1.240/32 irb1001
!
nvo vxlan access-if port-vlan xe34 2001
 map vnid 1001
!
nvo vxlan access-if port-vlan xe34 2002
 map vnid 1002
!
nvo vxlan access-if port-vlan xe34 2003
 map vnid 1003
!
```

R1

```
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.21/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe1/1
 load-interval 30
 ip address 17.20.20.2/24
!
interface xe1/2
 load-interval 30
 ip address 17.19.19.2/24
!
interface xe6/1
 ip address 17.24.24.2/24
!
interface xe31/1
```

```
 load-interval 30
 ip address 17.22.22.2/24
!
router ospf 100
 bfd all-interfaces
 network 10.10.10.21/32 area 0.0.0.0
 network 17.19.19.0/24 area 0.0.0.0
 network 17.20.20.0/24 area 0.0.0.0
 network 17.22.22.0/24 area 0.0.0.0
 network 17.24.24.0/24 area 0.0.0.0
!
router bgp 65005
 bgp router-id 10.10.10.21
 no bgp inbound-route-filter
 neighbor 10.10.10.5 remote-as 65004
 neighbor 10.10.10.5 update-source lo
 neighbor 10.10.10.5 as-origination-interval 1
 neighbor 10.10.10.5 advertisement-interval 0
 neighbor 10.10.10.6 remote-as 65004
 neighbor 10.10.10.6 update-source lo
 neighbor 10.10.10.6 as-origination-interval 1
 neighbor 10.10.10.6 advertisement-interval 0
 neighbor 10.10.10.7 remote-as 65004
 neighbor 10.10.10.7 update-source lo
 neighbor 10.10.10.7 as-origination-interval 1
 neighbor 10.10.10.7 advertisement-interval 0
 neighbor 10.10.10.22 remote-as 65006
 neighbor 10.10.10.22 update-source lo
 neighbor 10.10.10.22 as-origination-interval 1
 neighbor 10.10.10.22 advertisement-interval 0
 address-family ipv4 unicast
  network 10.10.10.21/32
  max-paths ebgp 4
  redistribute connected
  exit-address-family
  !
 address-family l2vpn evpn
  neighbor 10.10.-10.5 activate
  neighbor 10.10.10.5 allowas-in 1
  neighbor 10.10.10.6 activate
  neighbor 10.10.10.6 allowas-in 1
  neighbor 10.10.10.7 activate
  neighbor 10.10.10.7 allowas-in 1
  neighbor 10.10.10.22 activate
  neighbor 10.10.10.22 allowas-in 1
  exit-address-family
  !
!
```

R2

```
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.22/32 secondary
 ipv6 address ::1/128
```

```
!
interface xe1/1
 load-interval 30
 ip address 17.21.21.2/24
!
interface xe1/2
 load-interval 30
 ip address 17.18.18.2/24
!
interface xe6/1
 ip address 17.24.24.1/24
!
interface xe31/1
 load-interval 30
 ip address 17.23.23.2/24
!
router ospf 100
 bfd all-interfaces
 network 10.10.10.22/32 area 0.0.0.0
 network 17.18.18.0/24 area 0.0.0.0
 network 17.21.21.0/24 area 0.0.0.0
 network 17.23.23.0/24 area 0.0.0.0
 network 17.24.24.0/24 area 0.0.0.0
!
router bgp 65006
 bgp router-id 10.10.10.22
 no bgp inbound-route-filter
 neighbor 10.10.10.5 remote-as 65004
 neighbor 10.10.10.5 update-source lo
 neighbor 10.10.10.5 as-origination-interval 1
 neighbor 10.10.10.5 advertisement-interval 0
 neighbor 10.10.10.6 remote-as 65004
 neighbor 10.10.10.6 update-source lo
 neighbor 10.10.10.6 as-origination-interval 1
 neighbor 10.10.10.6 advertisement-interval 0
 neighbor 10.10.10.7 remote-as 65004
 neighbor 10.10.10.7 update-source lo
 neighbor 10.10.10.7 as-origination-interval 1
 neighbor 10.10.10.7 advertisement-interval 0
 neighbor 10.10.10.21 remote-as 65005
 neighbor 10.10.10.21 update-source lo
 neighbor 10.10.10.21 as-origination-interval 1
 neighbor 10.10.10.21 advertisement-interval 0
 address-family ipv4 unicast
  network 10.10.10.22/32
  redistribute connected
  exit-address-family
  !
 address-family l2vpn evpn
  neighbor 10.10.10.5 activate
  neighbor 10.10.10.5 allowas-in 1
  neighbor 10.10.10.6 activate
  neighbor 10.10.10.6 allowas-in 1
  neighbor 10.10.10.7 activate
  neighbor 10.10.10.7 allowas-in 1
  neighbor 10.10.10.21 activate
```

```
  neighbor 10.10.10.21 allowas-in 1
  exit-address-family
  !
!
```

SPINE1

```
!
interface ce49
 load-interval 30
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.11/32 secondary
 ipv6 address ::1/128
!
interface xe8
 load-interval 30
!
interface xe45
 load-interval 30
!
interface xe47
 load-interval 30
!
router bgp 65002
 bgp router-id 10.10.10.11
 no bgp inbound-route-filter
 bgp unnumbered-mode
  neighbor xe3 as-origination-interval 1
  neighbor xe3 advertisement-interval 0
  neighbor xe3 fall-over bfd
  neighbor ce49 as-origination-interval 1
  neighbor ce49 advertisement-interval 0
  neighbor ce49 fall-over bfd
  neighbor xe8 as-origination-interval 1
  neighbor xe8 advertisement-interval 0
  neighbor xe8 fall-over bfd
  neighbor xe45 as-origination-interval 1
  neighbor xe45 advertisement-interval 0
  neighbor xe45 fall-over bfd
  exit-unnumbered-mode
  !
 address-family ipv4 unicast
  network 10.10.10.11/32
  max-paths ebgp 4
  redistribute connected
  redistribute static
   bgp v4-unnumbered-mode
   neighbor xe3 activate
   neighbor xe3 allowas-in 1
   neighbor ce49 activate
   neighbor ce49 allowas-in 1
   neighbor xe8 activate
   neighbor xe8 allowas-in 1
   neighbor xe45 activate
```

```
   neighbor xe45 allowas-in 1
   exit-v4-unnumbered-mode
   !
  exit-address-family
  !
 address-family l2vpn evpn
  bgp l2vpn-unnumbered-mode
   neighbor xe3 activate
   neighbor xe3 allowas-in 1
   neighbor ce49 activate
   neighbor ce49 allowas-in 1
   neighbor xe8 activate
   neighbor xe8 allowas-in 1
   neighbor xe45 activate
   neighbor xe45 allowas-in 1
   exit-l2vpn-unnumbered-mode
   !
  exit-address-family
  !
!
```

SPINE2

```
!
interface ce50
 load-interval 30
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.12/32 secondary
 ipv6 address ::1/128
!
interface xe1
 load-interval 30
!
interface xe47
 load-interval 30
!
interface xe57
 load-interval 30
!
router bgp 65002
 bgp router-id 10.10.10.12
 no bgp inbound-route-filter
 bgp unnumbered-mode
  neighbor xe57 as-origination-interval 1
  neighbor xe57 advertisement-interval 0
  neighbor xe57 fall-over bfd
  neighbor xe47 as-origination-interval 1
  neighbor xe47 advertisement-interval 0
  neighbor xe47 fall-over bfd
  neighbor ce50 as-origination-interval 1
  neighbor ce50 advertisement-interval 0
  neighbor ce50 fall-over bfd
  neighbor xe1 as-origination-interval 1
  neighbor xe1 advertisement-interval 0
```

```
  neighbor xe1 fall-over bfd
  exit-unnumbered-mode
  !
  address-family ipv4 unicast
   network 10.10.10.12/32
   max-paths ebgp 4
   redistribute connected
   redistribute static
   bgp v4-unnumbered-mode
    neighbor xe57 activate
    neighbor xe57 allowas-in 1
    neighbor xe47 activate
    neighbor xe47 allowas-in 1
    neighbor ce50 activate
    neighbor ce50 allowas-in 1
    neighbor xe1 activate
    neighbor xe1 allowas-in 1
    exit-v4-unnumbered-mode
    !
   exit-address-family
   !
 address-family l2vpn evpn
  bgp l2vpn-unnumbered-mode
   neighbor xe57 activate
   neighbor xe57 allowas-in 1
   neighbor xe47 activate
   neighbor xe47 allowas-in 1
   neighbor ce50 activate
   neighbor ce50 allowas-in 1
   neighbor xe1 activate
   neighbor xe1 allowas-in 1
   exit-l2vpn-unnumbered-mode
   !
  exit-address-family
  !
!
```

SPINE3

```
!
interface ce49
 load-interval 30
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.16/32 secondary
 ipv6 address ::1/128
!
interface xe4
 load-interval 30
!
router bgp 65002
 bgp router-id 10.10.10.16
 no bgp inbound-route-filter
 bgp unnumbered-mode
  neighbor xe7 as-origination-interval 1
```

```
  neighbor xe7 advertisement-interval 0
  neighbor xe7 fall-over bfd
  neighbor ce49 as-origination-interval 1
  neighbor ce49 advertisement-interval 0
  neighbor ce49 fall-over bfd
  exit-unnumbered-mode
  !
 address-family ipv4 unicast
  network 10.10.10.12/32
  max-paths ebgp 4
  redistribute connected
  redistribute static
   bgp v4-unnumbered-mode
   neighbor xe57 activate
   neighbor xe57 allowas-in 1
   neighbor xe47 activate
   neighbor xe47 allowas-in 1
   neighbor ce50 activate
   neighbor ce50 allowas-in 1
   neighbor xe1 activate
   neighbor xe1 allowas-in 1
   exit-v4-unnumbered-mode
   !
  exit-address-family
  !
 address-family l2vpn evpn
  bgp l2vpn-unnumbered-mode
   neighbor xe7 activate
   neighbor xe7 allowas-in 1
   neighbor ce49 activate
   neighbor ce49 allowas-in 1
   exit-l2vpn-unnumbered-mode
   !
  exit-address-family
  !
!
```

SuperSpine1

```
!
interface ce1/3
 load-interval 30
!
interface ce1/4
 load-interval 30
!
interface ce2/1
 load-interval 30
!
interface ce29/1
 load-interval 30
!
interface ce31/1
 load-interval 30
!
interface lo
```

```
 ip address 127.0.0.1/8
 ip address 10.10.10.14/32 secondary
 ipv6 address ::1/128
!
router bgp 65003
 bgp router-id 10.10.10.14
 no bgp inbound-route-filter
 bgp unnumbered-mode
  neighbor ce1/4 as-origination-interval 1
  neighbor ce1/4 advertisement-interval 0
  neighbor ce1/4 fall-over bfd
  neighbor ce31/1 as-origination-interval 1
  neighbor ce31/1 advertisement-interval 0
  neighbor ce31/1 fall-over bfd
  neighbor ce29/1 as-origination-interval 1
  neighbor ce29/1 advertisement-interval 0
  neighbor ce29/1 fall-over bfd
  neighbor ce1/3 as-origination-interval 1
  neighbor ce1/3 advertisement-interval 0
  neighbor ce1/3 fall-over bfd
  exit-unnumbered-mode
  !
 address-family ipv4 unicast
  network 10.10.10.14/32
  max-paths ebgp 4
  redistribute connected
  redistribute static
  bgp v4-unnumbered-mode
   neighbor ce1/4 activate
   neighbor ce1/4 allowas-in 1
   neighbor ce31/1 activate
   neighbor ce31/1 allowas-in 1
   neighbor ce29/1 activate
   neighbor ce29/1 allowas-in 1
   neighbor ce1/3 activate
   neighbor ce1/3 allowas-in 1
   exit-v4-unnumbered-mode
   !
  exit-address-family
  !
 address-family l2vpn evpn
  bgp l2vpn-unnumbered-mode
   neighbor ce1/4 activate
   neighbor ce1/4 allowas-in 1
   neighbor ce31/1 activate
   neighbor ce31/1 allowas-in 1
   neighbor ce29/1 activate
   neighbor ce29/1 allowas-in 1
   neighbor ce1/3 activate
   neighbor ce1/3 allowas-in 1
   exit-l2vpn-unnumbered-mode
   exit-address-family
   !
!
```

SuperSpine2

```
!
interface ce9/1
 load-interval 30
!
interface ce11/1
 load-interval 30
!
interface ce14/1
 load-interval 30
!
interface ce15/1
 load-interval 30
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.10.10.15/32 secondary
 ipv6 address ::1/128
!
router bgp 65003
 bgp router-id 10.10.10.15
 no bgp inbound-route-filter
 bgp unnumbered-mode
  neighbor ce14/1 fall-over bfd
  neighbor ce15/1 fall-over bfd
  neighbor ce9/1 fall-over bfd
  neighbor ce11/1 fall-over bfd
  neighbor ce14/1 as-origination-interval 1
  neighbor ce14/1 advertisement-interval 0
  neighbor ce15/1 as-origination-interval 1
  neighbor ce15/1 advertisement-interval 0
  neighbor ce9/1 as-origination-interval 1
  neighbor ce9/1 advertisement-interval 0
  neighbor ce11/1 as-origination-interval 1
  neighbor ce11/1 advertisement-interval 0
  exit-unnumbered-mode
  !
 address-family ipv4 unicast
  network 10.10.10.15/32
  max-paths ebgp 4
  redistribute connected
  redistribute static
   bgp v4-unnumbered-mode
   neighbor ce14/1 activate
   neighbor ce14/1 allowas-in 1
   neighbor ce15/1 activate
   neighbor ce15/1 allowas-in 1
   neighbor ce9/1 activate
   neighbor ce9/1 allowas-in 1
   neighbor ce11/1 activate
   neighbor ce11/1 allowas-in 1
   exit-v4-unnumbered-mode
   !
  exit-address-family
  !
```

```
 address-family l2vpn evpn
  bgp l2vpn-unnumbered-mode
   neighbor ce14/1 activate
   neighbor ce14/1 allowas-in 1
   neighbor ce15/1 activate
   neighbor ce15/1 allowas-in 1
   neighbor ce9/1 activate
   neighbor ce9/1 allowas-in 1
   neighbor ce11/1 activate
   neighbor ce11/1 allowas-in 1
   exit-l2vpn-unnumbered-mode
   !
  exit-address-family
!
```

SW1

```
!
interface po1
 switchport
 bridge-group 1 spanning-tree disable
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan all
 load-interval 30
!
interface ce53
!
interface ce54
!
interface xe5
 channel-group 1 mode active
!
interface xe10
 channel-group 1 mode active
!
interface xe57
 switchport
 bridge-group 1 spanning-tree disable
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan all
 load-interval 30
!
```

Validation

Note: xw denotes EVPN routes learned from external sites or via a stitched tunnel.

BLEAF1

```
BLEAF1#show nvo vxlan
VXLAN Information
```

```
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI
VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
______________________________________________________________________________
_________________________________________________________________
1000     ----         L3       NW   ----      ----                          --
--      ----      10.10.10.5       10.10.10.1       1444.8fc8.0fc9
1000     ----         L3       NW   ----      ----                          --
--      ----      10.10.10.5       10.10.10.2       9819.2c9b.37d3
2000     ----         L3       NW   ----      ----                          --
--      ----      10.10.10.5       10.10.10.7       1444.8fc9.07c9

Total number of entries are 3
BLEAF1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
====================================================================
10.10.10.5       10.10.10.1       Installed     00:40:21      00:40:21
10.10.10.5       10.10.10.2       Installed     00:40:21      00:40:21
10.10.10.5       10.10.10.7       Installed     00:40:22      00:40:22

Total number of entries are 3
BLEAF1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID                Ip-Addr          Mac-Addr       Type        Age-Out
Retries-Left
______________________________________________________________________________
_________
Total number of entries are 0
BLEAF1#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                              Mac-Addr       Type
Age-Out    Retries-Left
______________________________________________________________________________
______________________
Total number of entries are 0
BLEAF1#show ip bgp summary
BGP router identifier 10.10.10.5, local AS number 65004
BGP table version is 41
7 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                              V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
cd31/1(fe80::230:abff:feea:da1e)       4 65003      3612       2961     40
0      0  00:40:22             6
cd2/1(fe80::36ef:b6ff:fe69:fc46)       4 65003      2873       2507     41
0      0  00:40:22             6

Total number of neighbors 2

Total number of Established sessions 2
```

```
BLEAF1#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.5, local AS number 65004
BGP table version is 269
7 BGP AS-PATH entries
0 BGP community entries

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
cd31/1(fe80::230:abff:feea:da1e)        4 65003      3612      2961    266
0      0  00:40:22            53      8     31      0      0     14
cd2/1(fe80::36ef:b6ff:fe69:fc46)        4 65003      2873      2507    269
0      0  00:40:22            53      8     31      0      0     14
10.10.10.21                             4 65005      3037      2725    269
0      0  00:40:22            11      0      0      0      0     11
10.10.10.22                             4 65006      2989      2727    269
0      0  00:40:22            11      0      0      0      0     11

Total number of neighbors 4

Total number of Established sessions 4
BLEAF1#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::36ef:b6ff:fe69:fc46, cd2/1,
installed 00:40:21, last update 00:40:21 ago
                           [20/0] via fe80::230:abff:feea:da1e, cd31/1
B            10.10.10.2/32 [20/0] via fe80::36ef:b6ff:fe69:fc46, cd2/1,
installed 00:40:21, last update 00:40:21 ago
                           [20/0] via fe80::230:abff:feea:da1e, cd31/1
C            10.10.10.5/32 is directly connected, lo, installed 18:47:39, last
update 18:47:39 ago
O            10.10.10.6/32 [110/3] via 17.20.20.2, xe33, installed 18:46:43,
last update 04:45:36 ago
                           [110/3] via 17.21.21.2, xe34
O            10.10.10.7/32 [110/3] via 17.20.20.2, xe33, installed 18:46:43,
last update 04:45:36 ago
                           [110/3] via 17.21.21.2, xe34
B            10.10.10.11/32 [20/0] via fe80::36ef:b6ff:fe69:fc46, cd2/1,
installed 00:40:22, last update 00:40:22 ago
                            [20/0] via fe80::230:abff:feea:da1e, cd31/1
B            10.10.10.12/32 [20/0] via fe80::36ef:b6ff:fe69:fc46, cd2/1,
installed 00:40:22, last update 00:40:22 ago
                            [20/0] via fe80::230:abff:feea:da1e, cd31/1
B            10.10.10.14/32 [20/0] via fe80::230:abff:feea:da1e, cd31/1,
installed 00:40:21, last update 00:40:21 ago
B            10.10.10.15/32 [20/0] via fe80::36ef:b6ff:fe69:fc46, cd2/1,
installed 00:40:21, last update 00:40:21 ago
O            10.10.10.21/32 [110/2] via 17.20.20.2, xe33, installed 18:46:43,
last update 18:46:43 ago
O            10.10.10.22/32 [110/2] via 17.21.21.2, xe34, installed 04:45:36,
last update 04:45:36 ago
```

```
O            17.18.18.0/24 [110/2] via 17.21.21.2, xe34, installed 04:45:36,
last update 04:45:36 ago
O            17.19.19.0/24 [110/2] via 17.20.20.2, xe33, installed 18:46:43,
last update 18:46:43 ago
C            17.20.20.0/24 is directly connected, xe33, installed 18:47:29,
last update 18:47:29 ago
C            17.21.21.0/24 is directly connected, xe34, installed 04:46:22,
last update 04:46:22 ago
O            17.22.22.0/24 [110/2] via 17.20.20.2, xe33, installed 18:46:43,
last update 18:46:43 ago
O            17.23.23.0/24 [110/2] via 17.21.21.2, xe34, installed 04:45:36,
last update 04:45:36 ago
O            17.24.24.0/24 [110/2] via 17.21.21.2, xe34, installed 18:46:43,
last update 04:45:36 ago
                           [110/2] via 17.20.20.2, xe33
C            127.0.0.0/8 is directly connected, lo, installed 18:48:20, last
update 18:48:20 ago
IP Route Table for VRF "management"
C            10.16.111.0/24 is directly connected, eth0, installed 18:47:34,
last update 18:47:34 ago
C            127.0.0.0/8 is directly connected, lo.management, installed
18:48:20, last update 18:48:20 ago
IP Route Table for VRF "l3_vrf_1"
B            10.10.10.1/32 [0/0] is directly connected, tunvxlan2, installed
00:40:21, last update 00:40:21 ago
B            10.10.10.2/32 [0/0] is directly connected, tunvxlan2, installed
00:40:21, last update 00:40:21 ago
B            10.10.10.7/32 [0/0] is directly connected, tunvxlanw2, installed
00:40:22, last update 00:40:22 ago
B            101.0.0.0/8 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:40:22, last update 00:40:22 ago
B            101.1.1.2/32 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:01:24, last update 00:01:24 ago
B            101.1.1.240/32 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:40:22, last update 00:40:22 ago
B            102.0.0.0/8 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:40:22, last update 00:40:22 ago
B            102.1.1.2/32 [20/0] via 10.10.10.2 (recursive is directly
connected, tunvxlan2), installed 00:11:24, last update 00:11:24 ago
B            103.0.0.0/8 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:40:22, last update 00:40:22 ago
B            103.1.1.2/32 [20/0] via 10.10.10.2 (recursive is directly
connected, tunvxlan2), installed 00:11:24, last update 00:11:24 ago
C            127.0.0.0/8 is directly connected, lo.l3_vrf_1, installed
18:47:53, last update 18:47:53 ago
B            201.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:40:22, last update 00:40:22 ago
B            201.1.1.2/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:02:49, last update 00:02:49 ago
B            201.1.1.240/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:40:22, last update 00:40:22 ago
B            202.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:40:22, last update 00:40:22 ago
B            202.1.1.2/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:02:49, last update 00:02:49 ago
B            203.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:40:22, last update 00:40:22 ago

Gateway of last resort is not set
BLEAF1#show ipv6 route vrf all
IPv6 Routing Table
```

```
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP, v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 18:48:20, last update 18:48:20 ago
C       fe80::/64 via ::, cd1/3, installed 03:55:43, last update 03:55:43 ago
IP Route Table for VRF "management"
C       ::1/128 via ::, lo.management, installed 18:48:20, last update
18:48:20 ago
C       fe80::/64 via ::, eth0, installed 18:48:20, last update 18:48:20 ago
IP Route Table for VRF "l3_vrf_1"
C       ::1/128 via ::, lo.l3_vrf_1, installed 18:47:53, last update 18:47:53
ago
B       ::ffff:a0a:a01/128 [0/0] via ::, tunvxlan2, installed 00:40:21, last
update 00:40:21 ago
B       ::ffff:a0a:a02/128 [0/0] via ::, tunvxlan2, installed 00:40:21, last
update 00:40:21 ago
B       ::ffff:a0a:a07/128 [0/0] via ::, tunvxlanw2, installed 00:40:22, last
update 00:40:22 ago
B       1001::/48 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:40:22, last update 00:40:22 ago
B       1001::2/128 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:01:24, last update 00:01:24 ago
B       1002::/48 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:40:22, last update 00:40:22 ago
B       1002::2/128 [20/0] via ::ffff:a0a:a02 (recursive via ::, tunvxlan2),
installed 00:11:24, last update 00:11:24 ago
B       1003::/48 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:40:22, last update 00:40:22 ago
B       1003::2/128 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:02:09, last update 00:02:09 ago
B       2001::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:40:22, last update 00:40:22 ago
B       2001::2/128 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:02:49, last update 00:02:49 ago
B       2002::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:40:22, last update 00:40:22 ago
B       2002::2/128 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:02:49, last update 00:02:49 ago
B       2003::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:40:22, last update 00:40:22 ago
C       fe80::/64 via ::, tunvxlanw2, installed 18:47:39, last update 18:47:39
ago
B       fe80::1/128 [20/0] via ::ffff:a0a:a02 (recursive via ::, tunvxlan2),
installed 00:11:25, last update 00:11:25 ago
BLEAF1#show bgp vrf all
BGP table version is 1, local router ID is 13.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-
Color

BGP Route Table for VRF l3_vrf_1
```

```
*>x 101.0.0.0        10.10.10.1            0        100        0   65003 65002
65001 ?
* x                  10.10.10.2            0        100        0   65003 65002
65001 ?
*>x 101.1.1.2/32     10.10.10.1            0        100        0   65003 65002
65001 i
*>x 101.1.1.240/32   10.10.10.1            0        100        0   65003 65002
65001 ?
* x                  10.10.10.2            0        100        0   65003 65002
65001 ?
*>x 102.0.0.0        10.10.10.1            0        100        0   65003 65002
65001 ?
* x                  10.10.10.2            0        100        0   65003 65002
65001 ?
*>x 102.1.1.2/32     10.10.10.2            0        100        0   65003 65002
65001 i
*>x 103.0.0.0        10.10.10.1            0        100        0   65003 65002
65001 ?
* x                  10.10.10.2            0        100        0   65003 65002
65001 ?
*>x 103.1.1.2/32     10.10.10.2            0        100        0   65003 65002
65001 i
*>xw201.0.0.0/8      10.10.10.7            0        100        0   65006 65004 ?
*>xw201.1.1.2/32     10.10.10.7            0        100        0   65005 65004 i
*>xw201.1.1.240/32   10.10.10.7            0        100        0   65006 65004 ?
*>xw202.0.0.0/8      10.10.10.7            0        100        0   65006 65004 ?
*>xw202.1.1.2/32     10.10.10.7            0        100        0   65005 65004 i
*>xw203.0.0.0/8      10.10.10.7            0        100        0   65006 65004 ?

Total number of prefixes 13
*>  10.10.10.1/32    fe80::36ef:b6ff:fe69:fc46     0        100        0
65003 65002 65001 i
*                    fe80::230:abff:feea:da1e     0        100        0   65003
65002 65001 i
*>  10.10.10.2/32    fe80::230:abff:feea:da1e     0        100        0   65003
65002 65001 i
*                    fe80::36ef:b6ff:fe69:fc46     0        100        0
65003 65002 65001 i
*>  10.10.10.5/32    0.0.0.0               0        100        32768  i
*                    0.0.0.0               0        100        32768  ?
*>  10.10.10.11/32   fe80::36ef:b6ff:fe69:fc46     0        100        0
65003 65002 i
*                    fe80::230:abff:feea:da1e     0        100        0   65003
65002 i
*>  10.10.10.12/32   fe80::36ef:b6ff:fe69:fc46     0        100        0
65003 65002 i
*                    fe80::230:abff:feea:da1e     0        100        0   65003
65002 i
*>  10.10.10.14/32   fe80::230:abff:feea:da1e     0        100        0   65003
i
*                    fe80::36ef:b6ff:fe69:fc46     0        100        0
65003 65002 65003 i
*>  10.10.10.15/32   fe80::36ef:b6ff:fe69:fc46     0        100        0
65003 i
*                    fe80::230:abff:feea:da1e     0        100        0   65003
65002 65003 i
*>  17.20.20.0/24    0.0.0.0               0        100        32768  ?
*>  17.21.21.0/24    0.0.0.0               0        100        32768  ?

Total number of prefixes 9
BLEAF1#show bgp ipv6 unicast vrf all
```

```
BGP table version is 1, local router ID is 13.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop                    Metric    LocPrf     Weight
Path

BGP Route Table for VRF l3_vrf_1
*>x 1001::/48        ::ffff:10.10.10.1             0        100       0
65003 65002 65001 ?
* x                  ::ffff:10.10.10.2             0        100       0
65003 65002 65001 ?
*>x 1001::2/128      ::ffff:10.10.10.1             0        100       0
65003 65002 65001 i
*>x 1002::/48        ::ffff:10.10.10.1             0        100       0
65003 65002 65001 ?
* x                  ::ffff:10.10.10.2             0        100       0
65003 65002 65001 ?
*>x 1002::2/128      ::ffff:10.10.10.2             0        100       0
65003 65002 65001 i
*>x 1003::/48        ::ffff:10.10.10.1             0        100       0
65003 65002 65001 ?
* x                  ::ffff:10.10.10.2             0        100       0
65003 65002 65001 ?
*>x 1003::2/128      ::ffff:10.10.10.1             0        100       0
65003 65002 65001 i
*>xw2001::/48        ::ffff:10.10.10.7             0        100       0
65006 65004 ?
*>xw2001::2/128      ::ffff:10.10.10.7             0        100       0
65006 65004 i
*>xw2002::/48        ::ffff:10.10.10.7             0        100       0
65006 65004 ?
*>xw2002::2/128      ::ffff:10.10.10.7             0        100       0
65005 65004 i
*>xw2003::/48        ::ffff:10.10.10.7             0        100       0
65006 65004 ?
*>x fe80::1/128      ::ffff:10.10.10.2             0        100       0
65003 65002 65001 i

Total number of prefixes 12
BLEAF1#show bgp l2vpn evpn
BGP table version is 269, local router ID is 10.10.10.5
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight
Path  Peer          Encap
```

```
RD[7001:1001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN

RD[7002:1001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[1000]
                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[1000]
```

```
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1e VXLAN

RD[7007:7001]
*>  [5]:[0]:[0]:[8]:[201.0.0.0]:[0.0.0.0]:[2000]
                        10.10.10.7           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.7           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[8]:[202.0.0.0]:[0.0.0.0]:[2000]
                        10.10.10.7           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.7           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[8]:[203.0.0.0]:[0.0.0.0]:[2000]
                        10.10.10.7           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.7           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.2]:[0.0.0.0]:[2000]
                        10.10.10.7           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                       10.10.10.7           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.240]:[0.0.0.0]:[2000]
                        10.10.10.7           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.7           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[32]:[202.1.1.2]:[0.0.0.0]:[2000]
                        10.10.10.7           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                       10.10.10.7           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[48]:[2001::]:[::]:[2000]
                        10.10.10.7           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.7           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[48]:[2002::]:[::]:[2000]
                        10.10.10.7           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.7           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[48]:[2003::]:[::]:[2000]
                        10.10.10.7           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.7           0        100        0   65005 65004
?  10.10.10.21     VXLAN
```

```
*>  [5]:[0]:[0]:[128]:[2001::2]:[::]:[2000]
                      10.10.10.7           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*                     10.10.10.7           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[128]:[2002::2]:[::]:[2000]
                      10.10.10.7           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                     10.10.10.7           0        100        0   65006 65004
i  10.10.10.22     VXLAN

RD[10.10.10.1:1]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[1001]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[32,101.1.1.1]:[1001]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,1001::1][1001]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1001]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[0]:[1001]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[32,101.1.1.2]:
[1001]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,1001::2][1
001]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN

RD[10.10.10.1:2]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[1002]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[32,102.1.1.1]:[1002]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
```

```
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,1002::1][1002]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1002]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN

RD[10.10.10.1:3]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[1003]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[32,103.1.1.1]:[1003]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,1003::1][1003]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1003]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[0]:[1003]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1003::2][1
003]:[1000]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN

RD[10.10.10.1:64512]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN

RD[10.10.10.2:1]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[1001]
                      10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[32,101.1.1.1]:[1001]
```

```
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,1001::1][1001]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1001]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[0]:[1001]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,fe80::1][1
001]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN

RD[10.10.10.2:2]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[1002]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[32,102.1.1.1]:[1002]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,1002::1][1002]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1002]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[0]:[1002]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[32,102.1.1.2]:
[1002]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
```

```
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[128,1002::2][1
002]:[1000]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN

RD[10.10.10.2:3]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[1003]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[32,103.1.1.1]:[1003]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,1003::1][1003]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1003]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[0]:[1003]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[32,103.1.1.2]:
[1003]:[1000]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1002::2][1
003]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN

RD[10.10.10.2:64512]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc46 VXLAN
*                    10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1e VXLAN

Total number of prefixes 64
BLEAF1#show ip bgp 101.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 101.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
```

```
  AS path: 65003 65002 65001
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.1 (IGP metric 0) from fe80::36ef:b6ff:fe69:fc46 (Remote
Id:10.10.10.15) Peer nexthop: fe80::36ef:b6ff:fe69:fc46
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External,
imported)
      Extended Community: RT:101:101 1001:1001  Encapsulation:VxLAN
EVPN_Router's_Mac:1444.8fc8.0fc9  SubType-0x0b:000000000065
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 08:05:46 2025,
00:02:12 ago

BLEAF1#show ip bgp 102.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 102.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65003 65002 65001
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.2 (IGP metric 0) from fe80::36ef:b6ff:fe69:fc46 (Remote
Id:10.10.10.15) Peer nexthop: fe80::36ef:b6ff:fe69:fc46
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External,
imported)
      Extended Community: RT:102:102 1001:1001  Encapsulation:VxLAN
EVPN_Router's_Mac:9819.2c9b.37d3  SubType-0x0b:000000000066
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 07:56:31 2025,
00:11:27 ago

BLEAF1#show ip bgp 201.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 201.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65005 65004
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.7 (IGP metric 0) from 10.10.10.21 (Remote Id:10.10.10.21)
Peer nexthop: 10.10.10.21
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Stitched: (Recieved from the external fabric peer)
      Extended Community: SOO:1005:1005  RT:7001:7001  Encapsulation:VxLAN
EVPN_Router's_Mac:1444.8fc9.07c9
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 08:05:06 2025,
00:02:53 ago

BLEAF1#show ip bgp 202.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 202.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65005 65004
  Path Selection reason: Nothing left to compare
```

```
   Nexthop:10.10.10.7 (IGP metric 0) from 10.10.10.21 (Remote Id:10.10.10.21)
Peer nexthop: 10.10.10.21
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Stitched: (Recieved from the external fabric peer)
      Extended Community: SOO:1005:1005  RT:7001:7001  Encapsulation:VxLAN
EVPN_Router's_Mac:1444.8fc9.07c9
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 08:05:06 2025,
00:02:53 ago

BLEAF1#show evpn l3vni-map
 L3VNI        Trans-L3VNI  L2VNI        IRB-interface
==================================================
 1000         2000         ---          ---

BLEAF1#show nvo vxlan mac-table
================================================================================
================================================================================
=====
                                                          VXLAN MAC Entries
================================================================================
================================================================================
=====
VNID    Interface VlanId   In-VlanId Mac-Addr    VTEP-Ip/ESI
Type            Status     MAC move AccessPortDesc LeafFlag
________________________________________________________________________________
________________________________________________________________________________
_____


Total number of entries are : 0

BLEAF1#
BLEAF1#
```

BLEAF2

```
BLEAF2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI
VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
________________________________________________________________________________
1000     ----         L3       NW   ----      ----                          --
--       ----      10.10.10.6       10.10.10.1       1444.8fc8.0fc9
1000     ----         L3       NW   ----      ----                          --
--       ----      10.10.10.6       10.10.10.2       9819.2c9b.37d3
2000     ----         L3       NW   ----      ----                          --
--       ----      10.10.10.6       10.10.10.7       1444.8fc9.07c9

Total number of entries are 3
```

```
BLEAF2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
10.10.10.6       10.10.10.7       Installed     00:40:31      00:40:31
10.10.10.6       10.10.10.1       Installed     00:40:30      00:40:30
10.10.10.6       10.10.10.2       Installed     00:40:30      00:40:30

Total number of entries are 3
BLEAF2#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID                Ip-Addr          Mac-Addr       Type        Age-Out
Retries-Left
_____________________________________________________________________________
_________
Total number of entries are 0
BLEAF2#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr                                  Mac-Addr       Type
Age-Out    Retries-Left
_____________________________________________________________________________
_____________________
Total number of entries are 0
BLEAF2#show ip bgp summary
BGP router identifier 10.10.10.6, local AS number 65004
BGP table version is 57
7 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
ce52(fe80::36ef:b6ff:fe69:fc5e)         4 65003      2848       2461     56
0       0  00:40:32            6
ce53(fe80::230:abff:feea:da1a)          4 65003      3577       2960     57
0       0  00:40:32            6

Total number of neighbors 2

Total number of Established sessions 2
BLEAF2#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.6, local AS number 65004
BGP table version is 306
7 BGP AS-PATH entries
0 BGP community entries

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
ce52(fe80::36ef:b6ff:fe69:fc5e)         4 65003      2848       2461     305
0       0  00:40:32            53      8     31      0       0     14
ce53(fe80::230:abff:feea:da1a)          4 65003      3577       2960     306
0       0  00:40:32            53      8     31      0       0     14
10.10.10.21                             4 65005      2942       2842     301
0       0  00:40:32            11      0      0      0       0     11
10.10.10.22                             4 65006      2928       2865     301
0       0  00:40:32            11      0      0      0       0     11
```

```
Total number of neighbors 4

Total number of Established sessions 4
BLEAF2#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::36ef:b6ff:fe69:fc5e, ce52,
installed 00:40:30, last update 00:40:30 ago
                           [20/0] via fe80::230:abff:feea:da1a, ce53
B            10.10.10.2/32 [20/0] via fe80::36ef:b6ff:fe69:fc5e, ce52,
installed 00:40:30, last update 00:40:30 ago
                           [20/0] via fe80::230:abff:feea:da1a, ce53
O            10.10.10.5/32 [110/3] via 17.22.22.2, ce50/1, installed 18:46:52,
last update 04:45:45 ago
                           [110/3] via 17.23.23.2, ce54/1
C            10.10.10.6/32 is directly connected, lo, installed 18:50:03, last
update 18:50:03 ago
O            10.10.10.7/32 [110/3] via 17.22.22.2, ce50/1, installed 18:49:08,
last update 18:49:03 ago
                           [110/3] via 17.23.23.2, ce54/1
B            10.10.10.11/32 [20/0] via fe80::36ef:b6ff:fe69:fc5e, ce52,
installed 00:40:31, last update 00:40:31 ago
                            [20/0] via fe80::230:abff:feea:da1a, ce53
B            10.10.10.12/32 [20/0] via fe80::36ef:b6ff:fe69:fc5e, ce52,
installed 00:40:31, last update 00:40:31 ago
                            [20/0] via fe80::230:abff:feea:da1a, ce53
B            10.10.10.14/32 [20/0] via fe80::230:abff:feea:da1a, ce53,
installed 00:40:30, last update 00:40:30 ago
B            10.10.10.15/32 [20/0] via fe80::36ef:b6ff:fe69:fc5e, ce52,
installed 00:40:30, last update 00:40:30 ago
O            10.10.10.21/32 [110/2] via 17.22.22.2, ce50/1, installed
18:49:08, last update 18:49:08 ago
O            10.10.10.22/32 [110/2] via 17.23.23.2, ce54/1, installed
18:49:03, last update 18:49:03 ago
O            17.18.18.0/24 [110/2] via 17.23.23.2, ce54/1, installed 18:49:03,
last update 18:49:03 ago
O            17.19.19.0/24 [110/2] via 17.22.22.2, ce50/1, installed 18:49:08,
last update 18:49:08 ago
O            17.20.20.0/24 [110/2] via 17.22.22.2, ce50/1, installed 18:47:37,
last update 18:47:37 ago
O            17.21.21.0/24 [110/2] via 17.23.23.2, ce54/1, installed 04:46:30,
last update 04:46:30 ago
C            17.22.22.0/24 is directly connected, ce50/1, installed 18:49:58,
last update 18:49:58 ago
C            17.23.23.0/24 is directly connected, ce54/1, installed 18:49:58,
last update 18:49:58 ago
O            17.24.24.0/24 [110/2] via 17.23.23.2, ce54/1, installed 18:49:08,
last update 18:49:03 ago
                           [110/2] via 17.22.22.2, ce50/1
C            127.0.0.0/8 is directly connected, lo, installed 18:50:29, last
update 18:50:29 ago
```

```
IP Route Table for VRF "management"
C            10.16.104.0/24 is directly connected, eth0, installed 18:50:01,
last update 18:50:01 ago
C            127.0.0.0/8 is directly connected, lo.management, installed
18:50:29, last update 18:50:29 ago
IP Route Table for VRF "l3_vrf_1"
B            10.10.10.1/32 [0/0] is directly connected, tunvxlan2, installed
00:40:30, last update 00:40:30 ago
B            10.10.10.2/32 [0/0] is directly connected, tunvxlan2, installed
00:40:30, last update 00:40:30 ago
B            10.10.10.7/32 [0/0] is directly connected, tunvxlanw2, installed
00:40:31, last update 00:40:31 ago
B            101.0.0.0/8 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:40:31, last update 00:40:31 ago
B            101.1.1.2/32 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:01:33, last update 00:01:33 ago
B            101.1.1.240/32 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:40:31, last update 00:40:31 ago
B            102.0.0.0/8 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:40:31, last update 00:40:31 ago
B            102.1.1.2/32 [20/0] via 10.10.10.2 (recursive is directly
connected, tunvxlan2), installed 00:11:34, last update 00:11:34 ago
B            103.0.0.0/8 [20/0] via 10.10.10.1 (recursive is directly
connected, tunvxlan2), installed 00:40:31, last update 00:40:31 ago
B            103.1.1.2/32 [20/0] via 10.10.10.2 (recursive is directly
connected, tunvxlan2), installed 00:11:34, last update 00:11:34 ago
C            127.0.0.0/8 is directly connected, lo.l3_vrf_1, installed
18:50:08, last update 18:50:08 ago
B            201.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:40:31, last update 00:40:31 ago
B            201.1.1.2/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:02:58, last update 00:02:58 ago
B            201.1.1.240/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:40:31, last update 00:40:31 ago
B            202.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:40:31, last update 00:40:31 ago
B            202.1.1.2/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:02:58, last update 00:02:58 ago
B            203.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlanw2), installed 00:40:31, last update 00:40:31 ago

Gateway of last resort is not set
BLEAF2#show ipv6 route vrf all
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 18:50:29, last update 18:50:29 ago
C       fe80::/64 via ::, xe11, installed 01:42:51, last update 01:42:51 ago
IP Route Table for VRF "management"
C       ::1/128 via ::, lo.management, installed 18:50:29, last update
18:50:29 ago
C       fe80::/64 via ::, eth0, installed 18:50:29, last update 18:50:29 ago
IP Route Table for VRF "l3_vrf_1"
```

```
C       ::1/128 via ::, lo.l3_vrf_1, installed 18:50:08, last update 18:50:08
ago
B       ::ffff:a0a:a01/128 [0/0] via ::, tunvxlan2, installed 00:40:30, last
update 00:40:30 ago
B       ::ffff:a0a:a02/128 [0/0] via ::, tunvxlan2, installed 00:40:30, last
update 00:40:30 ago
B       ::ffff:a0a:a07/128 [0/0] via ::, tunvxlanw2, installed 00:40:31, last
update 00:40:31 ago
B       1001::/48 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:40:31, last update 00:40:31 ago
B       1001::2/128 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:01:33, last update 00:01:33 ago
B       1002::/48 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:40:31, last update 00:40:31 ago
B       1002::2/128 [20/0] via ::ffff:a0a:a02 (recursive via ::, tunvxlan2),
installed 00:11:34, last update 00:11:34 ago
B       1003::/48 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:40:31, last update 00:40:31 ago
B       1003::2/128 [20/0] via ::ffff:a0a:a01 (recursive via ::, tunvxlan2),
installed 00:02:18, last update 00:02:18 ago
B       2001::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:40:31, last update 00:40:31 ago
B       2001::2/128 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:02:58, last update 00:02:58 ago
B       2002::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:40:31, last update 00:40:31 ago
B       2002::2/128 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:02:58, last update 00:02:58 ago
B       2003::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlanw2),
installed 00:40:31, last update 00:40:31 ago
C       fe80::/64 via ::, tunvxlanw2, installed 18:50:02, last update 18:50:02
ago
B       fe80::1/128 [20/0] via ::ffff:a0a:a02 (recursive via ::, tunvxlan2),
installed 00:11:34, last update 00:11:34 ago
BLEAF2#show bgp vrf all
BGP table version is 1, local router ID is 33.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-
Color

BGP Route Table for VRF l3_vrf_1
*>x 101.0.0.0        10.10.10.1           0        100       0   65003 65002
65001 ?
* x                  10.10.10.2           0        100       0   65003 65002
65001 ?
*>x 101.1.1.2/32     10.10.10.1           0        100       0   65003 65002
65001 i
*>x 101.1.1.240/32   10.10.10.1           0        100       0   65003 65002
65001 ?
* x                  10.10.10.2           0        100       0   65003 65002
65001 ?
*>x 102.0.0.0        10.10.10.1           0        100       0   65003 65002
65001 ?
* x                  10.10.10.2           0        100       0   65003 65002
65001 ?
```

```
*>x 102.1.1.2/32     10.10.10.2           0        100        0   65003 65002
65001 i
*>x 103.0.0.0        10.10.10.1           0        100        0   65003 65002
65001 ?
* x                  10.10.10.2           0        100        0   65003 65002
65001 ?
*>x 103.1.1.2/32     10.10.10.2           0        100        0   65003 65002
65001 i
*>xw201.0.0.0/8      10.10.10.7           0        100        0   65006 65004 ?
*>xw201.1.1.2/32     10.10.10.7           0        100        0   65005 65004 i
*>xw201.1.1.240/32   10.10.10.7           0        100        0   65006 65004 ?
*>xw202.0.0.0/8      10.10.10.7           0        100        0   65006 65004 ?
*>xw202.1.1.2/32     10.10.10.7           0        100        0   65005 65004 i
*>xw203.0.0.0/8      10.10.10.7           0        100        0   65006 65004 ?

Total number of prefixes 13
*>  10.10.10.1/32    fe80::36ef:b6ff:fe69:fc5e     0        100        0
65003 65002 65001 i
*                     fe80::230:abff:feea:da1a     0        100        0   65003
65002 65001 i
*>  10.10.10.2/32    fe80::230:abff:feea:da1a     0        100        0   65003
65002 65001 i
*                     fe80::36ef:b6ff:fe69:fc5e     0        100        0
65003 65002 65001 i
*>  10.10.10.6/32    0.0.0.0              0        100        32768  i
*                     0.0.0.0              0        100        32768  ?
*>  10.10.10.11/32   fe80::36ef:b6ff:fe69:fc5e     0        100        0
65003 65002 i
*                     fe80::230:abff:feea:da1a     0        100        0   65003
65002 i
*>  10.10.10.12/32   fe80::36ef:b6ff:fe69:fc5e     0        100        0
65003 65002 i
*                     fe80::230:abff:feea:da1a     0        100        0   65003
65002 i
*>  10.10.10.14/32   fe80::230:abff:feea:da1a     0        100        0   65003
i
*                     fe80::36ef:b6ff:fe69:fc5e     0        100        0
65003 65002 65003 i
*>  10.10.10.15/32   fe80::36ef:b6ff:fe69:fc5e     0        100        0
65003 i
*                     fe80::230:abff:feea:da1a     0        100        0   65003
65002 65003 i
*>  17.22.22.0/24    0.0.0.0              0        100        32768  ?
*>  17.23.23.0/24    0.0.0.0              0        100        32768  ?

Total number of prefixes 9
BLEAF2#show bgp ipv6 unicast vrf all
BGP table version is 1, local router ID is 33.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop                    Metric    LocPrf     Weight
Path

BGP Route Table for VRF l3_vrf_1
*>x 1001::/48        ::ffff:10.10.10.1            0        100        0
65003 65002 65001 ?
```

```
* x                  ::ffff:10.10.10.2             0        100       0
65003 65002 65001 ?
*>x 1001::2/128      ::ffff:10.10.10.1             0        100       0
65003 65002 65001 i
*>x 1002::/48        ::ffff:10.10.10.1             0        100       0
65003 65002 65001 ?
* x                  ::ffff:10.10.10.2             0        100       0
65003 65002 65001 ?
*>x 1002::2/128      ::ffff:10.10.10.2             0        100       0
65003 65002 65001 i
*>x 1003::/48        ::ffff:10.10.10.1             0        100       0
65003 65002 65001 ?
* x                  ::ffff:10.10.10.2             0        100       0
65003 65002 65001 ?
*>x 1003::2/128      ::ffff:10.10.10.1             0        100       0
65003 65002 65001 i
*>xw2001::/48        ::ffff:10.10.10.7             0        100       0
65005 65004 ?
*>xw2001::2/128      ::ffff:10.10.10.7             0        100       0
65006 65004 i
*>xw2002::/48        ::ffff:10.10.10.7             0        100       0
65005 65004 ?
*>xw2002::2/128      ::ffff:10.10.10.7             0        100       0
65005 65004 i
*>xw2003::/48        ::ffff:10.10.10.7             0        100       0
65005 65004 ?
*>x fe80::1/128      ::ffff:10.10.10.2             0        100       0
65003 65002 65001 i

Total number of prefixes 12
BLEAF2#show bgp l2vpn evpn
BGP table version is 306, local router ID is 10.10.10.6
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight
Path  Peer          Encap

RD[7001:1001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.1           0        100       0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100       0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.1           0        100       0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100       0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[1000]
```

```
                        10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[1000]
                        10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[1000]
                        10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[1000]
                        10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[1000]
                        10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.1           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN

RD[7002:1001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[1000]
                        10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                       10.10.10.2           0        100        0   65003 65002
65001 ?  fe80::230:abff:feea:da1a VXLAN
```

```
RD[7007:7001]
*>  [5]:[0]:[0]:[8]:[201.0.0.0]:[0.0.0.0]:[2000]
                      10.10.10.7            0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                     10.10.10.7            0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[8]:[202.0.0.0]:[0.0.0.0]:[2000]
                      10.10.10.7            0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                     10.10.10.7            0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[8]:[203.0.0.0]:[0.0.0.0]:[2000]
                      10.10.10.7            0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                     10.10.10.7            0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.2]:[0.0.0.0]:[2000]
                      10.10.10.7            0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                     10.10.10.7            0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.240]:[0.0.0.0]:[2000]
                      10.10.10.7            0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                     10.10.10.7            0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[32]:[202.1.1.2]:[0.0.0.0]:[2000]
                      10.10.10.7            0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                     10.10.10.7            0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[48]:[2001::]:[::]:[2000]
                      10.10.10.7            0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.7            0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[48]:[2002::]:[::]:[2000]
                      10.10.10.7            0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.7            0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[48]:[2003::]:[::]:[2000]
                      10.10.10.7            0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.7            0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[128]:[2001::2]:[::]:[2000]
                      10.10.10.7            0        100        0   65006 65004
i  10.10.10.22     VXLAN
*                     10.10.10.7            0        100        0   65005 65004
i  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[128]:[2002::2]:[::]:[2000]
                      10.10.10.7            0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                     10.10.10.7            0        100        0   65006 65004
i  10.10.10.22     VXLAN

RD[10.10.10.1:1]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[1001]
```

```
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[32,101.1.1.1]:[1001]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,1001::1][1001]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1001]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[0]:[1001]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[32,101.1.1.2]:
[1001]:[1000]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,1001::2][1
001]:[1000]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN

RD[10.10.10.1:2]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[1002]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[32,102.1.1.1]:[1002]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,1002::1][1002]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1002]
                        10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
```

```
RD[10.10.10.1:3]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[1003]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[32,103.1.1.1]:[1003]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,1003::1][1003]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1003]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[0]:[1003]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1003::2][1
003]:[1000]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN

RD[10.10.10.1:64512]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                       10.10.10.1           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.1           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN

RD[10.10.10.2:1]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[1001]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[32,101.1.1.1]:[1001]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,1001::1][1001]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1001]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
```

```
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[0]:[1001]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,fe80::1][1
001]:[1000]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN

RD[10.10.10.2:2]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[1002]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[32,102.1.1.1]:[1002]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,1002::1][1002]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1002]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[0]:[1002]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[32,102.1.1.2]:
[1002]:[1000]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[128,1002::2][1
002]:[1000]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                     10.10.10.2           0        100        0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN

RD[10.10.10.2:3]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[1003]
                       10.10.10.2           0        100        0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
```

```
*                      10.10.10.2           0        100       0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[32,103.1.1.1]:[1003]
                       10.10.10.2           0        100       0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.2           0        100       0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,1003::1][1003]
                       10.10.10.2           0        100       0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.2           0        100       0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1003]
                       10.10.10.2           0        100       0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.2           0        100       0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[0]:[1003]
                       10.10.10.2           0        100       0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.2           0        100       0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[32,103.1.1.2]:
[1003]:[1000]
                       10.10.10.2           0        100       0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.2           0        100       0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1002::2][1
003]
                       10.10.10.2           0        100       0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.2           0        100       0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN

RD[10.10.10.2:64512]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                       10.10.10.2           0        100       0   65003 65002
65001 i  fe80::36ef:b6ff:fe69:fc5e VXLAN
*                      10.10.10.2           0        100       0   65003 65002
65001 i  fe80::230:abff:feea:da1a VXLAN

Total number of prefixes 64
BLEAF2#show ip bgp 101.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 101.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65003 65002 65001
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.1 (IGP metric 0) from fe80::36ef:b6ff:fe69:fc5e (Remote
Id:10.10.10.15) Peer nexthop: fe80::36ef:b6ff:fe69:fc5e
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External,
imported)
      Extended Community: RT:101:101 1001:1001  Encapsulation:VxLAN
EVPN_Router's_Mac:1444.8fc8.0fc9  SubType-0x0b:000000000065
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
```

```
      Origin-AS validity: disabled      Last update: Mon Dec 15 08:05:46 2025,
00:02:21 ago

BLEAF2#show ip bgp 102.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 102.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65003 65002 65001
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.2 (IGP metric 0) from fe80::36ef:b6ff:fe69:fc5e (Remote
Id:10.10.10.15) Peer nexthop: fe80::36ef:b6ff:fe69:fc5e
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External,
imported)
      Extended Community: RT:102:102 1001:1001  Encapsulation:VxLAN
EVPN_Router's_Mac:9819.2c9b.37d3  SubType-0x0b:000000000066
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 07:56:31 2025,
00:11:36 ago

BLEAF2#show ip bgp 201.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 201.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65005 65004
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.7 (IGP metric 0) from 10.10.10.21 (Remote Id:10.10.10.21)
Peer nexthop: 10.10.10.21
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Stitched: (Recieved from the external fabric peer)
      Extended Community: SOO:1005:1005  RT:7001:7001  Encapsulation:VxLAN
EVPN_Router's_Mac:1444.8fc9.07c9
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 08:05:06 2025,
00:03:01 ago

BLEAF2#show ip bgp 202.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 202.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65005 65004
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.7 (IGP metric 0) from 10.10.10.21 (Remote Id:10.10.10.21)
Peer nexthop: 10.10.10.21
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Stitched: (Recieved from the external fabric peer)
      Extended Community: SOO:1005:1005  RT:7001:7001  Encapsulation:VxLAN
EVPN_Router's_Mac:1444.8fc9.07c9
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 08:05:06 2025,
00:03:01 ago

BLEAF2#show evpn l3vni-map
```

```
 L3VNI        Trans-L3VNI  L2VNI        IRB-interface
=================================================
 1000         2000         ---          ---

BLEAF2#show nvo vxlan mac-table
================================================================================
================================================================================
=====
                                                          VXLAN MAC Entries
================================================================================
================================================================================
=====
VNID    Interface VlanId   In-VlanId Mac-Addr    VTEP-Ip/ESI
Type            Status     MAC move AccessPortDesc LeafFlag
________________________________________________________________________________
________________________________________________________________________________
_____


Total number of entries are : 0

BLEAF2#
BLEAF2#
BLEAF2#
```

BLEAF3

```
BLEAF3#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI
VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
________________________________________________________________________________
________________________________________                                        --
1000     ----         L3       NW   ----       ----
--       ----      10.10.10.7       10.10.10.4       e001.a6a9.4b01
2000     ----         L3       NW   ----       ----                             --
--       ----      10.10.10.7       10.10.10.5       5c17.83ff.2255
2000     ----         L3       NW   ----       ----                             --
--       ----      10.10.10.7       10.10.10.6       7c8d.9caa.5231

Total number of entries are 3
BLEAF3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
10.10.10.7       10.10.10.5       Installed     00:40:41      00:40:41
10.10.10.7       10.10.10.4       Installed     00:40:41      00:40:41
10.10.10.7       10.10.10.6       Installed     00:40:42      00:40:42

Total number of entries are 3
BLEAF3#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
```

```
============================
VNID                Ip-Addr          Mac-Addr       Type        Age-Out
Retries-Left
_____________________________________________________________________________
_________
Total number of entries are 0
BLEAF3#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr                                  Mac-Addr       Type
Age-Out    Retries-Left
_____________________________________________________________________________
_____________________
Total number of entries are 0
BLEAF3#show ip bgp summary
BGP router identifier 10.10.10.7, local AS number 65004
BGP table version is 31
5 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe7(fe80::6eb9:c5ff:fe15:8765)          4 65002      2790      2771     28
0      0  00:40:43            2

Total number of neighbors 1

Total number of Established sessions 1
BLEAF3#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.7, local AS number 65004
BGP table version is 189
5 BGP AS-PATH entries
0 BGP community entries

Neighbor                                   V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
xe7(fe80::6eb9:c5ff:fe15:8765)          4 65002      2790      2771     181
0      0  00:40:43            29      0     19      3      0       7
10.10.10.21                            4 65005      2975      2809     188
0      0  00:40:43            22      0      0      0      0      22
10.10.10.22                            4 65006      2970      2808     188
0      0  00:40:43            22      0      0      0      0      22

Total number of neighbors 3

Total number of Established sessions 3
BLEAF3#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
```

```
B            10.10.10.4/32 [20/0] via fe80::6eb9:c5ff:fe15:8765, xe7,
installed 00:40:41, last update 00:40:41 ago
O            10.10.10.5/32 [110/3] via 17.18.18.2, xe3, installed 18:47:03,
last update 04:45:56 ago
                           [110/3] via 17.19.19.2, xe2
O            10.10.10.6/32 [110/3] via 17.18.18.2, xe3, installed 18:49:24,
last update 18:49:14 ago
                           [110/3] via 17.19.19.2, xe2
C            10.10.10.7/32 is directly connected, lo, installed 18:50:54, last
update 18:50:54 ago
B            10.10.10.16/32 [20/0] via fe80::6eb9:c5ff:fe15:8765, xe7,
installed 00:40:41, last update 00:40:41 ago
O            10.10.10.21/32 [110/2] via 17.19.19.2, xe2, installed 18:49:55,
last update 18:49:55 ago
O            10.10.10.22/32 [110/2] via 17.18.18.2, xe3, installed 18:50:00,
last update 18:50:00 ago
C            17.18.18.0/24 is directly connected, xe3, installed 18:50:51,
last update 18:50:51 ago
C            17.19.19.0/24 is directly connected, xe2, installed 18:50:51,
last update 18:50:51 ago
O            17.20.20.0/24 [110/2] via 17.19.19.2, xe2, installed 18:47:48,
last update 18:47:48 ago
O            17.21.21.0/24 [110/2] via 17.18.18.2, xe3, installed 04:46:41,
last update 04:46:41 ago
O            17.22.22.0/24 [110/2] via 17.19.19.2, xe2, installed 18:49:55,
last update 18:49:55 ago
O            17.23.23.0/24 [110/2] via 17.18.18.2, xe3, installed 18:50:00,
last update 18:50:00 ago
O            17.24.24.0/24 [110/2] via 17.19.19.2, xe2, installed 18:50:00,
last update 18:49:55 ago
                           [110/2] via 17.18.18.2, xe3
C            127.0.0.0/8 is directly connected, lo, installed 18:51:27, last
update 18:51:27 ago
IP Route Table for VRF "management"
C            10.16.112.0/24 is directly connected, eth0, installed 18:50:49,
last update 18:50:49 ago
C            127.0.0.0/8 is directly connected, lo.management, installed
18:51:27, last update 18:51:27 ago
IP Route Table for VRF "l3_vrf_1"
B            10.10.10.4/32 [0/0] is directly connected, tunvxlan2, installed
00:40:41, last update 00:40:41 ago
B            10.10.10.5/32 [0/0] is directly connected, tunvxlanw2, installed
00:40:41, last update 00:40:41 ago
B            10.10.10.6/32 [0/0] is directly connected, tunvxlanw2, installed
00:40:42, last update 00:40:42 ago
B            101.0.0.0/8 [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlanw2), installed 00:40:42, last update 00:40:42 ago
B            101.1.1.2/32 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlanw2), installed 00:11:45, last update 00:11:45 ago
B            101.1.1.240/32 [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlanw2), installed 00:40:42, last update 00:40:42 ago
B            102.0.0.0/8 [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlanw2), installed 00:40:42, last update 00:40:42 ago
B            102.1.1.2/32 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlanw2), installed 00:11:45, last update 00:11:45 ago
B            103.0.0.0/8 [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlanw2), installed 00:40:42, last update 00:40:42 ago
C            127.0.0.0/8 is directly connected, lo.l3_vrf_1, installed
18:51:00, last update 18:51:00 ago
B            201.0.0.0/8 [20/0] via 10.10.10.4 (recursive is directly
connected, tunvxlan2), installed 00:40:42, last update 00:40:42 ago
```

```
B            201.1.1.2/32 [20/0] via 10.10.10.4 (recursive is directly
connected, tunvxlan2), installed 00:03:09, last update 00:03:09 ago
B            201.1.1.240/32 [20/0] via 10.10.10.4 (recursive is directly
connected, tunvxlan2), installed 00:40:42, last update 00:40:42 ago
B            202.0.0.0/8 [20/0] via 10.10.10.4 (recursive is directly
connected, tunvxlan2), installed 00:40:42, last update 00:40:42 ago
B            202.1.1.2/32 [20/0] via 10.10.10.4 (recursive is directly
connected, tunvxlan2), installed 00:03:09, last update 00:03:09 ago
B            203.0.0.0/8 [20/0] via 10.10.10.4 (recursive is directly
connected, tunvxlan2), installed 00:40:42, last update 00:40:42 ago
B            203.1.1.2/32 [20/0] via 10.10.10.4 (recursive is directly
connected, tunvxlan2), installed 00:03:09, last update 00:03:09 ago

Gateway of last resort is not set
BLEAF3#show ipv6 route vrf all
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP, v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 18:51:27, last update 18:51:27 ago
C       fe80::/64 via ::, xe1, installed 05:55:08, last update 05:55:08 ago
IP Route Table for VRF "management"
C       ::1/128 via ::, lo.management, installed 18:51:27, last update
18:51:27 ago
C       fe80::/64 via ::, eth0, installed 18:51:27, last update 18:51:27 ago
IP Route Table for VRF "l3_vrf_1"
C       ::1/128 via ::, lo.l3_vrf_1, installed 18:51:00, last update 18:51:00
ago
B       ::ffff:a0a:a04/128 [0/0] via ::, tunvxlan2, installed 00:40:41, last
update 00:40:41 ago
B       ::ffff:a0a:a05/128 [0/0] via ::, tunvxlanw2, installed 00:40:41, last
update 00:40:41 ago
B       ::ffff:a0a:a06/128 [0/0] via ::, tunvxlanw2, installed 00:40:42, last
update 00:40:42 ago
B       1001::/48 [20/0] via ::ffff:a0a:a06 (recursive via ::, tunvxlanw2),
installed 00:40:42, last update 00:40:42 ago
B       1001::2/128 [20/0] via ::ffff:a0a:a05 (recursive via ::, tunvxlanw2),
installed 00:11:45, last update 00:11:45 ago
B       1002::/48 [20/0] via ::ffff:a0a:a06 (recursive via ::, tunvxlanw2),
installed 00:40:42, last update 00:40:42 ago
B       1002::2/128 [20/0] via ::ffff:a0a:a05 (recursive via ::, tunvxlanw2),
installed 00:11:45, last update 00:11:45 ago
B       1003::/48 [20/0] via ::ffff:a0a:a06 (recursive via ::, tunvxlanw2),
installed 00:40:42, last update 00:40:42 ago
B       2001::/48 [20/0] via ::ffff:a0a:a04 (recursive via ::, tunvxlan2),
installed 00:40:42, last update 00:40:42 ago
B       2001::2/128 [20/0] via ::ffff:a0a:a04 (recursive via ::, tunvxlan2),
installed 00:03:09, last update 00:03:09 ago
B       2002::/48 [20/0] via ::ffff:a0a:a04 (recursive via ::, tunvxlan2),
installed 00:40:42, last update 00:40:42 ago
B       2002::2/128 [20/0] via ::ffff:a0a:a04 (recursive via ::, tunvxlan2),
installed 00:03:09, last update 00:03:09 ago
B       2003::/48 [20/0] via ::ffff:a0a:a04 (recursive via ::, tunvxlan2),
installed 00:40:42, last update 00:40:42 ago
B       2003::2/128 [20/0] via ::ffff:a0a:a04 (recursive via ::, tunvxlan2),
installed 00:03:09, last update 00:03:09 ago
```

```
C       fe80::/64 via ::, tunvxlanw2, installed 18:50:53, last update 18:50:53
ago
B       fe80::1/128 [20/0] via ::ffff:a0a:a04 (recursive via ::, tunvxlan2),
installed 00:03:09, last update 00:03:09 ago
BLEAF3#show bgp vrf all
BGP table version is 1, local router ID is 63.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-
Color

BGP Route Table for VRF l3_vrf_1
*>xw101.0.0.0        10.10.10.6           0        100       0   65006 65004 ?
*>xw101.1.1.2/32     10.10.10.5           0        100       0   65005 65004 i
*>xw101.1.1.240/32   10.10.10.6           0        100       0   65006 65004 ?
*>xw102.0.0.0        10.10.10.6           0        100       0   65006 65004 ?
*>xw102.1.1.2/32     10.10.10.5           0        100       0   65005 65004 i
*>xw103.0.0.0        10.10.10.6           0        100       0   65006 65004 ?
*>x 201.0.0.0/8      10.10.10.4           0        100       0   65002 65001 ?
*>x 201.1.1.2/32     10.10.10.4           0        100       0   65002 65001 i
*>x 201.1.1.240/32   10.10.10.4           0        100       0   65002 65001 ?
*>x 202.0.0.0/8      10.10.10.4           0        100       0   65002 65001 ?
*>x 202.1.1.2/32     10.10.10.4           0        100       0   65002 65001 i
*>x 203.0.0.0/8      10.10.10.4           0        100       0   65002 65001 ?
*>x 203.1.1.2/32     10.10.10.4           0        100       0   65002 65001 i

Total number of prefixes 13
*>  10.10.10.4/32    fe80::6eb9:c5ff:fe15:8765     0        100       0
65002 65001 i
*>  10.10.10.7/32    0.0.0.0              0        100       32768  i
*                    0.0.0.0              0        100       32768  ?
*>  10.10.10.16/32   fe80::6eb9:c5ff:fe15:8765     0        100       0
65002 i
*>  17.18.18.0/24    0.0.0.0              0        100       32768  ?
*>  17.19.19.0/24    0.0.0.0              0        100       32768  ?

Total number of prefixes 5
BLEAF3#show bgp ipv6 unicast vrf all
BGP table version is 1, local router ID is 63.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop                    Metric    LocPrf     Weight
Path

BGP Route Table for VRF l3_vrf_1
*>xw1001::/48        ::ffff:10.10.10.6            0        100       0
65006 65004 ?
*>xw1001::2/128      ::ffff:10.10.10.5            0        100       0
65005 65004 i
*>xw1002::/48        ::ffff:10.10.10.6            0        100       0
65006 65004 ?
```

```
*>xw1002::2/128      ::ffff:10.10.10.5            0        100       0
65005 65004 i
*>xw1003::/48        ::ffff:10.10.10.6            0        100       0
65006 65004 ?
*>x 2001::/48        ::ffff:10.10.10.4            0        100       0
65002 65001 ?
*>x 2001::2/128      ::ffff:10.10.10.4            0        100       0
65002 65001 i
*>x 2002::/48        ::ffff:10.10.10.4            0        100       0
65002 65001 ?
*>x 2002::2/128      ::ffff:10.10.10.4            0        100       0
65002 65001 i
*>x 2003::/48        ::ffff:10.10.10.4            0        100       0
65002 65001 ?
*>x 2003::2/128      ::ffff:10.10.10.4            0        100       0
65002 65001 i
*>x fe80::1/128      ::ffff:10.10.10.4            0        100       0
65002 65001 i

Total number of prefixes 12
BLEAF3#show bgp l2vpn evpn
BGP table version is 189, local router ID is 10.10.10.7
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight
Path  Peer          Encap

RD[7004:5001]
*>  [5]:[0]:[0]:[8]:[201.0.0.0]:[0.0.0.0]:[1000]
                        10.10.10.4           0        100       0   65002 65001
?  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [5]:[0]:[0]:[8]:[202.0.0.0]:[0.0.0.0]:[1000]
                        10.10.10.4           0        100       0   65002 65001
?  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [5]:[0]:[0]:[8]:[203.0.0.0]:[0.0.0.0]:[1000]
                        10.10.10.4           0        100       0   65002 65001
?  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.240]:[0.0.0.0]:[1000]
                        10.10.10.4           0        100       0   65002 65001
?  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [5]:[0]:[0]:[48]:[2001::]:[::]:[1000]
                        10.10.10.4           0        100       0   65002 65001
?  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [5]:[0]:[0]:[48]:[2002::]:[::]:[1000]
                        10.10.10.4           0        100       0   65002 65001
?  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [5]:[0]:[0]:[48]:[2003::]:[::]:[1000]
                        10.10.10.4           0        100       0   65002 65001
?  fe80::6eb9:c5ff:fe15:8765 VXLAN
```

```
RD[7005:7001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[2000]
                      10.10.10.5           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[2000]
                      10.10.10.5           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[2000]
                      10.10.10.5           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[32]:[101.1.1.2]:[0.0.0.0]:[2000]
                      10.10.10.5           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[2000]
                      10.10.10.5           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[32]:[102.1.1.2]:[0.0.0.0]:[2000]
                      10.10.10.5           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[2000]
                      10.10.10.5           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[2000]
                      10.10.10.5           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[2000]
                      10.10.10.5           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[128]:[1001::2]:[::]:[2000]
                      10.10.10.5           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[128]:[1002::2]:[::]:[2000]
                      10.10.10.5           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                     10.10.10.5           0        100        0   65006 65004
i  10.10.10.22     VXLAN

RD[7006:7001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[2000]
```

```
                        10.10.10.6           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.6           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[2000]
                        10.10.10.6           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.6           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[2000]
                        10.10.10.6           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.6           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[32]:[101.1.1.2]:[0.0.0.0]:[2000]
                        10.10.10.6           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                       10.10.10.6           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[2000]
                        10.10.10.6           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.6           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[32]:[102.1.1.2]:[0.0.0.0]:[2000]
                        10.10.10.6           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                       10.10.10.6           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[2000]
                        10.10.10.6           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.6           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[2000]
                        10.10.10.6           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.6           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[2000]
                        10.10.10.6           0        100        0   65006 65004
?  10.10.10.22     VXLAN
*                       10.10.10.6           0        100        0   65005 65004
?  10.10.10.21     VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[128]:[1001::2]:[::]:[2000]
                        10.10.10.6           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                       10.10.10.6           0        100        0   65006 65004
i  10.10.10.22     VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[128]:[1002::2]:[::]:[2000]
                        10.10.10.6           0        100        0   65005 65004
i  10.10.10.21     VXLAN
*                       10.10.10.6           0        100        0   65006 65004
i  10.10.10.22     VXLAN

RD[10.10.10.4:1]
*>  [2]:[0]:[1001]:[48,0000:0201:0001]:[0]:[1001]
                        10.10.10.4           0        100        0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1001]:[48,0000:0201:0001]:[32,201.1.1.2]:[1001]:[1000]
```

```
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1001]:[48,0000:0201:0001]:[128,2001::2][1001]:[1000]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1001]:[48,0000:0201:0001]:[128,fe80::1][1001]:[1000]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1001]:[48,e001:a6a9:4b01]:[32,201.1.1.1]:[1001]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1001]:[48,e001:a6a9:4b01]:[128,2001::1][1001]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1001]:[48,e001:a6a9:4b01]:[128,fe80::e201:a6ff:fea9:4b01][1001]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [3]:[1001]:[32,10.10.10.4]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN

RD[10.10.10.4:2]
*>  [2]:[0]:[1002]:[48,0000:0202:0001]:[0]:[1002]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1002]:[48,0000:0202:0001]:[32,202.1.1.2]:[1002]:[1000]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1002]:[48,0000:0202:0001]:[128,2002::2][1002]:[1000]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1002]:[48,e001:a6a9:4b01]:[32,202.1.1.1]:[1002]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1002]:[48,e001:a6a9:4b01]:[128,2002::1][1002]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1002]:[48,e001:a6a9:4b01]:[128,fe80::e201:a6ff:fea9:4b01][1002]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [3]:[1002]:[32,10.10.10.4]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN

RD[10.10.10.4:3]
*>  [2]:[0]:[1003]:[48,0000:0203:0001]:[0]:[1003]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1003]:[48,0000:0203:0001]:[32,203.1.1.2]:[1003]:[1000]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1003]:[48,0000:0203:0001]:[128,2003::2][1003]:[1000]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1003]:[48,e001:a6a9:4b01]:[32,203.1.1.1]:[1003]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [2]:[0]:[1003]:[48,e001:a6a9:4b01]:[128,2003::1][1003]
                         10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
```

```
*>  [2]:[0]:[1003]:[48,e001:a6a9:4b01]:[128,fe80::e201:a6ff:fea9:4b01][1003]
                       10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN
*>  [3]:[1003]:[32,10.10.10.4]
                       10.10.10.4           0        100       0   65002 65001
i  fe80::6eb9:c5ff:fe15:8765 VXLAN

Total number of prefixes 51
BLEAF3#show ip bgp 101.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 101.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65005 65004
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.5 (IGP metric 0) from 10.10.10.21 (Remote Id:10.10.10.21)
Peer nexthop: 10.10.10.21
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Stitched: (Recieved from the external fabric peer)
      Extended Community: SOO:1004:1004  RT:7001:7001  Encapsulation:VxLAN
EVPN_Router's_Mac:5c17.83ff.2255
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 07:56:30 2025,
00:11:47 ago

BLEAF3#show ip bgp 102.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 102.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65005 65004
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.5 (IGP metric 0) from 10.10.10.21 (Remote Id:10.10.10.21)
Peer nexthop: 10.10.10.21
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Stitched: (Recieved from the external fabric peer)
      Extended Community: SOO:1004:1004  RT:7001:7001  Encapsulation:VxLAN
EVPN_Router's_Mac:5c17.83ff.2255
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 07:56:30 2025,
00:11:47 ago

BLEAF3#show ip bgp 201.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 201.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65002 65001
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.4 (IGP metric 0) from fe80::6eb9:c5ff:fe15:8765 (Remote
Id:10.10.10.16) Peer nexthop: fe80::6eb9:c5ff:fe15:8765
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External,
imported)
      Extended Community: RT:201:201 5001:5001  Encapsulation:VxLAN
EVPN_Router's_Mac:e001.a6a9.4b01
      rx path_id: -1      tx path_id: -1
```

```
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 08:05:06 2025,
00:03:12 ago

BLEAF3#show ip bgp 202.1.1.2/32 vrf l3_vrf_1
BGP routing table entry for 202.1.1.2/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  AS path: 65002 65001
  Path Selection reason: Nothing left to compare
   Nexthop:10.10.10.4 (IGP metric 0) from fe80::6eb9:c5ff:fe15:8765 (Remote
Id:10.10.10.16) Peer nexthop: fe80::6eb9:c5ff:fe15:8765
      Origin IGP, metric 0, localpref 100      valid, external, best, source-
safi: 70
      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External,
imported)
      Extended Community: RT:202:202 5001:5001  Encapsulation:VxLAN
EVPN_Router's_Mac:e001.a6a9.4b01
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Origin-AS validity: disabled      Last update: Mon Dec 15 08:05:06 2025,
00:03:12 ago

BLEAF3#show evpn l3vni-map
 L3VNI        Trans-L3VNI  L2VNI        IRB-interface
==================================================
 1000         2000         ---          ---

BLEAF3#show nvo vxlan mac-table
================================================================================
================================================================================
=====
                                                          VXLAN MAC Entries
================================================================================
================================================================================
=====
VNID    Interface VlanId   In-VlanId Mac-Addr    VTEP-Ip/ESI
Type            Status     MAC move AccessPortDesc LeafFlag
_____________________________________________________________________________
_____________________________________________________________________________
_____


Total number of entries are : 0

BLEAF3#
BLEAF3#
```

LEAF1

```
LEAF1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged
```

```
VNID      VNI-Name     VNI-Type Type Interface ESI
VLAN       DF-Status Src-Addr          Dst-Addr         Router-Mac
________________________________________________________________________________
_________________________________________________________________
1000      ----         L3       NW   ----      ----                            --
--       ----      10.10.10.1        10.10.10.2       9819.2c9b.37d3
1000      ----         L3       NW   ----      ----                            --
--       ----      10.10.10.1        10.10.10.6       7c8d.9caa.5231
1000      ----         L3       NW   ----      ----                            --
--       ----      10.10.10.1        10.10.10.5       5c17.83ff.2255
1001      ----         L2       NW   ----      ----                            --
--       ----      10.10.10.1        10.10.10.2       --------------
1001      ----         --       AC   po1       00:00:00:aa:aa:bb:bb:00:00:00
101        NON-DF    ----              ----
1002      ----         L2       NW   ----      ----                            --
--       ----      10.10.10.1        10.10.10.2       --------------
1002      ----         --       AC   po1       00:00:00:aa:aa:bb:bb:00:00:00
102        DF        ----              ----
1003      ----         L2       NW   ----      ----                            --
--       ----      10.10.10.1        10.10.10.2       --------------
1003      ----         --       AC   po1       00:00:00:aa:aa:bb:bb:00:00:00
103        NON-DF    ----              ----

Total number of entries are 9
LEAF1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
==========================================================================
10.10.10.1       10.10.10.2       Installed     00:41:01      00:41:01
10.10.10.1       10.10.10.6       Installed     00:41:02      00:41:02
10.10.10.1       10.10.10.5       Installed     00:41:02      00:41:02

Total number of entries are 3
LEAF1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID                Ip-Addr          Mac-Addr       Type        Age-Out
Retries-Left
________________________________________________________________________________
_________
1001               101.1.1.1        0000.0000.abcd Static Local   ----
1001               101.1.1.2        0000.0101.0001 Dynamic Local  ----
1002               102.1.1.1        0000.0000.abcd Static Local   ----
1002               102.1.1.2        0000.0102.0001 Dynamic Remote ----
1003               103.1.1.1        0000.0000.abcd Static Local   ----
1003               103.1.1.2        0000.0103.0001 Dynamic Remote ----
Total number of entries are 6
LEAF1#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                 Mac-Addr       Type
Age-Out    Retries-Left
________________________________________________________________________________
______________________
1001     1001::1                               0000.0000.abcd Static Local
----
1001     1001::2                               0000.0101.0001 Dynamic Local
----
1001     fe80::1                               0000.0101.0001 Dynamic Remote
----
```

```
1001     fe80::200:ff:fe00:abcd                0000.0000.abcd Static Local
----
1002     1002::1                               0000.0000.abcd Static Local
----
1002     1002::2                               0000.0102.0001 Dynamic Remote
----
1002     fe80::200:ff:fe00:abcd                0000.0000.abcd Static Local
----
1003     1002::2                               0000.0103.0001 Dynamic Remote
----
1003     1003::1                               0000.0000.abcd Static Local
----
1003     1003::2                               0000.0103.0001 Dynamic Local
----
1003     fe80::200:ff:fe00:abcd                0000.0000.abcd Static Local
----
Total number of entries are 11
LEAF1#show ip bgp summary
BGP router identifier 10.10.10.1, local AS number 65001
BGP table version is 52
6 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe47(fe80::ba6a:97ff:fee2:4bce)         4 65002       927        754     51
0      0  00:41:05            11
xe8(fe80::1644:8fff:fe45:8dd5)          4 65002       911        682     52
0      0  00:41:05            11

Total number of neighbors 2

Total number of Established sessions 2
LEAF1#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.1, local AS number 65001
BGP table version is 116
6 BGP AS-PATH entries
0 BGP community entries

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
xe47(fe80::ba6a:97ff:fee2:4bce)         4 65002       927        754     114
0      0  00:41:05            54      4     17      3      1     29
xe8(fe80::1644:8fff:fe45:8dd5)          4 65002       911        682     116
0      0  00:41:05            54      4     17      3      1     29

Total number of neighbors 2

Total number of Established sessions 2
LEAF1#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
```

```
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.1/32 is directly connected, lo, installed 03:52:36, last
update 03:52:36 ago
B            10.10.10.2/32 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:02, last update 00:41:02 ago
                           [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
B            10.10.10.5/32 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:02, last update 00:41:02 ago
                           [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
B            10.10.10.6/32 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:02, last update 00:41:02 ago
                           [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
B            10.10.10.11/32 [20/0] via fe80::1644:8fff:fe45:8dd5, xe8,
installed 00:41:02, last update 00:41:02 ago
B            10.10.10.12/32 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:02, last update 00:41:02 ago
B            10.10.10.14/32 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:03, last update 00:41:03 ago
                            [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
B            10.10.10.15/32 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:03, last update 00:41:03 ago
                            [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
B            17.20.20.0/24 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:02, last update 00:41:02 ago
                           [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
B            17.21.21.0/24 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:02, last update 00:41:02 ago
                           [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
B            17.22.22.0/24 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:02, last update 00:41:02 ago
                           [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
B            17.23.23.0/24 [20/0] via fe80::ba6a:97ff:fee2:4bce, xe47,
installed 00:41:02, last update 00:41:02 ago
                           [20/0] via fe80::1644:8fff:fe45:8dd5, xe8
C            127.0.0.0/8 is directly connected, lo, installed 03:52:52, last
update 03:52:52 ago
IP Route Table for VRF "management"
C            10.16.111.0/24 is directly connected, eth0, installed 03:52:34,
last update 03:52:34 ago
C            127.0.0.0/8 is directly connected, lo.management, installed
03:52:52, last update 03:52:52 ago
IP Route Table for VRF "l3_vrf_1"
B            10.10.10.2/32 [0/0] is directly connected, tunvxlan2, installed
00:41:02, last update 00:41:02 ago
B            10.10.10.5/32 [0/0] is directly connected, tunvxlan2, installed
00:41:02, last update 00:41:02 ago
B            10.10.10.6/32 [0/0] is directly connected, tunvxlan2, installed
00:41:02, last update 00:41:02 ago
C            101.0.0.0/8 is directly connected, irb1001, installed 03:52:36,
last update 03:52:36 ago
S            101.1.1.240/32 [1/0] is directly connected, irb1001, installed
03:52:36, last update 03:52:36 ago
C            102.0.0.0/8 is directly connected, irb1002, installed 03:52:36,
last update 03:52:36 ago
C            103.0.0.0/8 is directly connected, irb1003, installed 03:52:36,
last update 03:52:36 ago
C            127.0.0.0/8 is directly connected, lo.l3_vrf_1, installed
03:52:40, last update 03:52:40 ago
```

```
B            201.0.0.0/8 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:41:03, last update 00:41:03 ago
                         [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
B            201.1.1.2/32 [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2), installed 00:03:30, last update 00:03:30 ago
                          [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2)
B            201.1.1.240/32 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:41:03, last update 00:41:03 ago
                            [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
B            202.0.0.0/8 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:41:03, last update 00:41:03 ago
                         [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
B            202.1.1.2/32 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:03:30, last update 00:03:30 ago
                          [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
B            203.0.0.0/8 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:41:03, last update 00:41:03 ago
                         [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
IP Route Table for VRF "l2_vrf_1"
IP Route Table for VRF "l2_vrf_2"
IP Route Table for VRF "l2_vrf_3"
IP Route Table for VRF "evpn-gvrf-1"

Gateway of last resort is not set
LEAF1#show ipv6 route vrf all
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 03:52:53, last update 03:52:53 ago
C       fe80::/64 via ::, xe10, installed 03:52:08, last update 03:52:08 ago
IP Route Table for VRF "management"
C       ::1/128 via ::, lo.management, installed 03:52:53, last update
03:52:53 ago
C       fe80::/64 via ::, eth0, installed 03:52:53, last update 03:52:53 ago
IP Route Table for VRF "l3_vrf_1"
C       ::1/128 via ::, lo.l3_vrf_1, installed 03:52:41, last update 03:52:41
ago
B       ::ffff:a0a:a02/128 [0/0] via ::, tunvxlan2, installed 00:41:03, last
update 00:41:03 ago
B       ::ffff:a0a:a05/128 [0/0] via ::, tunvxlan2, installed 00:41:03, last
update 00:41:03 ago
B       ::ffff:a0a:a06/128 [0/0] via ::, tunvxlan2, installed 00:41:03, last
update 00:41:03 ago
C       1001::/48 via ::, irb1001, installed 03:52:37, last update 03:52:37
ago
C       1002::/48 via ::, irb1002, installed 03:52:37, last update 03:52:37
ago
```

```
C       1003::/48 via ::, irb1003, installed 03:52:37, last update 03:52:37
ago
B       2001::/48 [20/0] via ::ffff:a0a:a06 (recursive via ::, tunvxlan2),
installed 00:41:04, last update 00:41:04 ago
                  [20/0] via ::ffff:a0a:a05 (recursive via ::, tunvxlan2)
B       2001::2/128 [20/0] via ::ffff:a0a:a06 (recursive via ::, tunvxlan2),
installed 00:03:31, last update 00:03:31 ago
                    [20/0] via ::ffff:a0a:a05 (recursive via ::, tunvxlan2)
B       2002::/48 [20/0] via ::ffff:a0a:a06 (recursive via ::, tunvxlan2),
installed 00:41:04, last update 00:41:04 ago
                  [20/0] via ::ffff:a0a:a05 (recursive via ::, tunvxlan2)
B       2002::2/128 [20/0] via ::ffff:a0a:a06 (recursive via ::, tunvxlan2),
installed 00:03:31, last update 00:03:31 ago
                    [20/0] via ::ffff:a0a:a05 (recursive via ::, tunvxlan2)
B       2003::/48 [20/0] via ::ffff:a0a:a06 (recursive via ::, tunvxlan2),
installed 00:41:04, last update 00:41:04 ago
                  [20/0] via ::ffff:a0a:a05 (recursive via ::, tunvxlan2)
C       fe80::/64 via ::, irb1003, installed 03:52:37, last update 03:52:37
ago
IP Route Table for VRF "l2_vrf_1"
IP Route Table for VRF "l2_vrf_2"
IP Route Table for VRF "l2_vrf_3"
IP Route Table for VRF "evpn-gvrf-1"
LEAF1#show bgp vrf all
BGP table version is 1, local router ID is 103.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-
Color

BGP Route Table for VRF l3_vrf_1
*>  101.0.0.0        0.0.0.0              0        100       32768  ?
* x                  10.10.10.2           0        100       0   65002 65001 ?
*>  101.1.1.240/32   0.0.0.0              0        100       32768  ?
* x                  10.10.10.2           0        100       0   65002 65001 ?
*>  102.0.0.0        0.0.0.0              0        100       32768  ?
* x                  10.10.10.2           0        100       0   65002 65001 ?
* x 102.1.1.2/32     10.10.10.2           0        100       0   65002 65001 i
*>  103.0.0.0        0.0.0.0              0        100       32768  ?
* x                  10.10.10.2           0        100       0   65002 65001 ?
* x 103.1.1.2/32     10.10.10.2           0        100       0   65002 65001 i
*>x 201.0.0.0/8      10.10.10.6           0        100       0   65002 65003
65004 ?
* x                  10.10.10.5           0        100       0   65002 65003
65004 ?
*>x 201.1.1.2/32     10.10.10.5           0        100       0   65002 65003
65004 i
* x                  10.10.10.6           0        100       0   65002 65003
65004 i
*>x 201.1.1.240/32   10.10.10.6           0        100       0   65002 65003
65004 ?
* x                  10.10.10.5           0        100       0   65002 65003
65004 ?
*>x 202.0.0.0/8      10.10.10.6           0        100       0   65002 65003
65004 ?
```

```
* x                  10.10.10.5           0        100        0   65002 65003
65004 ?
*>x 202.1.1.2/32     10.10.10.6           0        100        0   65002 65003
65004 i
* x                  10.10.10.5           0        100        0   65002 65003
65004 i
*>x 203.0.0.0/8      10.10.10.6           0        100        0   65002 65003
65004 ?
* x                  10.10.10.5           0        100        0   65002 65003
65004 ?

Total number of prefixes 12
*>  10.10.10.1/32    0.0.0.0              0        100       32768  i
*                   0.0.0.0              0        100       32768  ?
*>  10.10.10.2/32    fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65001 i
*                   fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65001 i
*>  10.10.10.5/32    fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 65004 i
*                   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 65004 i
*>  10.10.10.6/32    fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 65004 i
*                   fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 65004 i
*>  10.10.10.11/32   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 i
*                   fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 65002 i
*>  10.10.10.12/32   fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 i
*                   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 65002 i
*>  10.10.10.14/32   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 i
*                   fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 i
*>  10.10.10.15/32   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 i
*                   fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 i
*>  17.20.20.0/24    fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 65004 ?
*                   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 65004 ?
*>  17.21.21.0/24    fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 65004 ?
*                   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 65004 ?
*>  17.22.22.0/24    fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 65004 ?
*                   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 65004 ?
*>  17.23.23.0/24    fe80::ba6a:97ff:fee2:4bce     0        100        0
65002 65003 65004 ?
*                   fe80::1644:8fff:fe45:8dd5     0        100        0
65002 65003 65004 ?

Total number of prefixes 12
LEAF1#show bgp l2vpn evpn
```

```
BGP table version is 116, local router ID is 10.10.10.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight
Path  Peer          Encap

RD[7002:1001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.2           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.2           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.2           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[1000]
                       10.10.10.2           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[1000]
                       10.10.10.2           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[1000]
                       10.10.10.2           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[1000]
                       10.10.10.2           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4bce VXLAN

RD[7005:1001]
*>  [5]:[0]:[0]:[8]:[201.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
```

```
*>  [5]:[0]:[0]:[8]:[202.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[8]:[203.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.240]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[32]:[202.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4bce VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>  [5]:[0]:[0]:[48]:[2001::]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[48]:[2002::]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[48]:[2003::]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[128]:[2001::2]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[128]:[2002::2]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4bce VXLAN
*                      10.10.10.5           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dd5 VXLAN

RD[7006:1001]
*>  [5]:[0]:[0]:[8]:[201.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[8]:[202.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
```

```
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[8]:[203.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.240]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[32]:[202.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[48]:[2001::]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[48]:[2002::]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[48]:[2003::]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[128]:[2001::2]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [5]:[0]:[0]:[128]:[2002::2]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4bce VXLAN

RD[10.10.10.1:1] VRF[l2_vrf_1]
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[1001]
                       10.10.10.2           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100        32768  i  ------
----      VXLAN
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                       10.10.10.2           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*   [2]:[0]:[1001]:[48,0000:0000:abcd]:[32,101.1.1.1]:[1001]
                       10.10.10.2           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
```

```
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,1001::1][1001]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1001]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[0]:[1001]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[32,101.1.1.2]:
[1001]:[1000]
                      10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,1001::2][1
001]:[1000]
                      10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,fe80::1][1
001]:[1000]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [3]:[1001]:[32,10.10.10.1]
                      10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*   [3]:[1001]:[32,10.10.10.2]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN

RD[10.10.10.1:2] VRF[l2_vrf_2]
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*   [2]:[0]:[1002]:[48,0000:0000:abcd]:[32,102.1.1.1]:[1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,1002::1][1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
```

```
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[0]:[1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[32,102.1.1.2]:
[1002]:[1000]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[128,1002::2][1
002]:[1000]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>  [3]:[1002]:[32,10.10.10.1]
                      10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*   [3]:[1002]:[32,10.10.10.2]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN

RD[10.10.10.1:3] VRF[l2_vrf_3]
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[1003]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*   [2]:[0]:[1003]:[48,0000:0000:abcd]:[32,103.1.1.1]:[1003]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,1003::1][1003]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1003]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>                    10.10.10.1           0        100       32768  i  ------
----      VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[0]:[1003]
                      10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[32,103.1.1.2]:
[1003]:[1000]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1002::2][1
003]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
```

```
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1003::2][1
003]:[1000]
                     10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*>  [3]:[1003]:[32,10.10.10.1]
                     10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*   [3]:[1003]:[32,10.10.10.2]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN

RD[10.10.10.1:64512] VRF[evpn-gvrf-1]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                     10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*>  [4]:[00:00:00:aa:aa:bb:bb:00:00:00]:[32,10.10.10.1]
                     10.10.10.1           0        100       32768  i  -----
-----      VXLAN
*   [4]:[00:00:00:aa:aa:bb:bb:00:00:00]:[32,10.10.10.2]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN

RD[10.10.10.2:1]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[1001]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                    10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[32,101.1.1.1]:[1001]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                    10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,1001::1][1001]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                    10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1001]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                    10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[0]:[1001]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*                    10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,fe80::1][1
001]:[1000]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*                    10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*>  [3]:[1001]:[32,10.10.10.2]
                     10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                    10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
```

```
RD[10.10.10.2:2]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[32,102.1.1.1]:[1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,1002::1][1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[0]:[1002]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[32,102.1.1.2]:
[1002]:[1000]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[128,1002::2][1
002]:[1000]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [3]:[1002]:[32,10.10.10.2]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN

RD[10.10.10.2:3]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[1003]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[32,103.1.1.1]:[1003]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                     10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,1003::1][1003]
                      10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
```

```
*                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1003]
                       10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[0]:[1003]
                       10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[32,103.1.1.2]:
[1003]:[1000]
                       10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1002::2][1
003]
                       10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [3]:[1003]:[32,10.10.10.2]
                       10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN

RD[10.10.10.2:64512]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                       10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN
*>  [4]:[00:00:00:aa:aa:bb:bb:00:00:00]:[32,10.10.10.2]
                       10.10.10.2           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dd5 VXLAN
*                      10.10.10.2           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4bce VXLAN

Total number of prefixes 89
LEAF1#show nvo vxlan mac-table
================================================================================
================================================================================
=====
                                                        VXLAN MAC Entries
================================================================================
================================================================================
=====
VNID    Interface VlanId   In-VlanId Mac-Addr    VTEP-Ip/ESI
Type            Status     MAC move AccessPortDesc LeafFlag
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
-----

1001       po1       101       ----      0000.0101.0001
00:00:00:aa:aa:bb:bb:00:00:00               Dynamic Local   -------    0
-------          ----
```

```
1001    irb1001  ----    ----    0000.0000.abcd 10.10.10.1
Static Local    -------    0        -------        ----
1002       ----      ----      ----      0000.0102.0001
00:00:00:aa:aa:bb:bb:00:00:00                Dynamic Remote  -------    0
-------        ----
1002    irb1002  ----    ----    0000.0000.abcd 10.10.10.1
Static Local    -------    0        -------        ----
1003       po1       103       ----      0000.0103.0001
00:00:00:aa:aa:bb:bb:00:00:00                Dynamic Local   -------    0
-------        ----
1003    irb1003  ----    ----    0000.0000.abcd 10.10.10.1
Static Local    -------    0        -------        ----

Total number of entries are : 6

LEAF1#
LEAF1#
LEAF1#
LEAF1#
LEAF1#
LEAF1#
```

LEAF2

```
LEAF2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI
VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
________________________________________________________________________________
__________________________________________________________________
1000     ----         L3      NW   ----      ----                           --
--      ----      10.10.10.2      10.10.10.6       7c8d.9caa.5231
1000     ----         L3      NW   ----      ----                           --
--      ----      10.10.10.2      10.10.10.5       5c17.83ff.2255
1000     ----         L3      NW   ----      ----                           --
--      ----      10.10.10.2      10.10.10.1       1444.8fc8.0fc9
1001     ----         L2      NW   ----      ----                           --
--      ----      10.10.10.2      10.10.10.1       --------------
1001     ----         --      AC   po1       00:00:00:aa:aa:bb:bb:00:00:00
101       DF        ----           ----
1002     ----         L2      NW   ----      ----                           --
--      ----      10.10.10.2      10.10.10.1       --------------
1002     ----         --      AC   po1       00:00:00:aa:aa:bb:bb:00:00:00
102       NON-DF    ----           ----
1003     ----         L2      NW   ----      ----                           --
--      ----      10.10.10.2      10.10.10.1       --------------
1003     ----         --      AC   po1       00:00:00:aa:aa:bb:bb:00:00:00
103       DF        ----           ----

Total number of entries are 9
LEAF2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
```

```
================================================================================
10.10.10.2       10.10.10.6       Installed     00:41:12      00:41:12
10.10.10.2       10.10.10.5       Installed     00:41:11      00:41:11
10.10.10.2       10.10.10.1       Installed     00:41:11      00:41:11

Total number of entries are 3
LEAF2#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID                Ip-Addr          Mac-Addr       Type        Age-Out
Retries-Left
_____________________________________________________________________________
_________
1001              101.1.1.1        0000.0000.abcd Static Local   ----
1001              101.1.1.2        0000.0101.0001 Dynamic Remote ----
1002              102.1.1.1        0000.0000.abcd Static Local   ----
1002              102.1.1.2        0000.0102.0001 Dynamic Local  ----
1003              103.1.1.1        0000.0000.abcd Static Local   ----
1003              103.1.1.2        0000.0103.0001 Dynamic Local  ----
Total number of entries are 6
LEAF2#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr                                  Mac-Addr       Type
Age-Out    Retries-Left
_____________________________________________________________________________
_____________________
1001     1001::1                                 0000.0000.abcd Static Local
----
1001     1001::2                                 0000.0101.0001 Dynamic Remote
----
1001     fe80::1                                 0000.0101.0001 Dynamic Local
----
1001     fe80::200:ff:fe00:abcd                  0000.0000.abcd Static Local
----
1002     1002::1                                 0000.0000.abcd Static Local
----
1002     1002::2                                 0000.0102.0001 Dynamic Local
----
1002     fe80::200:ff:fe00:abcd                  0000.0000.abcd Static Local
----
1003     1002::2                                 0000.0103.0001 Dynamic Local
----
1003     1003::1                                 0000.0000.abcd Static Local
----
1003     1003::2                                 0000.0103.0001 Dynamic Remote
----
1003     fe80::200:ff:fe00:abcd                  0000.0000.abcd Static Local
----
Total number of entries are 11
LEAF2#show ip bgp summary
BGP router identifier 10.10.10.2, local AS number 65001
BGP table version is 32
6 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
```

```
xe1(fe80::ba6a:97ff:fee2:4ba0)          4 65002      1023       681     32
0       0  00:41:13            11
xe3(fe80::1644:8fff:fe45:8dcd)          4 65002       961       782     32
0       0  00:41:13            11

Total number of neighbors 2

Total number of Established sessions 2
LEAF2#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.2, local AS number 65001
BGP table version is 117
6 BGP AS-PATH entries
0 BGP community entries

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
xe1(fe80::ba6a:97ff:fee2:4ba0)          4 65002      1023       681    117
0       0  00:41:13            51      4     14      3       1     29
xe3(fe80::1644:8fff:fe45:8dcd)          4 65002       961       782    117
0       0  00:41:13            51      4     14      3       1     29

Total number of neighbors 2

Total number of Established sessions 2
LEAF2#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:12, last update 00:41:12 ago
                           [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
C            10.10.10.2/32 is directly connected, lo, installed 03:56:46, last
update 03:56:46 ago
B            10.10.10.5/32 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:12, last update 00:41:12 ago
                           [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
B            10.10.10.6/32 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:12, last update 00:41:12 ago
                           [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
B            10.10.10.11/32 [20/0] via fe80::1644:8fff:fe45:8dcd, xe3,
installed 00:41:12, last update 00:41:12 ago
B            10.10.10.12/32 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:12, last update 00:41:12 ago
B            10.10.10.14/32 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:13, last update 00:41:13 ago
                            [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
B            10.10.10.15/32 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:13, last update 00:41:13 ago
                            [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
B            17.20.20.0/24 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:12, last update 00:41:12 ago
```

```
                              [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
B            17.21.21.0/24 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:12, last update 00:41:12 ago
                              [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
B            17.22.22.0/24 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:12, last update 00:41:12 ago
                              [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
B            17.23.23.0/24 [20/0] via fe80::ba6a:97ff:fee2:4ba0, xe1,
installed 00:41:12, last update 00:41:12 ago
                              [20/0] via fe80::1644:8fff:fe45:8dcd, xe3
C            127.0.0.0/8 is directly connected, lo, installed 03:57:00, last
update 03:57:00 ago
IP Route Table for VRF "management"
C            10.16.111.0/24 is directly connected, eth0, installed 03:56:43,
last update 03:56:43 ago
C            127.0.0.0/8 is directly connected, lo.management, installed
03:57:00, last update 03:57:00 ago
IP Route Table for VRF "l3_vrf_1"
B            10.10.10.1/32 [0/0] is directly connected, tunvxlan2, installed
00:41:12, last update 00:41:12 ago
B            10.10.10.5/32 [0/0] is directly connected, tunvxlan2, installed
00:41:12, last update 00:41:12 ago
B            10.10.10.6/32 [0/0] is directly connected, tunvxlan2, installed
00:41:12, last update 00:41:12 ago
C            101.0.0.0/8 is directly connected, irb1001, installed 03:56:46,
last update 03:56:46 ago
S            101.1.1.240/32 [1/0] is directly connected, irb1001, installed
03:56:45, last update 03:56:45 ago
C            102.0.0.0/8 is directly connected, irb1002, installed 03:56:46,
last update 03:56:46 ago
C            103.0.0.0/8 is directly connected, irb1003, installed 03:56:46,
last update 03:56:46 ago
C            127.0.0.0/8 is directly connected, lo.l3_vrf_1, installed
03:56:47, last update 03:56:47 ago
B            201.0.0.0/8 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:41:13, last update 00:41:13 ago
                           [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
B            201.1.1.2/32 [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2), installed 00:03:40, last update 00:03:40 ago
                            [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2)
B            201.1.1.240/32 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:41:13, last update 00:41:13 ago
                              [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
B            202.0.0.0/8 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:41:13, last update 00:41:13 ago
                           [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
B            202.1.1.2/32 [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2), installed 00:03:40, last update 00:03:40 ago
                            [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2)
B            203.0.0.0/8 [20/0] via 10.10.10.5 (recursive is directly
connected, tunvxlan2), installed 00:41:13, last update 00:41:13 ago
                           [20/0] via 10.10.10.6 (recursive is directly
connected, tunvxlan2)
IP Route Table for VRF "l2_vrf_1"
IP Route Table for VRF "l2_vrf_2"
IP Route Table for VRF "l2_vrf_3"
```

```
IP Route Table for VRF "evpn-gvrf-1"

Gateway of last resort is not set
LEAF2#
LEAF2#show bgp l2vpn evpn
BGP table version is 117, local router ID is 10.10.10.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight
Path  Peer          Encap

RD[7001:1001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[1000]
                      10.10.10.1           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.1           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[1000]
                      10.10.10.1           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.1           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[1000]
                      10.10.10.1           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.1           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[1000]
                      10.10.10.1           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.1           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[1000]
                      10.10.10.1           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.1           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[1000]
                      10.10.10.1           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.1           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[1000]
                      10.10.10.1           0        100        0   65002 65001
?  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.1           0        100        0   65002 65001
?  fe80::ba6a:97ff:fee2:4ba0 VXLAN

RD[7005:1001]
```

```
*>  [5]:[0]:[0]:[8]:[201.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[8]:[202.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[8]:[203.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.240]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[32]:[202.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[48]:[2001::]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[48]:[2002::]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[48]:[2003::]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[128]:[2001::2]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[128]:[2002::2]:[::]:[1000]
                       10.10.10.5           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.5           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4ba0 VXLAN

RD[7006:1001]
*>  [5]:[0]:[0]:[8]:[201.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
```

```
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[8]:[202.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[8]:[203.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[32]:[201.1.1.240]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[32]:[202.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[48]:[2001::]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[48]:[2002::]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[48]:[2003::]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 ?  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [5]:[0]:[0]:[128]:[2001::2]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [5]:[0]:[0]:[128]:[2002::2]:[::]:[1000]
                       10.10.10.6           0        100        0   65002 65003
65004 i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.6           0        100        0   65002 65003
65004 i  fe80::ba6a:97ff:fee2:4ba0 VXLAN

RD[10.10.10.1:1]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[1001]
                       10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                     10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[32,101.1.1.1]:[1001]
```

```
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,1001::1][1001]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1001]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[0]:[1001]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[32,101.1.1.2]:
[1001]:[1000]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,1001::2][1
001]:[1000]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [3]:[1001]:[32,10.10.10.1]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN

RD[10.10.10.1:2]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[1002]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[32,102.1.1.1]:[1002]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,1002::1][1002]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1002]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                      10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [3]:[1002]:[32,10.10.10.1]
```

```
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN

RD[10.10.10.1:3]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[1003]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[32,103.1.1.1]:[1003]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,1003::1][1003]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1003]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[0]:[1003]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1003::2][1
003]:[1000]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [3]:[1003]:[32,10.10.10.1]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN

RD[10.10.10.1:64512]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>  [4]:[00:00:00:aa:aa:bb:bb:00:00:00]:[32,10.10.10.1]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*                       10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN

RD[10.10.10.2:1] VRF[l2_vrf_1]
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[1001]
                        10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
```

```
*>                    10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*   [2]:[0]:[1001]:[48,0000:0000:abcd]:[32,101.1.1.1]:[1001]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                    10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,1001::1][1001]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                    10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1001]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1001]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                    10.10.10.2           0        100       32768  i  ------
----      VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[0]:[1001]
                       10.10.10.2           0        100       32768  i  -----
-----      VXLAN
*                     10.10.10.1           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[32,101.1.1.2]:
[1001]:[1000]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,1001::2][1
001]:[1000]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1001]:[48,0000:0101:0001]:[128,fe80::1][1
001]:[1000]
                       10.10.10.2           0        100       32768  i  -----
-----      VXLAN
*   [3]:[1001]:[32,10.10.10.1]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [3]:[1001]:[32,10.10.10.2]
                       10.10.10.2           0        100       32768  i  -----
-----      VXLAN

RD[10.10.10.2:2] VRF[l2_vrf_2]
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[1002]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                    10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*   [2]:[0]:[1002]:[48,0000:0000:abcd]:[32,102.1.1.1]:[1002]
                       10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                    10.10.10.2           0        100       32768  i  ------
----      VXLAN
```

```
*   [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,1002::1][1002]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                     10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1002]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1002]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                     10.10.10.2           0        100       32768  i  ------
----      VXLAN
*>  [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[0]:[1002]
                         10.10.10.2           0        100       32768  i  -----
-----      VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[32,102.1.1.2]:
[1002]:[1000]
                         10.10.10.2           0        100       32768  i  -----
-----      VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1002]:[48,0000:0102:0001]:[128,1002::2][1
002]:[1000]
                         10.10.10.2           0        100       32768  i  -----
-----      VXLAN
*   [3]:[1002]:[32,10.10.10.1]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [3]:[1002]:[32,10.10.10.2]
                         10.10.10.2           0        100       32768  i  -----
-----      VXLAN

RD[10.10.10.2:3] VRF[l2_vrf_3]
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[1003]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                     10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*   [2]:[0]:[1003]:[48,0000:0000:abcd]:[32,103.1.1.1]:[1003]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                     10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,1003::1][1003]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                     10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [2]:[0]:[1003]:[48,0000:0000:abcd]:[128,fe80::200:ff:fe00:abcd][1003]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>                     10.10.10.2           0        100       32768  i  ------
----      VXLAN
*   [2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[0]:[1003]
                         10.10.10.1           0        100        0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*>                     10.10.10.2           0        100       32768  i  ------
----      VXLAN
```

```
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[32,103.1.1.2]:
[1003]:[1000]
                      10.10.10.2           0        100       32768  i  -----
-----      VXLAN
*>
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1002::2][1
003]
                      10.10.10.2           0        100       32768  i  -----
-----      VXLAN
*
[2]:[00:00:00:aa:aa:bb:bb:00:00:00]:[1003]:[48,0000:0103:0001]:[128,1003::2][1
003]:[1000]
                      10.10.10.1           0        100       0   65002 65001
i  fe80::ba6a:97ff:fee2:4ba0 VXLAN
*   [3]:[1003]:[32,10.10.10.1]
                      10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [3]:[1003]:[32,10.10.10.2]
                      10.10.10.2           0        100       32768  i  -----
-----      VXLAN

RD[10.10.10.2:64512] VRF[evpn-gvrf-1]
*>  [1]:[00:00:00:aa:aa:bb:bb:00:00:00]:[4294965295]:[0]
                      10.10.10.2           0        100       32768  i  -----
-----      VXLAN
*   [4]:[00:00:00:aa:aa:bb:bb:00:00:00]:[32,10.10.10.1]
                      10.10.10.1           0        100       0   65002 65001
i  fe80::1644:8fff:fe45:8dcd VXLAN
*>  [4]:[00:00:00:aa:aa:bb:bb:00:00:00]:[32,10.10.10.2]
                      10.10.10.2           0        100       32768  i  -----
-----      VXLAN

Total number of prefixes 86
LEAF2#terminal length 0
LEAF2#
LEAF2#
LEAF2#
LEAF2#show nvo vxlan mac-table
================================================================================
================================================================================
=====
                                                          VXLAN MAC Entries
================================================================================
================================================================================
=====
VNID    Interface VlanId   In-VlanId Mac-Addr    VTEP-Ip/ESI
Type            Status     MAC move AccessPortDesc LeafFlag

--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
------

1001       po1       101       ----      0000.0101.0001
00:00:00:aa:aa:bb:bb:00:00:00               Dynamic Local   -------    0
-------           ----
1001    irb1001  ----    ----    0000.0000.abcd 10.10.10.2
Static Local    -------    0        -------          ----
1002       po1       102       ----      0000.0102.0001
00:00:00:aa:aa:bb:bb:00:00:00               Dynamic Local   -------    0
-------           ----
```

```
1002    irb1002  ----    ----    0000.0000.abcd 10.10.10.2
Static Local    -------    0        -------          ----
1003       po1       103       ----      0000.0103.0001
00:00:00:aa:aa:bb:bb:00:00:00              Dynamic Local   -------    0
-------        ----
1003    irb1003  ----    ----    0000.0000.abcd 10.10.10.2
Static Local    -------    0        -------          ----

Total number of entries are : 6

LEAF2#
```

LEAF3

```
LEAF3#terminal length 0
LEAF3#
LEAF3#
LEAF3#
LEAF3#
LEAF3#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI
VLAN      DF-Status Src-Addr         Dst-Addr        Router-Mac
_______________________________________________________________________________
_________________________________________________________________
1000     ----         L3      NW   ----       ----                          --
--      ----      10.10.10.4       10.10.10.7       1444.8fc9.07c9
1001     ----         --      AC   xe34       --- Single Homed Port ---
2001      ----     ----              ----
1002     ----         --      AC   xe34       --- Single Homed Port ---
2002      ----     ----              ----
1003     ----         --      AC   xe34       --- Single Homed Port ---
2003      ----     ----              ----

Total number of entries are 4
LEAF3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
========================================================================
10.10.10.4       10.10.10.7       Installed     00:41:34      00:41:34

Total number of entries are 1
LEAF3#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID                Ip-Addr          Mac-Addr       Type        Age-Out
Retries-Left
_______________________________________________________________________________
_________
1001               201.1.1.1       e001.a6a9.4b01 Static Local   ----
1001               201.1.1.2       0000.0201.0001 Dynamic Local  ----
1002               202.1.1.1       e001.a6a9.4b01 Static Local   ----
```

```
1002                202.1.1.2         0000.0202.0001 Dynamic Local  ----
1003                203.1.1.1         e001.a6a9.4b01 Static Local   ----
1003                203.1.1.2         0000.0203.0001 Dynamic Local  ----
Total number of entries are 6
LEAF3#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
==========================
VNID     Ip-Addr                                  Mac-Addr       Type
Age-Out    Retries-Left
_____________________________________________________________________________
_____________________
1001     2001::1                                  e001.a6a9.4b01 Static Local
----
1001     2001::2                                  0000.0201.0001 Dynamic Local
----
1001     fe80::1                                  0000.0201.0001 Dynamic Local
----
1001     fe80::e201:a6ff:fea9:4b01                e001.a6a9.4b01 Static Local
----
1002     2002::1                                  e001.a6a9.4b01 Static Local
----
1002     2002::2                                  0000.0202.0001 Dynamic Local
----
1002     fe80::e201:a6ff:fea9:4b01                e001.a6a9.4b01 Static Local
----
1003     2003::1                                  e001.a6a9.4b01 Static Local
----
1003     2003::2                                  0000.0203.0001 Dynamic Local
----
1003     fe80::e201:a6ff:fea9:4b01                e001.a6a9.4b01 Static Local
----
Total number of entries are 10
LEAF3#show ip bgp summary
BGP router identifier 10.10.10.4, local AS number 65001
BGP table version is 4
3 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
ce4(fe80::6eb9:c5ff:fe15:878f)          4 65002      3095      3142      4
0      0  00:41:35            4

Total number of neighbors 1

Total number of Established sessions 1
LEAF3#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.4, local AS number 65001
BGP table version is 34
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
ce4(fe80::6eb9:c5ff:fe15:878f)          4 65002      3095      3142     34
0      0  00:41:35           11      0      0      0      0     11

Total number of neighbors 1
```

```
Total number of Established sessions 1
LEAF3#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.4/32 is directly connected, lo, installed 21:28:10, last
update 21:28:10 ago
B            10.10.10.7/32 [20/0] via fe80::6eb9:c5ff:fe15:878f, ce4,
installed 00:41:34, last update 00:41:34 ago
B            10.10.10.16/32 [20/0] via fe80::6eb9:c5ff:fe15:878f, ce4,
installed 00:41:34, last update 00:41:34 ago
B            17.18.18.0/24 [20/0] via fe80::6eb9:c5ff:fe15:878f, ce4,
installed 00:41:34, last update 00:41:34 ago
B            17.19.19.0/24 [20/0] via fe80::6eb9:c5ff:fe15:878f, ce4,
installed 00:41:34, last update 00:41:34 ago
C            127.0.0.0/8 is directly connected, lo, installed 21:28:46, last
update 21:28:46 ago
IP Route Table for VRF "management"
C            10.16.111.0/24 is directly connected, eth0, installed 21:28:07,
last update 21:28:07 ago
C            127.0.0.0/8 is directly connected, lo.management, installed
21:28:46, last update 21:28:46 ago
IP Route Table for VRF "l3_vrf_1"
B            10.10.10.7/32 [0/0] is directly connected, tunvxlan2, installed
00:41:34, last update 00:41:34 ago
B            101.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlan2), installed 00:41:35, last update 00:41:35 ago
B            101.1.1.2/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlan2), installed 00:22:43, last update 00:22:43 ago
B            101.1.1.240/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlan2), installed 00:41:35, last update 00:41:35 ago
B            102.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlan2), installed 00:41:35, last update 00:41:35 ago
B            102.1.1.2/32 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlan2), installed 00:22:43, last update 00:22:43 ago
B            103.0.0.0/8 [20/0] via 10.10.10.7 (recursive is directly
connected, tunvxlan2), installed 00:41:35, last update 00:41:35 ago
C            127.0.0.0/8 is directly connected, lo.l3_vrf_1, installed
21:28:19, last update 21:28:19 ago
C            201.0.0.0/8 is directly connected, irb1001, installed 01:06:59,
last update 01:06:59 ago
S            201.1.1.240/32 [1/0] is directly connected, irb1001, installed
01:04:14, last update 01:04:14 ago
C            202.0.0.0/8 is directly connected, irb1002, installed 01:06:59,
last update 01:06:59 ago
C            203.0.0.0/8 is directly connected, irb1003, installed 01:06:59,
last update 01:06:59 ago
IP Route Table for VRF "l2_vrf_1"
IP Route Table for VRF "l2_vrf_2"
IP Route Table for VRF "l2_vrf_3"

Gateway of last resort is not set
```

```
LEAF3#show ipv6 route vrf all
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 21:28:46, last update 21:28:46 ago
C       fe80::/64 via ::, xe1/3, installed 01:43:56, last update 01:43:56 ago
IP Route Table for VRF "management"
C       ::1/128 via ::, lo.management, installed 21:28:46, last update
21:28:46 ago
C       fe80::/64 via ::, eth0, installed 21:28:46, last update 21:28:46 ago
IP Route Table for VRF "l3_vrf_1"
C       ::1/128 via ::, lo.l3_vrf_1, installed 21:28:19, last update 21:28:19
ago
B       ::ffff:a0a:a07/128 [0/0] via ::, tunvxlan2, installed 00:41:34, last
update 00:41:34 ago
B       1001::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlan2),
installed 00:41:35, last update 00:41:35 ago
B       1001::2/128 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlan2),
installed 00:22:43, last update 00:22:43 ago
B       1002::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlan2),
installed 00:41:35, last update 00:41:35 ago
B       1002::2/128 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlan2),
installed 00:22:43, last update 00:22:43 ago
B       1003::/48 [20/0] via ::ffff:a0a:a07 (recursive via ::, tunvxlan2),
installed 00:41:35, last update 00:41:35 ago
C       2001::/48 via ::, irb1001, installed 01:06:59, last update 01:06:59
ago
C       2002::/48 via ::, irb1002, installed 01:06:59, last update 01:06:59
ago
C       2003::/48 via ::, irb1003, installed 01:06:59, last update 01:06:59
ago
C       fe80::/64 via ::, tunvxlan2, installed 01:04:25, last update 01:04:25
ago
IP Route Table for VRF "l2_vrf_1"
IP Route Table for VRF "l2_vrf_2"
IP Route Table for VRF "l2_vrf_3"
LEAF3#show bgp vrf all
BGP table version is 1, local router ID is 203.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-
Color

BGP Route Table for VRF l3_vrf_1
*>x 101.0.0.0        10.10.10.7           0        100       0   65002 65004 ?
*>x 101.1.1.2/32     10.10.10.7           0        100       0   65002 65004 i
*>x 101.1.1.240/32   10.10.10.7           0        100       0   65002 65004 ?
*>x 102.0.0.0        10.10.10.7           0        100       0   65002 65004 ?
*>x 102.1.1.2/32     10.10.10.7           0        100       0   65002 65004 i
```

```
*>x 103.0.0.0        10.10.10.7            0        100        0   65002 65004 ?
*>  201.0.0.0/8      0.0.0.0               0        100        32768  ?
*>  201.1.1.240/32   0.0.0.0               0        100        32768  ?
*>  202.0.0.0/8      0.0.0.0               0        100        32768  ?
*>  203.0.0.0/8      0.0.0.0               0        100        32768  ?

Total number of prefixes 10
*>  10.10.10.4/32    0.0.0.0               0        100        32768  i
*                    0.0.0.0               0        100        32768  ?
*>  10.10.10.7/32    fe80::6eb9:c5ff:fe15:878f     0        100        0
65002 65004 i
*>  10.10.10.16/32   fe80::6eb9:c5ff:fe15:878f     0        100        0
65002 i
*>  17.18.18.0/24    fe80::6eb9:c5ff:fe15:878f     0        100        0
65002 65004 ?
*>  17.19.19.0/24    fe80::6eb9:c5ff:fe15:878f     0        100        0
65002 65004 ?

Total number of prefixes 5
LEAF3#show bgp l2vpn evpn
BGP table version is 34, local router ID is 10.10.10.4
Status codes: s suppressed, d damped, h history, a add-path, b back-up, *
valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight
Path  Peer          Encap

RD[7007:5001]
*>  [5]:[0]:[0]:[8]:[101.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.7            0        100        0   65002 65004
?  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[0]:[0]:[8]:[102.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.7            0        100        0   65002 65004
?  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[0]:[0]:[8]:[103.0.0.0]:[0.0.0.0]:[1000]
                       10.10.10.7            0        100        0   65002 65004
?  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[32]:[101.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.7            0        100        0   65002 65004
i  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[0]:[0]:[32]:[101.1.1.240]:[0.0.0.0]:[1000]
                       10.10.10.7            0        100        0   65002 65004
?  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[32]:[102.1.1.2]:[0.0.0.0]:[1000]
                       10.10.10.7            0        100        0   65002 65004
i  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[0]:[0]:[48]:[1001::]:[::]:[1000]
                       10.10.10.7            0        100        0   65002 65004
?  fe80::6eb9:c5ff:fe15:878f VXLAN
```

```
*>  [5]:[0]:[0]:[48]:[1002::]:[::]:[1000]
                       10.10.10.7           0        100       0   65002 65004
?  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[0]:[0]:[48]:[1003::]:[::]:[1000]
                       10.10.10.7           0        100       0   65002 65004
?  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[128]:[1001::2]:[::]:[1000]
                       10.10.10.7           0        100       0   65002 65004
i  fe80::6eb9:c5ff:fe15:878f VXLAN
*>  [5]:[00:00:00:aa:aa:bb:bb:00:00:00]:[0]:[128]:[1002::2]:[::]:[1000]
                       10.10.10.7           0        100       0   65002 65004
i  fe80::6eb9:c5ff:fe15:878f VXLAN

RD[10.10.10.4:1] VRF[l2_vrf_1]
*>  [2]:[0]:[1001]:[48,0000:0201:0001]:[0]:[1001]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1001]:[48,0000:0201:0001]:[32,201.1.1.2]:[1001]:[1000]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1001]:[48,0000:0201:0001]:[128,2001::2][1001]:[1000]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1001]:[48,0000:0201:0001]:[128,fe80::1][1001]:[1000]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1001]:[48,e001:a6a9:4b01]:[32,201.1.1.1]:[1001]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1001]:[48,e001:a6a9:4b01]:[128,2001::1][1001]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1001]:[48,e001:a6a9:4b01]:[128,fe80::e201:a6ff:fea9:4b01][1001]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [3]:[1001]:[32,10.10.10.4]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN

RD[10.10.10.4:2] VRF[l2_vrf_2]
*>  [2]:[0]:[1002]:[48,0000:0202:0001]:[0]:[1002]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1002]:[48,0000:0202:0001]:[32,202.1.1.2]:[1002]:[1000]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1002]:[48,0000:0202:0001]:[128,2002::2][1002]:[1000]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1002]:[48,e001:a6a9:4b01]:[32,202.1.1.1]:[1002]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1002]:[48,e001:a6a9:4b01]:[128,2002::1][1002]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1002]:[48,e001:a6a9:4b01]:[128,fe80::e201:a6ff:fea9:4b01][1002]
                       10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [3]:[1002]:[32,10.10.10.4]
```

```
                         10.10.10.4           0        100       32768  i  -----
-----      VXLAN

RD[10.10.10.4:3] VRF[l2_vrf_3]
*>  [2]:[0]:[1003]:[48,0000:0203:0001]:[0]:[1003]
                         10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1003]:[48,0000:0203:0001]:[32,203.1.1.2]:[1003]:[1000]
                         10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1003]:[48,0000:0203:0001]:[128,2003::2][1003]:[1000]
                         10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1003]:[48,e001:a6a9:4b01]:[32,203.1.1.1]:[1003]
                         10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1003]:[48,e001:a6a9:4b01]:[128,2003::1][1003]
                         10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [2]:[0]:[1003]:[48,e001:a6a9:4b01]:[128,fe80::e201:a6ff:fea9:4b01][1003]
                         10.10.10.4           0        100       32768  i  -----
-----      VXLAN
*>  [3]:[1003]:[32,10.10.10.4]
                         10.10.10.4           0        100       32768  i  -----
-----      VXLAN

Total number of prefixes 33
LEAF3#show nvo vxlan mac-table
================================================================================
================================================================================
=====
                                                          VXLAN MAC Entries
================================================================================
================================================================================
=====
VNID    Interface VlanId   In-VlanId Mac-Addr    VTEP-Ip/ESI
Type              Status     MAC move AccessPortDesc LeafFlag

--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
------

1001    xe34    2001    ----    0000.0201.0001 10.10.10.4
Dynamic Local   -------    0        -------           ----
1001    irb1001  ----    ----    e001.a6a9.4b01 10.10.10.4
Static Local    -------    0        -------           ----
1002    xe34    2002    ----    0000.0202.0001 10.10.10.4
Dynamic Local   -------    0        -------           ----
1002    irb1002  ----    ----    e001.a6a9.4b01 10.10.10.4
Static Local    -------    0        -------           ----
1003    xe34    2003    ----    0000.0203.0001 10.10.10.4
Dynamic Local   -------    0        -------           ----
1003    irb1003  ----    ----    e001.a6a9.4b01 10.10.10.4
Static Local    -------    0        -------           ----

Total number of entries are : 6

LEAF3#
LEAF3#
LEAF3#
```

R1

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            10.10.10.5/32 [110/2] via 17.20.20.1, xe1/1, installed 14:19:43,
last update 14:19:43 ago
O            10.10.10.6/32 [110/2] via 17.22.22.1, xe31/1, installed 14:22:04,
last update 14:22:04 ago
O            10.10.10.7/32 [110/2] via 17.19.19.1, xe1/2, installed 14:22:35,
last update 14:22:35 ago
C            10.10.10.21/32 is directly connected, lo, installed 1d16h57m,
last update 1d16h57m ago
O            10.10.10.22/32 [110/2] via 17.24.24.1, xe6/1, installed 1d16h56m,
last update 1d16h56m ago
O            17.18.18.0/24 [110/2] via 17.24.24.1, xe6/1, installed 14:23:30,
last update 14:22:35 ago
                           [110/2] via 17.19.19.1, xe1/2
C            17.19.19.0/24 is directly connected, xe1/2, installed 14:23:30,
last update 14:23:30 ago
C            17.20.20.0/24 is directly connected, xe1/1, installed 14:20:29,
last update 14:20:29 ago
O            17.21.21.0/24 [110/2] via 17.20.20.1, xe1/1, installed 00:19:21,
last update 00:19:21 ago
                           [110/2] via 17.24.24.1, xe6/1
C            17.22.22.0/24 is directly connected, xe31/1, installed 14:22:49,
last update 14:22:49 ago
O            17.23.23.0/24 [110/2] via 17.22.22.1, xe31/1, installed 14:22:45,
last update 14:22:04 ago
                           [110/2] via 17.24.24.1, xe6/1
C            17.24.24.0/24 is directly connected, xe6/1, installed 1d16h57m,
last update 1d16h57m ago
C            127.0.0.0/8 is directly connected, lo, installed 1d16h58m, last
update 1d16h58m ago

Gateway of last resort is not set
R1#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.21, local AS number 65005
BGP table version is 350
4 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/
Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
10.10.10.5        4 65004       5796       5965     346      0      0
14:19:40           13      0      0      0      0     13
10.10.10.6        4 65004       5840       5925     350      0      0
14:21:54           13      0      0      0      0     13
10.10.10.7        4 65004       5851       5952     350      0      0
14:22:31           14      0      0      0      0     14
10.10.10.22       4 65006       5985       6001     350      0      0
1d16h56m           30      0      0      0      0     30
```

```
Total number of neighbors 4

Total number of Established sessions 4
R1#show ip ospf neighbor

Total number of full neighbors: 4
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.10.10.7        1   Full/Backup     00:00:31    17.19.19.1      xe1/2
0
10.10.10.5        1   Full/Backup     00:00:32    17.20.20.1      xe1/1
0
10.10.10.6        1   Full/Backup     00:00:31    17.22.22.1      xe31/1
0
10.10.10.22       1   Full/DR         00:00:30    17.24.24.1      xe6/1
0
R1#
```

R2

```
R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            10.10.10.5/32 [110/2] via 17.21.21.1, xe1/1, installed 00:18:49,
last update 00:18:49 ago
O            10.10.10.6/32 [110/2] via 17.23.23.1, xe31/1, installed 14:22:06,
last update 14:22:06 ago
O            10.10.10.7/32 [110/2] via 17.18.18.1, xe1/2, installed 14:22:58,
last update 14:22:58 ago
O            10.10.10.21/32 [110/2] via 17.24.24.2, xe6/1, installed 1d16h56m,
last update 1d16h56m ago
C            10.10.10.22/32 is directly connected, lo, installed 1d16h58m,
last update 1d16h58m ago
C            17.18.18.0/24 is directly connected, xe1/2, installed 14:23:43,
last update 14:23:43 ago
O            17.19.19.0/24 [110/2] via 17.24.24.2, xe6/1, installed 14:23:43,
last update 14:22:58 ago
                           [110/2] via 17.18.18.1, xe1/2
O            17.20.20.0/24 [110/2] via 17.24.24.2, xe6/1, installed 14:20:41,
last update 00:18:49 ago
                           [110/2] via 17.21.21.1, xe1/1
C            17.21.21.0/24 is directly connected, xe1/1, installed 00:19:34,
last update 00:19:34 ago
O            17.22.22.0/24 [110/2] via 17.23.23.1, xe31/1, installed 14:22:57,
last update 14:22:06 ago
                           [110/2] via 17.24.24.2, xe6/1
C            17.23.23.0/24 is directly connected, xe31/1, installed 14:23:02,
last update 14:23:02 ago
```

```
C            17.24.24.0/24 is directly connected, xe6/1, installed 1d16h57m,
last update 1d16h57m ago
C            127.0.0.0/8 is directly connected, lo, installed 1d16h58m, last
update 1d16h58m ago

Gateway of last resort is not set
R2#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.22, local AS number 65006
BGP table version is 318
4 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/
Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
10.10.10.5         4 65004       5793       5951     316       0      0
00:18:48            17      0      0      0       0     17
10.10.10.6         4 65004       5830       5915     318       0      0
14:22:05            17      0      0      0       0     17
10.10.10.7         4 65004       5831       5948     318       0      0
14:22:52            26      0      0      0       0     26
10.10.10.21        4 65005       6001       5986     318       0      0
1d16h56m            30      0      0      0       0     30

Total number of neighbors 4

Total number of Established sessions 4
R2#show ip ospf neighbor

Total number of full neighbors: 4
OSPF process 100 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address         Interface
Instance ID
10.10.10.7       1   Full/Backup     00:00:35   17.18.18.1     xe1/2
0
10.10.10.5       1   Full/Backup     00:00:31   17.21.21.1     xe1/1
0
10.10.10.6       1   Full/Backup     00:00:32   17.23.23.1     xe31/1
0
10.10.10.21      1   Full/Backup     00:00:29   17.24.24.2     xe6/1
0
R2#
```

SPINE1

```
SPINE1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::1644:8fff:fec8:fd5, xe8, installed
14:21:54, last update 14:21:54 ago
B            10.10.10.2/32 [20/0] via fe80::9a19:2cff:fe9b:37d7, xe3,
installed 16:58:02, last update 16:58:02 ago
```

```
B            10.10.10.5/32 [20/0] via fe80::36ef:b6ff:fe69:fc42, ce49,
installed 14:19:13, last update 10:09:17 ago
                           [20/0] via fe80::230:abff:feea:da39, xe45
B            10.10.10.6/32 [20/0] via fe80::36ef:b6ff:fe69:fc42, ce49,
installed 14:21:30, last update 10:09:17 ago
                           [20/0] via fe80::230:abff:feea:da39, xe45
C            10.10.10.11/32 is directly connected, lo, installed 17:06:36,
last update 17:06:36 ago
B            10.10.10.12/32 [20/0] via fe80::9a19:2cff:fe9b:37d7, xe3,
installed 17:06:15, last update 07:13:32 ago
                            [20/0] via fe80::36ef:b6ff:fe69:fc42, ce49
                            [20/0] via fe80::1644:8fff:fec8:fd5, xe8
                            [20/0] via fe80::230:abff:feea:da39, xe45
B            10.10.10.14/32 [20/0] via fe80::230:abff:feea:da39, xe45,
installed 17:05:02, last update 17:05:02 ago
B            10.10.10.15/32 [20/0] via fe80::36ef:b6ff:fe69:fc42, ce49,
installed 10:09:17, last update 10:09:17 ago
B            17.20.20.0/24 [20/0] via fe80::36ef:b6ff:fe69:fc42, ce49,
installed 14:19:13, last update 10:09:17 ago
                           [20/0] via fe80::230:abff:feea:da39, xe45
B            17.21.21.0/24 [20/0] via fe80::36ef:b6ff:fe69:fc42, ce49,
installed 00:18:22, last update 00:18:22 ago
                           [20/0] via fe80::230:abff:feea:da39, xe45
B            17.22.22.0/24 [20/0] via fe80::36ef:b6ff:fe69:fc42, ce49,
installed 14:21:30, last update 10:09:17 ago
                           [20/0] via fe80::230:abff:feea:da39, xe45
B            17.23.23.0/24 [20/0] via fe80::36ef:b6ff:fe69:fc42, ce49,
installed 14:21:30, last update 10:09:17 ago
                           [20/0] via fe80::230:abff:feea:da39, xe45
C            127.0.0.0/8 is directly connected, lo, installed 17:07:05, last
update 17:07:05 ago

Gateway of last resort is not set
SPINE1#show ip bgp summary
BGP router identifier 10.10.10.11, local AS number 65002
BGP table version is 32
10 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe3(fe80::9a19:2cff:fe9b:37d7)          4 65001      2580       2817     32
0      0  16:58:04            5
ce49(fe80::36ef:b6ff:fe69:fc42)         4 65003      1548       1570     32
0      0  10:09:19            8
xe8(fe80::1644:8fff:fec8:fd5)           4 65001      2568       2766     32
0      0  14:21:56            2
xe45(fe80::230:abff:feea:da39)          4 65003      2798       2745     32
0      0  17:05:04            9

Total number of neighbors 4

Total number of Established sessions 4
SPINE1#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.11, local AS number 65002
BGP table version is 345
10 BGP AS-PATH entries
0 BGP community entries
```

```
Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd      AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
xe3(fe80::9a19:2cff:fe9b:37d7)            4 65001      2580      2817    342
0      0  16:58:06            35      4     16      3      1     11
ce49(fe80::36ef:b6ff:fe69:fc42)           4 65003      1548      1570    345
0      0  10:09:21            25      2      3      0      0     20
xe8(fe80::1644:8fff:fec8:fd5)             4 65001      2568      2766    345
0      0  14:21:58            31      4      9      3      1     14
xe45(fe80::230:abff:feea:da39)            4 65003      2798      2745    345
0      0  17:05:06            20      0      0      0      0     20

Total number of neighbors 4

Total number of Established sessions 4
SPINE1#
```

SPINE2

```
SPINE2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::1644:8fff:fec8:ffa, xe47,
installed 14:22:01, last update 14:22:01 ago
B            10.10.10.2/32 [20/0] via fe80::9a19:2cff:fe9b:37d6, xe1,
installed 16:58:10, last update 16:58:10 ago
B            10.10.10.5/32 [20/0] via fe80::36ef:b6ff:fe69:fc66, ce50,
installed 14:19:20, last update 10:09:24 ago
                           [20/0] via fe80::230:abff:feea:da38, xe57
B            10.10.10.6/32 [20/0] via fe80::36ef:b6ff:fe69:fc66, ce50,
installed 14:21:37, last update 10:09:24 ago
                           [20/0] via fe80::230:abff:feea:da38, xe57
B            10.10.10.11/32 [20/0] via fe80::9a19:2cff:fe9b:37d6, xe1,
installed 17:05:09, last update 07:14:11 ago
                            [20/0] via fe80::36ef:b6ff:fe69:fc66, ce50
                            [20/0] via fe80::1644:8fff:fec8:ffa, xe47
                            [20/0] via fe80::230:abff:feea:da38, xe57
C            10.10.10.12/32 is directly connected, lo, installed 1d16h57m,
last update 1d16h57m ago
B            10.10.10.14/32 [20/0] via fe80::230:abff:feea:da38, xe57,
installed 17:05:09, last update 17:05:09 ago
B            10.10.10.15/32 [20/0] via fe80::36ef:b6ff:fe69:fc66, ce50,
installed 10:09:24, last update 10:09:24 ago
B            17.20.20.0/24 [20/0] via fe80::36ef:b6ff:fe69:fc66, ce50,
installed 14:19:20, last update 10:09:24 ago
                           [20/0] via fe80::230:abff:feea:da38, xe57
B            17.21.21.0/24 [20/0] via fe80::36ef:b6ff:fe69:fc66, ce50,
installed 00:18:30, last update 00:18:29 ago
                           [20/0] via fe80::230:abff:feea:da38, xe57
```

```
B            17.22.22.0/24 [20/0] via fe80::36ef:b6ff:fe69:fc66, ce50,
installed 14:21:37, last update 10:09:24 ago
                           [20/0] via fe80::230:abff:feea:da38, xe57
B            17.23.23.0/24 [20/0] via fe80::36ef:b6ff:fe69:fc66, ce50,
installed 14:21:37, last update 10:09:24 ago
                           [20/0] via fe80::230:abff:feea:da38, xe57
C            127.0.0.0/8 is directly connected, lo, installed 1d16h57m, last
update 1d16h57m ago

Gateway of last resort is not set
SPINE2#show ip bgp summary
BGP router identifier 10.10.10.12, local AS number 65002
BGP table version is 39
11 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe57(fe80::230:abff:feea:da38)            4 65003      3123       3461     39
0      0  17:05:11            10
xe47(fe80::1644:8fff:fec8:ffa)            4 65001      6065       6251     39
0      0  14:22:03            11
ce50(fe80::36ef:b6ff:fe69:fc66)           4 65003      1797       2110     39
0      0  10:09:26            11
xe1(fe80::9a19:2cff:fe9b:37d6)            4 65001      2727       3096     39
0      0  16:58:12             8

Total number of neighbors 4

Total number of Established sessions 4
SPINE2#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.12, local AS number 65002
BGP table version is 406
11 BGP AS-PATH entries
0 BGP community entries

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
xe57(fe80::230:abff:feea:da38)            4 65003      3123       3461     406
0      0  17:05:13           75      8     25      6      2     34
xe47(fe80::1644:8fff:fec8:ffa)            4 65001      6068       6251     406
0      0  14:22:05           68      8     25      6      2     27
ce50(fe80::36ef:b6ff:fe69:fc66)           4 65003      1797       2111     406
0      0  10:09:28           75      8     25      6      2     34
xe1(fe80::9a19:2cff:fe9b:37d6)            4 65001      2727       3096     406
0      0  16:58:14           71      8     25      6      2     30

Total number of neighbors 4

Total number of Established sessions 4
SPINE2#
SPINE2#
```

SPINE3

```
SPINE3#show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.4/32 [20/0] via fe80::e201:a6ff:fea9:4b05, ce49,
installed 16:58:48, last update 16:58:48 ago
B            10.10.10.7/32 [20/0] via fe80::1644:8fff:fec9:7ce, xe7, installed
14:22:31, last update 14:22:31 ago
C            10.10.10.16/32 is directly connected, lo, installed 1d16h56m,
last update 1d16h56m ago
B            17.18.18.0/24 [20/0] via fe80::1644:8fff:fec9:7ce, xe7, installed
14:22:31, last update 14:22:31 ago
B            17.19.19.0/24 [20/0] via fe80::1644:8fff:fec9:7ce, xe7, installed
14:22:31, last update 14:22:31 ago
C            127.0.0.0/8 is directly connected, lo, installed 1d16h56m, last
update 1d16h56m ago

Gateway of last resort is not set
SPINE3#show ip bgp summary
BGP router identifier 10.10.10.16, local AS number 65002
BGP table version is 9
3 BGP AS-PATH entries
0 BGP community entries
4  Configured ibgp ECMP multipath: Currently set at 4

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
xe7(fe80::1644:8fff:fec9:7ce)           4 65004      5834       5808       9
0      0  14:22:35            3
ce49(fe80::e201:a6ff:fea9:4b05)         4 65001      2611       3116       9
0      0  16:58:52            1

Total number of neighbors 2

Total number of Established sessions 2
SPINE3#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.16, local AS number 65002
BGP table version is 73
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
xe7(fe80::1644:8fff:fec9:7ce)           4 65004      5834       5808     70
0      0  14:22:36           10      0      0      0      0     10
ce49(fe80::e201:a6ff:fea9:4b05)         4 65001      2611       3116     73
0      0  16:58:53           26      0     16      3      0      7

Total number of neighbors 2

Total number of Established sessions 2
SPINE3#
```

SuperSpine1

```
SuperSpine1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::ba6a:97ff:fee2:4bc6, ce1/3,
installed 14:22:16, last update 07:52:59 ago
                           [20/0] via fe80::1644:8fff:fe45:8df5, ce1/4
B            10.10.10.2/32 [20/0] via fe80::ba6a:97ff:fee2:4bc6, ce1/3,
installed 16:58:25, last update 07:52:59 ago
                           [20/0] via fe80::1644:8fff:fe45:8df5, ce1/4
B            10.10.10.5/32 [20/0] via fe80::5e17:83ff:feff:2276, ce31/1,
installed 14:19:36, last update 14:19:36 ago
B            10.10.10.6/32 [20/0] via fe80::7e8d:9cff:feaa:524e, ce29/1,
installed 14:21:53, last update 14:21:53 ago
B            10.10.10.11/32 [20/0] via fe80::1644:8fff:fe45:8df5, ce1/4,
installed 17:05:25, last update 17:05:25 ago
B            10.10.10.12/32 [20/0] via fe80::ba6a:97ff:fee2:4bc6, ce1/3,
installed 17:05:25, last update 17:05:25 ago
C            10.10.10.14/32 is directly connected, lo, installed 17:05:59,
last update 17:05:59 ago
B            10.10.10.15/32 [20/0] via fe80::ba6a:97ff:fee2:4bc6, ce1/3,
installed 07:42:10, last update 07:41:47 ago
                            [20/0] via fe80::7e8d:9cff:feaa:524e, ce29/1
                            [20/0] via fe80::5e17:83ff:feff:2276, ce31/1
                            [20/0] via fe80::1644:8fff:fe45:8df5, ce1/4
B            17.20.20.0/24 [20/0] via fe80::5e17:83ff:feff:2276, ce31/1,
installed 14:19:36, last update 14:19:36 ago
B            17.21.21.0/24 [20/0] via fe80::5e17:83ff:feff:2276, ce31/1,
installed 00:18:45, last update 00:18:44 ago
B            17.22.22.0/24 [20/0] via fe80::7e8d:9cff:feaa:524e, ce29/1,
installed 14:21:53, last update 14:21:53 ago
B            17.23.23.0/24 [20/0] via fe80::7e8d:9cff:feaa:524e, ce29/1,
installed 14:21:53, last update 14:21:53 ago
C            127.0.0.0/8 is directly connected, lo, installed 17:06:48, last
update 17:06:48 ago

Gateway of last resort is not set
SuperSpine1#show ip bgp summary
BGP router identifier 10.10.10.14, local AS number 65003
BGP table version is 30
9 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                                  V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd   Desc
ce1/4(fe80::1644:8fff:fe45:8df5)        4 65002      2746      2798     30
0      0  17:05:28            6
```

```
ce31/1(fe80::5e17:83ff:feff:2276)       4 65004      2510      2910     30
0       0  14:19:39             4
ce29/1(fe80::7e8d:9cff:feaa:524e)       4 65004      2490      2813     30
0       0  14:21:56             4
ce1/3(fe80::ba6a:97ff:fee2:4bc6)        4 65002      2742      2764     30
0       0  17:05:28             5

Total number of neighbors 4

Total number of Established sessions 4
SuperSpine1#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.14, local AS number 65003
BGP table version is 337
9 BGP AS-PATH entries
0 BGP community entries

Neighbor                                 V    AS     MsgRcv     MsgSen  TblVer
InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
Desc
ce1/4(fe80::1644:8fff:fe45:8df5)         4 65002      2746      2799     335
0       0  17:05:29            69      8     25      6      2     28
ce31/1(fe80::5e17:83ff:feff:2276)        4 65004      2510      2910     336
0       0  14:19:40            16      0      6      0      0     10
ce29/1(fe80::7e8d:9cff:feaa:524e)        4 65004      2490      2813     336
0       0  14:21:57            16      0      6      0      0     10
ce1/3(fe80::ba6a:97ff:fee2:4bc6)         4 65002      2742      2764     336
0       0  17:05:29            66      8     25      6      2     25

Total number of neighbors 4

Total number of Established sessions 4
SuperSpine1#
SuperSpine1#
```

SuperSpine2

```
SuperSpine2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            10.10.10.1/32 [20/0] via fe80::ba6a:97ff:fee2:4bd0, ce11/1,
10:15:29
                           [20/0] via fe80::1644:8fff:fe45:8dfb, ce14/1
B            10.10.10.2/32 [20/0] via fe80::ba6a:97ff:fee2:4bd0, ce11/1,
10:15:29
                           [20/0] via fe80::1644:8fff:fe45:8dfb, ce14/1
B            10.10.10.5/32 [20/0] via fe80::5e17:83ff:feff:225e, ce15/1,
10:15:28
B            10.10.10.6/32 [20/0] via fe80::7e8d:9cff:feaa:524f, ce9/1,
10:15:28
```

```
B           10.10.10.11/32 [20/0] via fe80::1644:8fff:fe45:8dfb, ce14/1,
10:15:28
B           10.10.10.12/32 [20/0] via fe80::ba6a:97ff:fee2:4bd0, ce11/1,
10:15:28
B           10.10.10.14/32 [20/0] via fe80::ba6a:97ff:fee2:4bd0, ce11/1,
10:15:29
                           [20/0] via fe80::7e8d:9cff:feaa:524f, ce9/1
                           [20/0] via fe80::5e17:83ff:feff:225e, ce15/1
                           [20/0] via fe80::1644:8fff:fe45:8dfb, ce14/1
C           10.10.10.15/32 is directly connected, lo, 10:15:34
B           17.20.20.0/24 [20/0] via fe80::5e17:83ff:feff:225e, ce15/1,
10:15:28
B           17.21.21.0/24 [20/0] via fe80::5e17:83ff:feff:225e, ce15/1,
00:24:34
B           17.22.22.0/24 [20/0] via fe80::7e8d:9cff:feaa:524f, ce9/1,
10:15:28
B           17.23.23.0/24 [20/0] via fe80::7e8d:9cff:feaa:524f, ce9/1,
10:15:28
C           127.0.0.0/8 is directly connected, lo, 1d15h48m

Gateway of last resort is not set

SuperSpine2#show ip bgp summary
BGP router identifier 10.10.10.15, local AS number 65003
BGP table version is 14
12 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor                 V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ
Up/Down  State/PfxRcd
ce14/1                    4 65002       1571       1550      14      0      0
10:10:01            10
ce15/1                    4 65004       1571       1651      14      0      0
10:10:01             8
ce9/1                     4 65004       1533       1601      14      0      0
10:10:01             8
ce11/1                    4 65002       1583       1601      14      0      0
10:10:01            10

Total number of neighbors 4

Total number of Established sessions 4
SuperSpine2#
SuperSpine2#
SuperSpine2#show bgp l2vpn evpn summary
BGP router identifier 10.10.10.15, local AS number 65003
BGP table version is 108
12 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ
Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
ce14/1                    4 65002       1570       1549     107      0      0
10:09:50            61      8     25      6      2        22
ce15/1                    4 65004       1571       1651     108      0      0
10:09:50            51      8     19      0      0        24
ce9/1                     4 65004       1532       1601     108      0      0
10:09:50            51      8     19      0      0        24
```

```
ce11/1                    4 65002       1582       1600     108       0       0
10:09:50            64       8     25       6       2         23

Total number of neighbors 4

Total number of Established sessions 4
SuperSpine2#
```

EVPN L3 Gateway with VXLAN Stitching - Commands

evpn-stitching

Use this command to enable EVPN L3 gateway with VXLAN stitching.

Use `no` parameter of this command to disable EVPN L3 gateway with VXLAN stitching.

Command Syntax

```
evpn-stitching
```

```
no evpn-stitching
```

Parameters

None

Default

None

Command Mode

VRF mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example shows how to enable EVPN L3 gateway with VXLAN stitching:

```
#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#commit
(config-vrf-evpn)#

#configure terminal
(config)#ip vrf red
(config-vrf)#no evpn-stitching
(config-vrf)#commit
```

rd (route distinguisher)

Use this command to assign a route distinguisher (RD).

Use `no` parameter of this command to remove the configured RD.

Command Syntax

```
rd ASN:nn_or_IP-address:nn
no rd ASN:nn_or_IP-address:nn
```

Parameters

`ASN:nn_or_IP-address:nn`	AS number and an arbitrary number (for example, 100:1). Otherwise, specify a 32-bit IP address and an arbitrary number (for example, 192.16.10.1:1).

•

Default

None

Command Mode

evpn-stitching mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example shows how to configure an EVPN instance with VXLAN Stitching:

```
#configure terminal
(config)#ip vrf red
(config-vrf)#rd 101:101
(config-vrf)#evpn-stitching
(config-vrf-evpn)#rd 102:102
(config-vrf-evpn)#commit

#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#no rd 102:102
(config-vrf-evpn)#commit
```

route-target

Use this command to configure a route-target of type import, export, or both to the EVPN instance.

Use `no` parameter of this command to remove the configured the route-target.

Command Syntax

```
route-target (import|export|both) (ASN:nn_or_IP-address:nn)
```

```
no route-target (import|export|both) (ASN:nn_or_IP-address:nn)
```

Parameters

`import`	AImport routing information.
`export`	Export routing information.
`both`	Import and export routing information.
`ASN:nn_or_IP-address:nn`	AS number and an arbitrary number (for example, 100:1). Otherwise, specify a 32-bit IP address and an arbitrary number (for example, 192.16.10.1:1).

Default

None

Command Mode

evpn-stitching mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example is for configuring an EVPN instance with VXLAN stitching, including route-target import/export::

```
#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#rd 102:102
(config-vrf-evpn)#route-target both 102:102
(config-vrf-evpn)#commit

#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#no route-target both 102:102
(config-vrf-evpn)#commit
```

host-route-extend

Use this command to enable host route extension for a given L2VNI or all L2VNIs in EVPN stitching mode.

Use `no` parameter of this command to disable host route extension.

Command Syntax

```
host-route-extend-l2vni <VNI_RANGE | all>
no host-route-extend-l2vni
```

Parameters

VNI_RANGE	Specifies L2VNIs in the range <1-16777215> for which host route extension is enabled.
all	Applies host route extension to all configured L2VNIs.

Default

None

Command Mode

evpn-stitching mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example shows how to configure a :

```
#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#host-route-extend-l2vni all
(config-vrf-evpn)#commit

#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#no host-route-extend-l2vni all
(config-vrf-evpn)#commit

#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#host-route-extend-l2vni 1-2,5,10-14
(config-vrf-evpn)#commit

#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#no host-route-extend-l2vni 1-2,5,10-14
(config-vrf-evpn)#commit
```

translation-l3vni

Use this command to configure L3VNI translation from a local L3VNI to a DCI L3VNI.

Use `no` parameter of this command to disable the L3VNI translation.

Command Syntax

```
translation-l3vni <1-16777215>

no translation-l3vni
```

Parameters

<1-16777215>	Specifies the L3VNI identifier to be used for translation.

Default

None

Command Mode

evpn-stitching mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example shows how to configure an L3VNI translation:

```
#configure terminal
(config)#ip vrf red
(config-vrf)#l3vni 1000
(config-vrf)#evpn-stitching
(config-vrf-evpn)#translation-l3vni 300
(config-vrf-evpn)#commit
%% After map/unmap of L3VNI from VRF red, interfaces are unbounded with VRF
(config-vrf-evpn)#

#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#no translation-l3vni
(config-vrf-evpn)#commit
%% After map/unmap of L3VNI from VRF red, interfaces are unbounded with VRF
(config-vrf-evpn)#
```

exit-config-vrf-evpn-mode

Use this command to exit the EVPN stitching.

Command Syntax

```
exit-config-vrf-evpn-mode
```

Parameters

None

Default

None

Command Mode

evpn-stitching mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example shows how to exit an EVPN Stitching:

```
#configure terminal
(config)#ip vrf red
(config-vrf)#evpn-stitching
(config-vrf-evpn)#exit-config-vrf-evpn-mode
(config-vrf)#
```

Implementation Examples

Scenario 1: A service provider needs to extend Layer 2 services from a VXLAN-based data center fabric to an MPLS-based WAN domain without protocols merging. Within this architecture, the Border Gateway (BGW) or Stitching Gateway performs encapsulation translation while preserving EVPN route semantics.

Use Case 1: By establishing separate BGP sessions for the DC and WAN, the BGW facilitates the exchange of EVPN Type-2 MAC routes. The BGW imports MAC routes from the VXLAN fabric and re-advertises them to the MPLS domain with appropriate service labels. This ensures MAC/IP mobility and route consistency for workloads spanning across heterogeneous data centers.

Scenario 2: A customer wants traffic from an L3VPN site to reach specific hosts in a VXLAN-based data center where reachability is discovered via Type-2 MAC-IP routes.

Use Case 2: The Border Gateway installs a /32 prefix route derived from the received Type-2 MAC-IP route into the associated IP-VRF. It then redistributes this prefix into the VPNv4/v6 address family toward the DCI domain. This allows seamless Layer 3 connectivity where the BGW acts as the next hop for both the VXLAN fabric and the MPLS WAN.

Scenario 3: A multi-site data center fabric requires scalable Layer 3 routing between sites across an MPLS backbone.

Use Case 3: The Border Gateway re-originates EVPN Type-5 routes imported from the VXLAN fabric into the VPNv4/v6 address family to be advertised to the DCI domain. In the reverse direction, VPNv4/v6 routes from the WAN are imported into the IP-VRF and re-advertised as Type-5 routes into the VXLAN fabric. This avoids protocol conversion while maintaining independent scaling per domain.

Troubleshooting

1. The gateway will not function unless specific global NVO parameters are active and the system has been restarted.
   - Verify `nvo vxlan enable` and `nvo vxlan irb` are configured in the running configuration.
   - Confirm the node was rebooted after these commands were committed.
2. If the translation tunnel does not establish between sites, verify that the peer type fabric external is configured on the BLEAFs facing the core.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
ACL	Access Control List
BGP	Border Gateway Protocol
BLEAF	Border Leaf
DPATH	Domain Path
EVPN	Ethernet Virtual Private Network
IRB	Integrated Routing and Bridging
L3VNI	Layer 3 Virtual Network Identifier
NVO	Network Virtualization Overlay
QoS	Quality of Service
SOO	Site-of-Origin
VRF	Virtual Routing and Forwarding
VTEP	VXLAN Tunnel End Point
VXLAN	Virtual Extensible LAN

VxLAN - EVPN with Integrated Routing and Bridging Deployment Mode

CHAPTER 1 EVPN-VxLAN with IRB

Overview

An EVPN-based Integrated Routing and Bridging solution used for forwarding of intra-subnets and inter-subnets traffic.

There are 2 modes of IRB.

Symmetric IRB

In this mode, both the ingress and egress VTEPs perform layer-2 and layer-3 lookups (switching and routing). In this case, a given VTEP needs to learn the ARP and MAC-address entries only for tenant systems (TSs) across the tenant VxLAN network belonging to VNIDs attached to that VTEP.

Asymmetric IRB

In this mode, the ingress VTEP perform layer-2 and layer-3 lookups and egress VTEPs perform layer-2 lookups only.

The disadvantage of this mode is the need for each VTEP in the tenant network to be configured with all the VNIDs for that tenant irrespective of whether a given VTEP has TS attached for that VNID or not.

Three approaches are available to achieve IRB solution.

- Centralized Gateway
- Anycast Gateway
- Distributed Gateway

Topology

The procedures in this section use the topology in Figure 1-11.

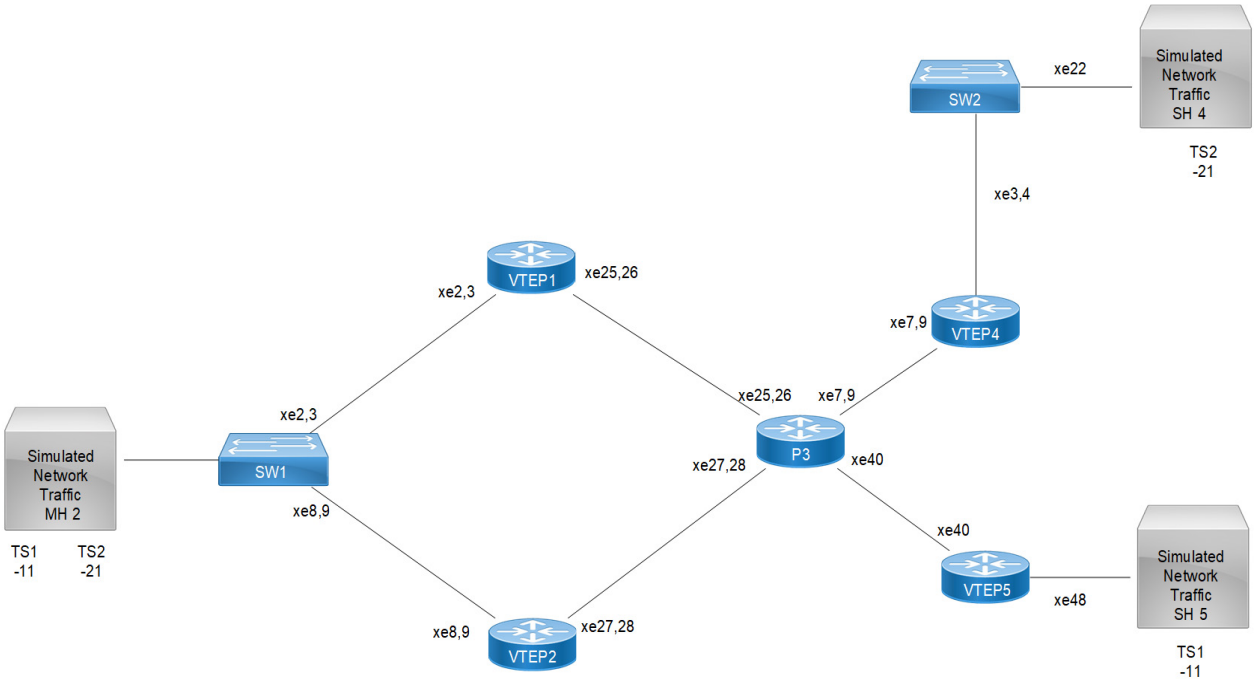

Figure 1-11: VxLAN EVPN IRB

Note: In the above topology TS1, TS2 are the tenant systems.

Base Configuration - L2 VXLAN

VTEP1

(Multi-homed group1) - Part of both Multi-homed with po1(MH2).

Generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(Config)#qos enable`	Enabling QoS
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

`(config)#interface po1`	Enter Interface mode for po1 (MH2)
`(config-if)#switchport`	Make it L2 interface
`(config-if)# evpn multi-homed system-mac 0000.0000.2222`	Configure system MAC as ESI value for Lag (po1) interface. VTEP1 and VTEP2 should have same ESI value
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe2`	Enter Interface mode for xe2
`(config-if)#channel-group 1 mode active`	Make it member port of po1

(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe3	Enter Interface mode for xe3
(config-if)#channel-group 1 mode active	Make it member port of po1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 1.1.1.1/32 secondary	Configure loopback IP address as 1.1.1.1 for VTEP1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po2	Enter Interface mode for po2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe25	Enter Interface mode for xe25
(config-if)#channel-group 2 mode active	Make it member port of po2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe26	Enter Interface mode for xe26
(config-if)#channel-group 2 mode active	Make it member port of po2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po2	Enter Interface mode for po2
(config-if)#ip address 100.11.11.1/24	Configure IP address as 100.11.11.1 on network side of Spine-P3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 1.1.1.1	Configure router-id as 1.1.1.1 (lo IP address)
(config-router)#network 1.1.1.1/32 area 0.0.0.0	Add 1.1.1.1 (lo IP address) network into area 0
(config-router)#network 100.11.11.0/24 area 0.0.0.0	Add 100.11.11.0 (Spine-P3) network into area 0
(config-router)#bfd all-interfaces	Enabling BFD on all OSPF interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

BGP configuration:

(Config)#router bgp 5000	Enter into Router BGP mode
(config-router)#bgp router-id 1.1.1.1	Configure router-id as 1.1.1.1 (lo IP address)
(config-router)#neighbor 2.2.2.2 remote-as 5000	Specify a VTEP2 loopback IP address and remote-as defined
(config-router)#neighbor 2.2.2.2 update-source lo	Configure update as loopback for VTEP2
(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP2

(config-router)#neighbor 4.4.4.4 remote-as 5000	Specify a VTEP4 loopback IP address and remote-as defined
(config-router)#neighbor 4.4.4.4 update-source lo	Configure update as loopback for VTEP4
(config-router)#neighbor 4.4.4.4 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP4
(config-router)#neighbor 5.5.5.5 remote-as 5000	Specify a VTEP5 loopback IP address and remote-as defined
(config-router)#neighbor 5.5.5.5 update-source lo	Configure update as loopback for VTEP5
(config-router)#neighbor 5.5.5.5 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP5
(config-router)#address-family l2vpn evpn	Enter into L2VPN EVPN address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2(VTEP2) into L2VPN EVPN address family mode
(config-router-af)#neighbor 4.4.4.4 activate	Activate 3.3.3.3(VTEP4) into L2VPN EVPN address family mode
(config-router-af)#neighbor 5.5.5.5 activate	Activate 5.5.5.5(VTEP5) into L2VPN EVPN address family mode
(config-router-af)#exit-address-family	Exit from L2VPN address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L2 MAC VRF Configuration:

(config)#mac vrf L2VRF1	Create MAC routing/forwarding instance with L2VRF1 name and enter into VRF mode
(config-vrf)#rd 1.1.1.1:11	Assign RD value
(config-vrf)#description MAC VRF RED	Give description to L2VRF1 as RED
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
(config-vrf)#exit	Exit from VRF mode
(config)#mac vrf L2VRF2	Create MAC routing/forwarding instance with L2VRF2 name and enter into VRF mode
(config-vrf)#rd 1.1.1.1:21	Assign RD value
(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#description MAC VRF BLUE	Give description to L2VRF2 as BLUE
(config-vrf)#exit	Exit from VRF mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L2 VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#evpn esi hold-time 90	Configure ESI hold time to allow tunnel to come up at the time of VXLAN initialization before making the ESI up. It should be same on both VTEP1 and VTEP2
(config)#nvo vxlan vtep-ip-global 1.1.1.1	Configure Source vtep-ip-global configuration - Use loopback IP address
(config)#nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# vni-name VNI-101	Configure VNI name as VNI-101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# vni-name VNI-201	Configure VNI name as VNI-201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan access-if port-vlan po1 10	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 101	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1010 ip 11.11.11.51	Configure static MAC-IP
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan po1 20	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 201	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1020 ip 21.21.21.51	Configure static MAC-IP
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#exit	Exit from configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VTEP2

(Multi-homed group1) - Part of both Multi-homed with po1(MH1).

Generic configuration:

#configure terminal	Enter Configure mode.
(config)#evpn vxlan multihoming enable	Enable Multihoming, save configs and reboot the board for multihoming to be effective

(Config)#qos enable	Enabling QoS
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

(config)#interface po1	Enter Interface mode for po1 (MH2)
(config-if)#switchport	Make it L2 interface
(config-if)# evpn multi-homed system-mac 0000.0000.2222	Configure system MAC as ESI value for Lag (po1) interface. VTEP1 and VTEP2 should have same ESI value
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe8	Enter Interface mode for xe2
(config-if)#channel-group 1 mode active	Make it member port of po1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe9	Enter Interface mode for xe3
(config-if)#channel-group 1 mode active	Make it member port of po1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 2.2.2.2/32 secondary	Configure loopback IP address as 2.2.2.2 for VTEP2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po3	Enter Interface mode for po3
(config-if)#switchport	Configure po3 as L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe27	Enter Interface mode for xe27
(config-if)#channel-group 3 mode active	Make it member port of po3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe28	Enter Interface mode for xe28
(config-if)#channel-group 3 mode active	Make it member port of po3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)# bridge 1 protocol rstp vlan-bridge	Configure bridge 1
(config)# vlan database	Enter into VLAN database mode
(config)#vlan 2 bridge 1 state enable	Configure VLAN 2 as part of bridge 1
(config)#interface po3	Enter Interface mode for po3
(config-if)# bridge-group 1	Configure bridge 1 for po3
(config-if)# switchport mode trunk	Switchport mode as trunk
(config-if)# switchport trunk allowed vlan add 2	Trunk allowed VLAN 2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface vlan1.2	Enter into SVI port VLAN1.2
(config-if)#ip address 100.12.12.1/24	Configure IP address as 100.12.12.1 on network side of Spine-P3

(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 2.2.2.2	Configure router-id as 2.2.2.2 (lo IP address)
(config-router)#network 2.2.2.2/32 area 0.0.0.0	Add 2.2.2.2 (lo IP address) network into area 0
(config-router)#network 100.12.12.0/24 area 0.0.0.0	Add 100.12.12.0(Spine-P3) network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.

BGP configuration:

(Config)#router bgp 5000	Enter into Router BGP mode
(config-router)#bgp router-id 2.2.2.2	Configure router-id as 2.2.2.2 (lo IP address)
(config-router)#neighbor 1.1.1.1 remote-as 5000	Specify a VTEP1 loopback IP address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Configure update as loopback for VTEP1
(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP1
(config-router)#neighbor 4.4.4.4 remote-as 5000	Specify a VTEP4 loopback IP address and remote-as defined
(config-router)#neighbor 4.4.4.4 update-source lo	Configure update as loopback for VTEP4
(config-router)#neighbor 4.4.4.4 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP4
(config-router)#neighbor 5.5.5.5 remote-as 5000	Specify a VTEP5 loopback IP address and remote-as defined
(config-router)#neighbor 5.5.5.5 update-source lo	Configure update as loopback for VTEP5
(config-router)#neighbor 5.5.5.5 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP5
(config-router)#address-family l2vpn evpn	Enter into L2VPN EVPN address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1(VTEP1) into L2VPN EVPN address family mode
(config-router-af)#neighbor 4.4.4.4 activate	Activate 4.4.4.4(VTEP4) into L2VPN EVPN address family mode
(config-router-af)#neighbor 5.5.5.5 activate	Activate 5.5.5.5(VTEP5) into L2VPN EVPN address family mode
(config-router-af)#exit-address-family	Exit from L2VPN address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VRF Configuration:

(config)#mac vrf L2VRF1	Create MAC routing/forwarding instance with L2VRF1 name and enter into VRF mode
(config-vrf)#rd 2.2.2.2:11	Assign RD value
(config-vrf)#description MAC VRF RED	Give description to L2VRF1 as RED
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
(config-vrf)#exit	Exit from VRF mode
(config)#mac vrf L2VRF2	Create MAC routing/forwarding instance with L2VRF2 name and enter into VRF mode
(config-vrf)#rd 2.2.2.2:21	Assign RD value
(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#description MAC VRF BLUE	Give description to L2VRF2 as BLUE
(config-vrf)#exit	Exit from VRF mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#evpn esi hold-time 90	Configure ESI hold time to allow tunnel to come up at the time of vxlan initialization before making the ESI up.It should be same on both VTEP1 and VTEP2
(config)#nvo vxlan vtep-ip-global 2.2.2.2	Configure Source vtep-IP-global configuration - Use loopback IP address
(config)#nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# vni-name VNI-101	Configure VNI name as VNI-101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# vni-name VNI-201	Configure VNI name as VNI-201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan access-if port-vlan po1 10	Enable port-VLAN mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 101	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1010 ip 11.11.11.51	Configure static MAC-IP
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode

(config)#nvo vxlan access-if port-vlan po1 20	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 201	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1020 ip 21.21.21.51	Configure static MAC-IP
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#exit	Exit from configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VTEP4

Single Home -SH5.

Generic configuration:

#configure terminal	Enter Configure mode.
(Config)#qos enable	Enabling qos

Interface and loopback configuration:

(config)#interface sa1	Enter Interface mode for sa1 (SH5)
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 4.4.4.4/32 secondary	Configure loopback IP address as 4.4.4.4 for VTEP4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po4	Enter Interface mode for po4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe7	Enter Interface mode for xe7
(config-if)#channel-group 4 mode active	Make it member port of po4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe9	Enter Interface mode for xe9
(config-if)#channel-group 4 mode active	Make it member port of po4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po4	Enter L3SI po4.4
(config-if)#ip address 100.14.14.1/24	Configure IP address as 100.14.14.1 on network side of Spine-P3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 4.4.4.4	Configure router-id as 4.4.4.4 (lo IP address)
(config-router)#network 4.4.4.4/32 area 0.0.0.0	Add 4.4.4.4 (lo IP address) network into area 0
(config-router)#network 100.14.14.0/24 area 0.0.0.0	Add 100.14.14.0(Spine-P3) network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

BGP configuration:

(Config)#router bgp 5000	Enter into Router BGP mode
(config-router)#bgp router-id 4.4.4.4	Configure router-id as 4.4.4.4 (lo IP address)
(config-router)#neighbor 1.1.1.1 remote-as 5000	Specify a VTEP1 loopback IP address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Configure update as loopback for VTEP1
(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP1
(config-router)#neighbor 2.2.2.2 remote-as 5000	Specify a VTEP2 loopback IP address and remote-as defined
(config-router)#neighbor 2.2.2.2 update-source lo	Configure update as loopback for VTEP2
(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP2
(config-router)#neighbor 5.5.5.5 remote-as 5000	Specify a VTEP5 loopback IP address and remote-as defined
(config-router)#neighbor 5.5.5.5 update-source lo	Configure update as loopback for VTEP5
(config-router)#neighbor 5.5.5.5 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP5
(config-router)#address-family l2vpn evpn	Enter into L2VPN EVPN address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1(VTEP1) into L2VPN EVPN address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2(VTEP2) into L2VPN EVPN address family mode
(config-router-af)#neighbor 5.5.5.5 activate	Activate 5.5.5.5(VTEP5) into L2VPN EVPN address family mode
(config-router-af)#exit-address-family	Exit from L2VPN address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VRF Configuration:

(config)#mac vrf L2VRF1	Create MAC routing/forwarding instance with L2VRF1 name and enter into VRF mode
(config-vrf)#rd 4.4.4.4:11	Assign RD value
(config-vrf)#description MAC VRF RED	Give description to L2VRF1 as RED
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
(config-vrf)#exit	Exit from VRF mode
(config)#mac vrf L2VRF2	Create MAC routing/forwarding instance with L2VRF2 name and enter into VRF mode
(config-vrf)#rd 4.4.4.4:21	Assign RD value
(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#description MAC VRF BLUE	Give description to L2VRF2 as BLUE
(config-vrf)#exit	Exit from VRF mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan vtep-ip-global 4.4.4.4	Configure Source vtep-IP-global configuration. Use loopback IP address
(config)#nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# vni-name VNI-101	Configure VNI name as VNI-101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# vni-name VNI-201	Configure VNI name as VNI-201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)# nvo vxlan access-if port-vlan sa1 20	Enable port-VLAN mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 201	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.5555.1020 ip 21.21.21.101	Configure static MAC-IP
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode

(config)#exit	Exit from configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VTEP5

Single Home -SH3

Generic configuration:

#configure terminal	Enter Configure mode.
(Config)#qos enable	Enabling qos
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

(config)#interface xe48	Enter Interface mode for xe48 (SH3)
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 5.5.5.5/32 secondary	Configure loopback IP address as 5.5.5.5 for VTEP5
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe40	Enter interface mode
(config-if)#ip address 100.15.15.1/24	Configure IP address as 100.15.15.1 on network side of Spine-P3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 5.5.5.5	Configure router-id as 5.5.5.5 (lo IP address)
(config-router)#network 5.5.5.5/32 area 0.0.0.0	Add 5.5.5.5 (lo IP address) network into area 0
(config-router)#network 100.15.15.0/24 area 0.0.0.0	Add 100.15.15.0(Spine-P3) network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

BGP configuration:

(Config)#router bgp 5000	Enter into Router BGP mode
(config-router)#bgp router-id 5.5.5.5	Configure router-id as 5.5.5.5(lo IP address)

(config-router)#neighbor 1.1.1.1 remote-as 5000	Specify a VTEP1 loopback IP address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Configure update as loopback for VTEP1
(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP1
(config-router)#neighbor 2.2.2.2 remote-as 5000	Specify a VTEP2 loopback IP address and remote-as defined
(config-router)#neighbor 2.2.2.2 update-source lo	Configure update as loopback for VTEP2
(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP2
(config-router)#neighbor 4.4.4.4 remote-as 5000	Specify a VTEP4 loopback IP address and remote-as defined
(config-router)#neighbor 4.4.4.4 update-source lo	Configure update as loopback for VTEP4
(config-router)#neighbor 4.4.4.4 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP4
(config-router)#address-family l2vpn evpn	Enter into L2VPN EVPN address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1(VTEP1) into L2VPN EVPN address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2(VTEP2) into L2VPN EVPN address family mode
(config-router-af)#neighbor 4.4.4.4 activate	Activate 4.4.4.4(VTEP4) into L2VPN EVPN address family mode
(config-router-af)#exit-address-family	Exit from L2VPN address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VRF Configuration:

(config)#mac vrf L2VRF1	Create MAC routing/forwarding instance with L2VRF1 name and enter into VRF mode
(config-vrf)#rd 5.5.5.5:11	Assign RD value
(config-vrf)#description MAC VRF RED	Give description to L2VRF1 as RED
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
(config-vrf)#exit	Exit from VRF mode
(config)#mac vrf L2VRF2	Create MAC routing/forwarding instance with L2VRF2 name and enter into VRF mode
(config-vrf)#rd 5.5.5.5:21	Assign RD value
(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#description MAC VRF BLUE	Give description to L2VRF2 as BLUE
(config-vrf)#exit	Exit from VRF mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan vtep-ip-global 5.5.5.5	Configure Source vtep-IP-global configuration. Use loopback IP address
(config)#nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# vni-name VNI-101	Configure VNI name as VNI-101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for EVPN-BGP to carry EVPN route
(config-nvo)# vni-name VNI-201	Configure VNI name as VNI-201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)# nvo vxlan access-if port-vlan xe48 10	Enable port-VLAN mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 101	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.4444.1010 ip 11.11.11.201	Configure static MAC-IP
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#exit	Exit from configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Switch1 (MH2)

Multihomed to 2-VTEPs (VTEP1 and VTEP2). It acts as Tenant system for VLAN1.20.

#configure terminal	Enter Configure mode.
(config)# bridge 1 protocol rstp vlan-bridge	Configure RSTP VLAN bridge
(config)# vlan database	Enter into VLAN database mode
(config)#vlan 2-20 bridge 1 state enable	Configure VLANs from 2-20 and associate with bridge 1
(config)#interface xe7	Enter Interface mode for xe7
(config-if)#switchport	Make as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate bridge 1 into interface
(config-if)# bridge-group 1 spanning-tree disable	Configure interface as STP disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 2,10,20	Trunk allowed VLAN as 2.10.20
(config-if)#switchport trunk native vlan 2	Native VLAN as 2

(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po1	Enter Interface mode for po1
(config-if)#switchport	Make po1 as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate po1 to bridge 1
(config-if)# bridge-group 1 spanning-tree disable	Configure po1 as STP disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 2,10,20	Trunk allowed VLAN as 2.10.20
(config-if)#switchport trunk native vlan 2	Native VLAN as 2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe3	Enter Interface mode for xe3
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#interface xe4	Enter Interface mode for xe4
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#exit	Exit from configuration mode
(config)#interface xe9	Enter Interface mode for xe9
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#interface xe10	Enter Interface mode for xe10
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#exit	Exit from configuration mode
(config)#interface vlan1.20	Enter Interface mode for VLAN1.20
(config-if)# ip address 21.21.21.2/24	Configure IP address
(config-if)#ipv6 address 21:21::21:2/48	Configure IPv6 address
(config)#exit	Exit from configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Switch2 (SH5)

#configure terminal	Enter Configure mode.
(config)# bridge 1 protocol rstp vlan-bridge	Configure RSTP VLAN bridge
(config)# vlan database	
(config)#vlan 2-20 bridge 1 state enable	Configure VLANs from 2-20 and associate with bridge 1
(config)#interface xe22	Enter Interface mode for xe22
(config-if)#switchport	Make xe22 as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate xe22 to bridge 1
(config-if)# bridge-group 1 spanning-tree disable	Configure xe22 as STP disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 6,10,20	Trunk allowed VLAN as 6,.10 &.20
(config-if)#switchport trunk native vlan 6	Native VLAN as 6

(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface sa1	Enter Interface mode for sa11
(config-if)#switchport	Make sa1 as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate sa1 to bridge 1
(config-if)# bridge-group 1 spanning-tree disable	Configure sa1 as STP disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 6,10,20	Trunk allowed VLAN as 2,.10.& 20
(config-if)#switchport trunk native vlan 6	Native VLAN as 6
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe3	Enter Interface mode for xe3
(config-if)# static-channel-group 1	Make it member port of sa1
(config)#interface xe4	Enter Interface mode for xe4
(config-if)# static-channel-group 1	Make it member port of sa1
(config)#exit	Exit from configuration mode
(config)#interface vlan1.20	Enter Interface mode for VLAN1.20
(config-if)# ip address 21.21.21.3/24	Configure IP address
(config-if)#ipv6 address 21:21::21:3/48	Configure IPv6 address
(config)#exit	Exit from configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Spine-P3

Spine node where all VTEPs are connected.

Generic configuration:

#configure terminal	Enter Configure mode.
(Config)#qos enable	Enabling qos
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 100.100.100.100/32 secondary	Configure loopback IP address as 100.100.100.100 for Spine-P3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po2	Enter Interface mode for po2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe25	Enter Interface mode for xe25
(config-if)#channel-group 2 mode active	Make it member port of po2
(config-if)#exit	Exit Interface mode and return to Configure mode.

(config)#interface xe26	Enter Interface mode for xe26
(config-if)#channel-group 2 mode active	Make it member port of po3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po2	Enter Interface mode for po2
(config-if)#ip address 100.11.11.2/24	Configure IP address as 100.11.11.2 on network side of VTEP1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po3	Enter Interface mode for po3
(config)#i switchport	Configure po3 as L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe27	Enter Interface mode for xe27
(config-if)#channel-group 3 mode active	Make it member port of po3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe28	Enter Interface mode for xe28
(config-if)#channel-group 3 mode active	Make it member port of po3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)# bridge 1 protocol rstp vlan-bridge	Configure bridge 1
(config)# vlan database	Enter into VLAN database mode
(config)#vlan 2 bridge 1 state enable	Configure VLAN 2 as part of bridge 1
(config)#interface po3	Enter Interface mode for po3
(config-if)# bridge-group 1	Configure bridge 1 for po3
(config-if)# switchport mode trunk	Switchport mode as trunk
(config-if)# switchport trunk allowed vlan add 2	Trunk allowed VLAN 2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface vlan1.2	Enter into SVI port VLAN1.2
(config-if)#ip address 100.12.12.2/24	Configure IP address as 100.12.12.2 on network side of VTEP2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po4	Enter Interface mode for po4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe7	Enter Interface mode for xe7
(config-if)#channel-group 4 mode active	Make it member port of po4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe9	Enter Interface mode for xe9
(config-if)#channel-group 4 mode active	Make it member port of po4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po4	Enter L3SI po4.4
(config-if)#ip address 100.14.14.2/24	Configure IP address as 100.14.14.12 on network side of VTEP4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe40	Enter interface mode

(config-if)#ip address 100.15.15.1/24	Configure IP address as 100.15.15.1 on network side of VTEP5
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 100.100.100.100	Configure router-id as 100.100.100.100 (lo IP address)
(config-router)#network 100.100.100.100/32 area 0.0.0.0	Add 100.100.100.100 (lo IP address) network into area 0
(config-router)#network 100.11.11.0/24 area 0.0.0.0	Add 100.11.11.0 (VTEP1) network into area 0
(config-router)#network 100.12.12.0/24 area 0.0.0.0	Add 100.12.12.0 (VTEP2) network into area 0
(config-router)#network 100.14.14.0/24 area 0.0.0.0	Add 100.14.14.0 (VTEP4) network into area 0
(config-router)#network 100.15.15.0/24 area 0.0.0.0	Add 100.15.15.0 (VTEP5) network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Centralized Gateway

In Centralized gateway approach, when two TS belonging to two different subnets connected to the same/different VTEP node, wanted to communicate with each other, their traffic needed to be back hauled from the VTEP node to the centralized gateway node where inter- subnet switching is performed and then back to the VTEP node.

IRB Configuration for Centralized Gateway

Configure from Base Configuration-L2 VXLAN section, then configure below commands for centralized gateway approach.

VTEP5

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#ip vrf L3VRF1	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
(config-vrf)#rd 51000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.

(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from VRF mode
(config)# interface irb1001	Configure IRB interface 1001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 11.11.11.1/24	Configure IP address
(config-if)#ipv6 address 11:11::11:1/48	Configure IPv6 address
(config-if)#exit	Exit from interface config mode
(config)# interface irb2001	Configure IRB interface 2001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 21.21.21.1/24	Configure IP address
(config-if)#ipv6 address 21:21::21:1/48	Configure IPv6 address
(config-if)#exit	Exit from interface config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#router bgp 5000	Enter into BGP router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit from address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb1001	Configure IRB1001 under VXLAN ID 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb2001	Configure IRB2001 under VXLAN ID 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Validation

VTEP5

```
TB2-VTEP5#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
5.5.5.5          2.2.2.2          Installed        00:26:30       00:26:30
5.5.5.5          4.4.4.4          Installed        00:26:30       00:26:30
5.5.5.5          1.1.1.1          Installed        00:26:30       00:26:30

Total number of entries are 3
```

```
TB2-VTEP5#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                           VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
101      VNI-101      L2       NW   ----      ----                          ---- ----      5.5.5.5          2.2.2.2
101      VNI-101      L2       NW   ----      ----                          ---- ----      5.5.5.5          4.4.4.4
101      VNI-101      L2       NW   ----      ----                          ---- ----      5.5.5.5          1.1.1.1
101      VNI-101      --       AC   xe48      --- Single Homed Port ---     10   ----      ----             ----
201      VNI-201      L2       NW   ----      ----                          ---- ----      5.5.5.5          2.2.2.2
201      VNI-201      L2       NW   ----      ----                          ---- ----      5.5.5.5          4.4.4.4
201      VNI-201      L2       NW   ----      ----                          ---- ----      5.5.5.5          1.1.1.1

Total number of entries are 7
TB2-VTEP5#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       3c2c.99d6.167a Static Local   ----
201      21.21.21.101     0000.4444.1020 Static Remote  ----
101      11.11.11.51      0000.2222.1010 Static Remote  ----
101      11.11.11.1       3c2c.99d6.167a Static Local   ----
101      11.11.11.201     0000.5555.1010 Static Local   ----
Total number of entries are 6
TB2-VTEP5#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
==========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________________________________
201      21:21::21:1                   3c2c.99d6.167a Static Local     ----
101      11:11::11:1                   3c2c.99d6.167a Static Local     ----
Total number of entries are 2
TB2-VTEP5#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001
 1000         201          irb2001

TB2-VTEP5#show ip route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:27:00
C            21.21.21.0/24 is directly connected, irb2001, 00:26:58
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:32:53

Gateway of last resort is not set
TB2-VTEP5#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C      ::1/128 via ::, lo.L3VRF1, 00:32:53
C      11:11::/48 via ::, irb1001, 00:27:00
C      21:21::/48 via ::, irb2001, 00:26:58
```

```
C      fe80::/64 via ::, irb2001, 00:04:38
TB2-VTEP5#show ip route summary

-------------------------------------------------
IP routing table name is Default-IP-Routing-Table(0)
-------------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 12
Total number of IPv4 paths       : 12
Pending routes (due to route max reached): 0
Route Source    Networks
connected       3
ospf            9
Total           12
FIB             12

ECMP statistics (active in ASIC):
 Total number of IPv4 ECMP routes : 0
 Total number of IPv4 ECMP paths  : 0
TB2-VTEP5#show ipv6 route summary

-------------------------------------------------
IPv6 routing table name is Default-IPv6-Routing-Table(0)
-------------------------------------------------
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 2
Total number of IPv6 paths       : 2
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
Total           2
FIB             2

ECMP statistics (active in ASIC):
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0
TB2-VTEP5#show bgp l2vpn evpn
BGP table version is 11, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[1.1.1.1:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,1.1.1.1]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[201]:[32,1.1.1.1]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
```

```
RD[2.2.2.2:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[101]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[4.4.4.4:11]
*>i  [3]:[101]:[32,4.4.4.4]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[4.4.4.4:21]
*>i  [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [3]:[201]:[32,4.4.4.4]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[5.5.5.5:11] VRF[L2VRF1]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
*>   [2]:[0]:[101]:[48,3c2c:99d6:167a]:[32,11.11.11.1]:[101]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
*>   [2]:[0]:[101]:[48,3c2c:99d6:167a]:[128,11:11::11:1]:[101]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
* i  [3]:[101]:[32,1.1.1.1]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[101]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [3]:[101]:[32,4.4.4.4]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>   [3]:[101]:[32,5.5.5.5]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN

RD[5.5.5.5:21] VRF[L2VRF2]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>   [2]:[0]:[201]:[48,3c2c:99d6:167a]:[32,21.21.21.1]:[201]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
*>   [2]:[0]:[201]:[48,3c2c:99d6:167a]:[128,21:21::21:1]:[201]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
```

```
* i  [3]:[201]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
* i  [3]:[201]:[32,4.4.4.4]
                     4.4.4.4             0        100       0    i  4.4.4.4         VXLAN
*>   [3]:[201]:[32,5.5.5.5]
                     5.5.5.5             0        100       32768  i  ----------      VXLAN

Total number of prefixes 39
TB2-VTEP5#
```

Anycast Gateway

For today's large multi-tenant data center, centralized L3 gateway scheme is very inefficient and sometimes impractical. In order to overcome the drawback of centralized L3GW approach, anycast mode is used.

In Anycast gateway approach, all the VTEPs acts as default gateway for all the VNIDs. We will configure same anycast MAC in all VTEPs.

IRB Configuration for Anycast

Configure from Base Configuration-L2 VXLAN section, then configure below commands for Anycast gateway approach.

VTEP1

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#ip vrf L3VRF1	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
(config-vrf)#rd 11000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from VRF mode
(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast MAC address
(config)# interface irb1001	Configure IRB interface 1001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 11.11.11.1/24 anycast	Configure IP address
(config-if)#ipv6 address 11:11::11:1/48	Configure IPv6 address
(config-if)#evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)#exit	Exit from interface config mode
(config)#interface irb2001	Configure IRB interface 2001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 21.21.21.1/24 anycast	Configure IP address
(config-if)#ipv6 address 21:21::21:1/48	Configure IPv6 address

(config-if)#evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)#exit	Exit from interface config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#router bgp 5000	Enter into BGP router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit from address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb1001	Configure IRB1001 under VXLAN ID 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb2001	Configure IRB2001 under VXLAN ID 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VTEP2

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#ip vrf L3VRF1	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
(config-vrf)#rd 21000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from VRF mode
(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast MAC address
(config)# interface irb1001	Configure IRB interface 1001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 11.11.11.1/24 anycast	Configure IP address
(config-if)#ipv6 address 11:11::11:1/48	Configure IPv6 address
(config-if)#evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address

(config-if)#exit	Exit from interface config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#interface irb2001	Configure IRB interface 2001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 21.21.21.1/24 anycast	Configure IP address
(config-if)#ipv6 address 21:21::21:1/48	Configure IPv6 address
(config-if)#evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)#exit	Exit from interface config mode
(config)#router bgp 5000	Enter into BGP router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit from address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb1001	Configure IRB1001 under VXLAN ID 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb2001	Configure IRB2001 under VXLAN ID 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VTEP4

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#ip vrf L3VRF1	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
(config-vrf)#rd 41000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from VRF mode
(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast MAC address
(config)# interface irb1001	Configure IRB interface 1001

(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 11.11.11.1/24 anycast	Configure IP address
(config-if)#ipv6 address 11:11::11:1/48	Configure IPv6 address
(config-if)#evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)#exit	Exit from interface config mode
(config)#interface irb2001	Configure IRB interface 2001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 21.21.21.1/24 anycast	Configure IP address
(config-if)#ipv6 address 21:21::21:1/48	Configure IPv6 address
(config-if)#evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)#exit	Exit from interface config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#router bgp 5000	Enter into BGP router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit from address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb1001	Configure IRB1001 under VXLAN ID 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb2001	Configure IRB2001 under VXLAN ID 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VTEP5

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#ip vrf L3VRF1	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
(config-vrf)#rd 51000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from VRF mode

(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast MAC address
(config)# interface irb1001	Configure IRB interface 1001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 11.11.11.1/24 anycast	Configure IP address
(config-if)#ipv6 address 11:11::11:1/48	Configure IPv6 address
(config-if)#evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)#exit	Exit from interface config mode
(config)#interface irb2001	Configure IRB interface 2001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 21.21.21.1/24 anycast	Configure IP address
(config-if)#ipv6 address 21:21::21:1/48	Configure IPv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)#exit	Exit from interface config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#router bgp 5000	Enter into BGP router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit from address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb1001	Configure IRB1001 under VXLAN ID 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb2001	Configure IRB2001 under VXLAN ID 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Validations

VTEP1

```
TB2-VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
```

```
1.1.1.1          5.5.5.5          Installed        00:13:05       00:13:05
1.1.1.1          4.4.4.4          Installed        00:18:33       00:18:33
1.1.1.1          2.2.2.2          Installed        00:18:34       00:18:34

Total number of entries are 3
TB2-VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
101      VNI-101      L2       NW   ----      ----                           ---- ----      1.1.1.1          5.5.5.5
101      VNI-101      L2       NW   ----      ----                           ---- ----      1.1.1.1          4.4.4.4
101      VNI-101      L2       NW   ----      ----                           ---- ----      1.1.1.1          2.2.2.2
101      VNI-101      --       AC   po1       00:00:00:00:00:22:22:00:00:00 10   DF        ----             ----
201      VNI-201      L2       NW   ----      ----                           ---- ----      1.1.1.1          5.5.5.5
201      VNI-201      L2       NW   ----      ----                           ---- ----      1.1.1.1          4.4.4.4
201      VNI-201      L2       NW   ----      ----                           ---- ----      1.1.1.1          2.2.2.2
201      VNI-201      --       AC   po1       00:00:00:00:00:22:22:00:00:00 20   DF        ----             ----

Total number of entries are 8
TB2-VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Local   ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.101     0000.4444.1020 Static Remote  ----
101      11.11.11.51      0000.2222.1010 Static Local   ----
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.201     0000.5555.1010 Static Remote  ----
Total number of entries are 6
TB2-VTEP1#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
==========================
VNID     Ip-Addr                                Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________________
201      21:21::21:1                    0000.0000.1111 Static Local     ----
101      11:11::11:1                    0000.0000.1111 Static Local     ----
Total number of entries are 2
TB2-VTEP1#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001
 1000         201          irb2001

TB2-VTEP1#show ip route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:19:26
C            21.21.21.0/24 is directly connected, irb2001, 00:19:26
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:19:28

Gateway of last resort is not set
TB2-VTEP1#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
```

```
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C      ::1/128 via ::, lo.L3VRF1, 00:19:28
C      11:11::/48 via ::, irb1001, 00:19:26
C      21:21::/48 via ::, irb2001, 00:19:26
C      fe80::/64 via ::, irb2001, 00:19:25
TB2-VTEP1#show ip route summary

-------------------------------------------------
IP routing table name is Default-IP-Routing-Table(0)
-------------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 12
Total number of IPv4 paths       : 12
Pending routes (due to route max reached): 0
Route Source    Networks
connected       3
ospf            9
Total           12
FIB             12

ECMP statistics (active in ASIC):
 Total number of IPv4 ECMP routes : 0
 Total number of IPv4 ECMP paths  : 0
TB2-VTEP1#show ipv6 route summary

-------------------------------------------------
IPv6 routing table name is Default-IPv6-Routing-Table(0)
-------------------------------------------------
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 2
Total number of IPv6 paths       : 2
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
Total           2
FIB             2

ECMP statistics (active in ASIC):
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0
TB2-VTEP1#show bgp l2vpn evpn
BGP table version is 6, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[21000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      2.2.2.2              0        100       0    ?  2.2.2.2         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                      2.2.2.2              0        100       0    ?  2.2.2.2         VXLAN

RD[41000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      4.4.4.4              0        100       0    ?  4.4.4.4         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                      4.4.4.4              0        100       0    ?  4.4.4.4         VXLAN

RD[51000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      5.5.5.5              0        100       0    ?  5.5.5.5         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
```

```
                      5.5.5.5              0        100       0    ?  5.5.5.5         VXLAN

RD[1.1.1.1:1] VRF[evpn-gvrf-1]:
*>   [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
*>   [4]:[00:00:00:00:00:22:22:00:00:00]:[32,1.1.1.1]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[1.1.1.1:11] VRF[L2VRF1]:
*>   [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                    4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>                    1.1.1.1              0        100       32768  i  ----------      VXLAN
* i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                    4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>                    1.1.1.1              0        100       32768  i  ----------      VXLAN
*>   [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>   [3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i  [3]:[101]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [3]:[101]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i  [3]:[101]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[1.1.1.1:21] VRF[L2VRF2]:
*>   [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                    4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>                    1.1.1.1              0        100       32768  i  ----------      VXLAN
* i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                    4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>                    1.1.1.1              0        100       32768  i  ----------      VXLAN
*>   [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>   [3]:[201]:[32,1.1.1.1]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [3]:[201]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i  [3]:[201]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[2.2.2.2:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
```

```
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,2.2.2.2]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[101]:[32,2.2.2.2]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[4.4.4.4:11]
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [3]:[101]:[32,4.4.4.4]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[4.4.4.4:21]
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [3]:[201]:[32,4.4.4.4]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[5.5.5.5:11]
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                       5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                       5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                       5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [3]:[101]:[32,5.5.5.5]
                       5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[5.5.5.5:21]
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                       5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                       5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [3]:[201]:[32,5.5.5.5]
                       5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

Total number of prefixes 55
TB2-VTEP1#


TB2-VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
2.2.2.2          4.4.4.4          Installed        00:18:42       00:18:42
2.2.2.2          1.1.1.1          Installed        00:18:43       00:18:43
```

```
2.2.2.2          5.5.5.5          Installed        00:13:14       00:13:14

Total number of entries are 3
TB2-VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
101      VNI-101      L2       NW   ----      ----                            ---- ----      2.2.2.2          4.4.4.4
101      VNI-101      L2       NW   ----      ----                            ---- ----      2.2.2.2          1.1.1.1
101      VNI-101      L2       NW   ----      ----                            ---- ----      2.2.2.2          5.5.5.5
101      VNI-101      --       AC   po1       00:00:00:00:00:22:22:00:00:00 10   NON-DF    ----             ----
201      VNI-201      L2       NW   ----      ----                            ---- ----      2.2.2.2          4.4.4.4
201      VNI-201      L2       NW   ----      ----                            ---- ----      2.2.2.2          1.1.1.1
201      VNI-201      L2       NW   ----      ----                            ---- ----      2.2.2.2          5.5.5.5
201      VNI-201      --       AC   po1       00:00:00:00:00:22:22:00:00:00 20   NON-DF    ----             ----

Total number of entries are 8
TB2-VTEP2#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Local   ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.101     0000.4444.1020 Static Remote  ----
101      11.11.11.51      0000.2222.1010 Static Local   ----
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.201     0000.5555.1010 Static Remote  ----
Total number of entries are 6
TB2-VTEP2#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________________________________
201      21:21::21:1                   0000.0000.1111 Static Local     ----
101      11:11::11:1                   0000.0000.1111 Static Local     ----
Total number of entries are 2
TB2-VTEP2#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001
 1000         201          irb2001

TB2-VTEP2#show ip route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:19:37
C            21.21.21.0/24 is directly connected, irb2001, 00:19:37
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:19:40

Gateway of last resort is not set
TB2-VTEP2#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime
```

```
IP Route Table for VRF "L3VRF1"
C      ::1/128 via ::, lo.L3VRF1, 00:19:40
C      11:11::/48 via ::, irb1001, 00:19:37
C      21:21::/48 via ::, irb2001, 00:19:37
C      fe80::/64 via ::, irb2001, 00:19:36
TB2-VTEP2#show ip route summary

-----------------------------------------------
IP routing table name is Default-IP-Routing-Table(0)
-----------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 12
Total number of IPv4 paths       : 12
Pending routes (due to route max reached): 0
Route Source    Networks
connected       3
ospf            9
Total           12
FIB             12

ECMP statistics (active in ASIC):
 Total number of IPv4 ECMP routes : 0
 Total number of IPv4 ECMP paths  : 0
TB2-VTEP2#show ipv6 route summary

-----------------------------------------------
IPv6 routing table name is Default-IPv6-Routing-Table(0)
-----------------------------------------------
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 2
Total number of IPv6 paths       : 2
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
Total           2
FIB             2

ECMP statistics (active in ASIC):
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0
TB2-VTEP2#show bgp l2vpn evpn
BGP table version is 6, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[11000:11]
*>i[5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                       1.1.1.1             0        100       0    ?  1.1.1.1         VXLAN
*>i[5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                       1.1.1.1             0        100       0    ?  1.1.1.1         VXLAN

RD[41000:11]
*>i[5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                       4.4.4.4             0        100       0    ?  4.4.4.4         VXLAN
*>i[5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                       4.4.4.4             0        100       0    ?  4.4.4.4         VXLAN

RD[51000:11]
*>i[5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                       5.5.5.5             0        100       0    ?  5.5.5.5         VXLAN
*>i[5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                       5.5.5.5             0        100       0    ?  5.5.5.5         VXLAN
```

```
RD[1.1.1.1:1]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[4]:[00:00:00:00:00:22:22:00:00:00]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:11]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:21]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i[3]:[201]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[2.2.2.2:1] VRF[evpn-gvrf-1]:
*> [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[4]:[00:00:00:00:00:22:22:00:00:00]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*> [4]:[00:00:00:00:00:22:22:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       32768  i  ----------      VXLAN

RD[2.2.2.2:11] VRF[L2VRF1]:
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>                   2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i[2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                   1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>                   2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                   1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>                   2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>                   2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i[3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*> [3]:[101]:[32,2.2.2.2]
                      2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[3]:[101]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[3]:[101]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[2.2.2.2:21] VRF[L2VRF2]:
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>                   2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
```

```
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i[2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                  4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                  1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>                   2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                  4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                  1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>                   2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>                   2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[3]:[201]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*> [3]:[201]:[32,2.2.2.2]
                      2.2.2.2              0        100       32768  i  ----------      VXLAN
* i[3]:[201]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[3]:[201]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[4.4.4.4:11]
*>i[2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[3]:[101]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[4.4.4.4:21]
*>i[2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[3]:[201]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[5.5.5.5:11]
*>i[2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i[2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i[2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i[3]:[101]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[5.5.5.5:21]
*>i[2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i[2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i[3]:[201]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

Total number of prefixes 55
TB2-VTEP2#
```

VTEP4

```
TB2-VTEP4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
4.4.4.4          2.2.2.2          Installed        00:18:55       00:18:55
```

```
4.4.4.4          1.1.1.1          Installed        00:18:55       00:18:55
4.4.4.4          5.5.5.5          Installed        00:13:27       00:13:27

Total number of entries are 3
TB2-VTEP4#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
101      VNI-101      L2       NW   ----      ----                           ---- ----      4.4.4.4          2.2.2.2
101      VNI-101      L2       NW   ----      ----                           ---- ----      4.4.4.4          1.1.1.1
101      VNI-101      L2       NW   ----      ----                           ---- ----      4.4.4.4          5.5.5.5
201      VNI-201      L2       NW   ----      ----                           ---- ----      4.4.4.4          2.2.2.2
201      VNI-201      L2       NW   ----      ----                           ---- ----      4.4.4.4          1.1.1.1
201      VNI-201      L2       NW   ----      ----                           ---- ----      4.4.4.4          5.5.5.5
201      VNI-201      --       AC   sa1       --- Single Homed Port ---      20   ----      ----             ----

Total number of entries are 7
TB2-VTEP4#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.101     0000.4444.1020 Static Local   ----
101      11.11.11.51      0000.2222.1010 Static Remote  ----
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.201     0000.5555.1010 Static Remote  ----
Total number of entries are 6
TB2-VTEP4#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________________________________
201      21:21::21:1                   0000.0000.1111 Static Local     ----
101      11:11::11:1                   0000.0000.1111 Static Local     ----
Total number of entries are 2
TB2-VTEP4#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001
 1000         201          irb2001

TB2-VTEP4#show ip route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:19:46
C            21.21.21.0/24 is directly connected, irb2001, 00:19:46
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:19:49

Gateway of last resort is not set
TB2-VTEP4#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime
```

```
IP Route Table for VRF "L3VRF1"
C      ::1/128 via ::, lo.L3VRF1, 00:19:49
C      11:11::/48 via ::, irb1001, 00:19:46
C      21:21::/48 via ::, irb2001, 00:19:46
C      fe80::/64 via ::, irb2001, 00:19:46
TB2-VTEP4#show ip route summary

-----------------------------------------------
IP routing table name is Default-IP-Routing-Table(0)
-----------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 12
Total number of IPv4 paths       : 12
Pending routes (due to route max reached): 0
Route Source    Networks
connected       3
ospf            9
Total           12
FIB             12

ECMP statistics (active in ASIC):
 Total number of IPv4 ECMP routes : 0
 Total number of IPv4 ECMP paths  : 0
TB2-VTEP4#show ipv6 route summary

-----------------------------------------------
IPv6 routing table name is Default-IPv6-Routing-Table(0)
-----------------------------------------------
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 2
Total number of IPv6 paths       : 2
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
Total           2
FIB             2

ECMP statistics (active in ASIC):
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0
TB2-VTEP4#show bgp l2vpn evpn
BGP table version is 4, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[11000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      1.1.1.1             0        100       0    ?  1.1.1.1         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                      1.1.1.1             0        100       0    ?  1.1.1.1         VXLAN

RD[21000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                      2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN

RD[51000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      5.5.5.5             0        100       0    ?  5.5.5.5         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                      5.5.5.5             0        100       0    ?  5.5.5.5         VXLAN
```

```
RD[1.1.1.1:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[201]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[2.2.2.2:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[101]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[4.4.4.4:11] VRF[L2VRF1]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                   1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>                    4.4.4.4              0        100       32768  i  ----------      VXLAN
* i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
```

```
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                   1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>                    4.4.4.4              0        100       32768  i  ----------      VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[101]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [3]:[101]:[32,4.4.4.4]
                      4.4.4.4              0        100       32768  i  ----------      VXLAN
* i  [3]:[101]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[4.4.4.4:21] VRF[L2VRF2]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                   1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>                    4.4.4.4              0        100       32768  i  ----------      VXLAN
* i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i                   1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>                    4.4.4.4              0        100       32768  i  ----------      VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                      4.4.4.4              0        100       32768  i  ----------      VXLAN
* i  [3]:[201]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [3]:[201]:[32,4.4.4.4]
                      4.4.4.4              0        100       32768  i  ----------      VXLAN
* i  [3]:[201]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[5.5.5.5:11]
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [3]:[101]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[5.5.5.5:21]
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [3]:[201]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

Total number of prefixes 57
TB2-VTEP4#
```

VTEP5

```
TB2-VTEP5#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
5.5.5.5          2.2.2.2          Installed        00:13:36       00:13:36
5.5.5.5          4.4.4.4          Installed        00:13:36       00:13:36
5.5.5.5          1.1.1.1          Installed        00:13:36       00:13:36

Total number of entries are 3
TB2-VTEP5#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
101      VNI-101      L2       NW   ----      ----                           ---- ----      5.5.5.5          2.2.2.2
101      VNI-101      L2       NW   ----      ----                           ---- ----      5.5.5.5          4.4.4.4
101      VNI-101      L2       NW   ----      ----                           ---- ----      5.5.5.5          1.1.1.1
101      VNI-101      --       AC   xe48      --- Single Homed Port ---      10   ----      ----             ----
201      VNI-201      L2       NW   ----      ----                           ---- ----      5.5.5.5          2.2.2.2
201      VNI-201      L2       NW   ----      ----                           ---- ----      5.5.5.5          4.4.4.4
201      VNI-201      L2       NW   ----      ----                           ---- ----      5.5.5.5          1.1.1.1

Total number of entries are 7
TB2-VTEP5#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.101     0000.4444.1020 Static Remote  ----
101      11.11.11.51      0000.2222.1010 Static Remote  ----
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.201     0000.5555.1010 Static Local   ----
Total number of entries are 6
TB2-VTEP5#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________________________________
201      21:21::21:1                    0000.0000.1111 Static Local     ----
101      11:11::11:1                    0000.0000.1111 Static Local     ----
Total number of entries are 2
TB2-VTEP5#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
===================================
 1000         101          irb1001
 1000         201          irb2001

TB2-VTEP5#show ip route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:14:07
C            21.21.21.0/24 is directly connected, irb2001, 00:14:05
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:20:00

Gateway of last resort is not set
TB2-VTEP5#show ipv6 route vrf L3VRF1
```

```
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C      ::1/128 via ::, lo.L3VRF1, 00:20:00
C      11:11::/48 via ::, irb1001, 00:14:07
C      21:21::/48 via ::, irb2001, 00:14:05
C      fe80::/64 via ::, irb2001, 00:14:05
TB2-VTEP5#show ip route summary

-------------------------------------------------
IP routing table name is Default-IP-Routing-Table(0)
-------------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 12
Total number of IPv4 paths       : 12
Pending routes (due to route max reached): 0
Route Source    Networks
connected       3
ospf            9
Total           12
FIB             12

ECMP statistics (active in ASIC):
 Total number of IPv4 ECMP routes : 0
 Total number of IPv4 ECMP paths  : 0
TB2-VTEP5#show ipv6 route summary

-------------------------------------------------
IPv6 routing table name is Default-IPv6-Routing-Table(0)
-------------------------------------------------
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 2
Total number of IPv6 paths       : 2
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
Total           2
FIB             2

ECMP statistics (active in ASIC):
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0
TB2-VTEP5#show bgp l2vpn evpn
BGP table version is 7, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[11000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      1.1.1.1             0        100       0    ?  1.1.1.1         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                      1.1.1.1             0        100       0    ?  1.1.1.1         VXLAN

RD[21000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                      2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN
```

```
RD[41000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                     4.4.4.4              0        100       0    ?  4.4.4.4         VXLAN
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                     4.4.4.4              0        100       0    ?  4.4.4.4         VXLAN

RD[1.1.1.1:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,1.1.1.1]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[101]:[32,1.1.1.1]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[201]:[32,1.1.1.1]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[2.2.2.2:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[101]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[4.4.4.4:11]
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [3]:[101]:[32,4.4.4.4]
```

```
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[4.4.4.4:21]
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [3]:[201]:[32,4.4.4.4]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[5.5.5.5:11] VRF[L2VRF1]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[101]:[48,0000:0000:1111]:[32,11.11.11.1]:[101]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
* i                  4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i                  1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>   [2]:[0]:[101]:[48,0000:0000:1111]:[128,11:11::11:1]:[101]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
* i                  4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i                  1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
* i  [3]:[101]:[32,1.1.1.1]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[101]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [3]:[101]:[32,4.4.4.4]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>   [3]:[101]:[32,5.5.5.5]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN

RD[5.5.5.5:21] VRF[L2VRF2]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,0000:0000:1111]:[32,21.21.21.1]:[201]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
* i                  4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i                  1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>   [2]:[0]:[201]:[48,0000:0000:1111]:[128,21:21::21:1]:[201]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
* i                  4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i                  1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                  2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i  [3]:[201]:[32,1.1.1.1]
                     1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [3]:[201]:[32,4.4.4.4]
                     4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>   [3]:[201]:[32,5.5.5.5]
                     5.5.5.5              0        100       32768  i  ----------      VXLAN
```

```
Total number of prefixes 57
TB2-VTEP5#
```

Distributed Gateway

In distributed gateway approach, VTEP will act as default gateways for one or more VNIDs,

Each VTEP having its own default gateway IP and MAC configuration for a given VNID.

IRB Configuration for Distributed

Configure from Base Configuration-L2 VXLAN section, then configure below commands for centralized distributed approach.

VTEP4

`(config)#nvo vxlan irb`	Enable VXLAN irb
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board
`(config)#ip vrf L3VRF1`	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
`(config-vrf)#rd 41000:11`	Assign RD value
`(config-vrf)# route-target both 100:100`	Assign route-target value for same for import and export.
`(config-vrf)# l3vni 1000`	Configure L3VNI as 1000 for L3VRF1
`(config-vrf)#exit`	Exit from VRF mode
`(config)# interface irb2001`	Configure IRB interface 2001
`(config-if)#ip vrf forwarding L3VRF1`	Configure L3VRF1
`(config-if)#ip address 21.21.21.1/24`	Configure IP address
`(config-if)#ipv6 address 21:21::21:1/48`	Configure IPv6 address
`(config-if)#exit`	Exit from interface config mode
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board
`(config)#router bgp 5000`	Enter into BGP router mode
`(config-router)#address-family ipv4 vrf L3VRF1`	Enter into address-family mode for L3VRF1
`(config-router-af)#redistribute connected`	Redistribute connected
`(config-router-af)#exit-address-family`	Exit from address-family
`(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled`	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2`	Assign VRF for EVPN-BGP to carry EVPN route
`(config-nvo)# evpn irb2001`	Configure IRB2001 under VXLAN id 201
`(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

VTEP5

Unconfigure vnid 201 from nvo vxlan.

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)#ip vrf L3VRF1	Create mac routing/forwarding instance with L3VRF1 name and enter into VRF mode
(config-vrf)#rd 51000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from VRF mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)# interface irb1001	Configure IRB interface 1001
(config-if)#ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)#ip address 11.11.11.1/24	Configure IP address
(config-if)#ipv6 address 11:11::11:1/48	Configure IPv6 address
(config-if)#exit	Exit from interface config mode
(config)#router bgp 5000	Enter into BGP router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit from address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)# evpn irb1001	Configure IRB under VXLAN id 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Validations

VTEP4

```
TB2-VTEP4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
4.4.4.4          2.2.2.2          Installed        00:01:17       00:01:17
4.4.4.4          1.1.1.1          Installed        00:01:17       00:01:17
4.4.4.4          5.5.5.5          Installed        00:02:22       00:02:22

Total number of entries are 3
TB2-VTEP4#show nvo vxlan
VXLAN Information
=================
```

```
   Codes: NW - Network Port
          AC - Access Port
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                           VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
201      VNI-201      L2       NW   ----      ----                          ---- ----      4.4.4.4          2.2.2.2
201      VNI-201      L2       NW   ----      ----                          ---- ----      4.4.4.4          1.1.1.1
201      VNI-201      --       AC   sa1       --- Single Homed Port ---     20   ----      ----             ----
1000     ----         L3       NW   ----      ----                          ---- ----      4.4.4.4          5.5.5.5

Total number of entries are 4
TB2-VTEP4#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       3c2c.99c7.077a Static Local   ----
201      21.21.21.101     0000.4444.1020 Static Local   ----
Total number of entries are 3
TB2-VTEP4#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         201          irb2001

TB2-VTEP4#show ip route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
B            5.5.5.5/32 [0/0] is directly connected, tunvxlan2, 00:02:23
B            11.11.11.0/24 [200/0] via 5.5.5.5 (recursive is directly connected, tunvxlan2), 00:01:26

C            21.21.21.0/24 is directly connected, irb2001, 00:01:18
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:02:23

Gateway of last resort is not set
TB2-VTEP4#show ip route summary

-----------------------------------------------
IP routing table name is Default-IP-Routing-Table(0)
-----------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 12
Total number of IPv4 paths       : 12
Pending routes (due to route max reached): 0
Route Source    Networks
connected       3
ospf            9
Total           12
FIB             12

ECMP statistics (active in ASIC):
 Total number of IPv4 ECMP routes : 0
 Total number of IPv4 ECMP paths  : 0
TB2-VTEP4#show bgp l2vpn evpn
BGP table version is 13, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
```

```
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[51000:11]
*>i  [5]:[0]:[1000]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      5.5.5.5              0        100       0    ?  5.5.5.5         VXLAN

RD[1.1.1.1:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[201]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[2.2.2.2:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[101]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[4.4.4.4:11] VRF[L2VRF1]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i  [2]:[0]:[101]:[48,3c2c:99d6:167a]:[32,11.11.11.1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i  [2]:[0]:[101]:[48,3c2c:99d6:167a]:[128,11:11::11:1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
* i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[101]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
```

```
* i  [3]:[101]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

RD[4.4.4.4:21] VRF[L2VRF2]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                      4.4.4.4              0        100       32768  i  ----------      VXLAN
*>   [2]:[0]:[201]:[48,3c2c:99c7:077a]:[32,21.21.21.1]:[201]
                      4.4.4.4              0        100       32768  i  ----------      VXLAN
*>   [2]:[0]:[201]:[48,3c2c:99c7:077a]:[128,21:21::21:1]:[201]
                      4.4.4.4              0        100       32768  i  ----------      VXLAN
* i  [3]:[201]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [3]:[201]:[32,4.4.4.4]
                      4.4.4.4              0        100       32768  i  ----------      VXLAN

RD[5.5.5.5:11]
*>i  [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [2]:[0]:[101]:[48,3c2c:99d6:167a]:[32,11.11.11.1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [2]:[0]:[101]:[48,3c2c:99d6:167a]:[128,11:11::11:1]:[101]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN
*>i  [3]:[101]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i  5.5.5.5         VXLAN

Total number of prefixes 39
TB2-VTEP4#
```

VTEP5

```
TB2-VTEP5#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
5.5.5.5          2.2.2.2          Installed        00:34:13       00:34:13
5.5.5.5          4.4.4.4          Installed        00:01:26       00:01:26
5.5.5.5          1.1.1.1          Installed        00:34:13       00:34:13

Total number of entries are 3
TB2-VTEP5#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status Src-Addr         Dst-Addr
--------------------------------------------------------------------------------------------------------------------
___
101      VNI-101      L2       NW   ----      ----                           ---- ----      5.5.5.5          2.2.2.2
101      VNI-101      L2       NW   ----      ----                           ---- ----      5.5.5.5          1.1.1.1
101      VNI-101      --       AC   xe48      --- Single Homed Port ---      10   ----      ----             ----
1000     ----         L3       NW   ----      ----                           ---- ----      5.5.5.5          4.4.4.4

Total number of entries are 4
TB2-VTEP5#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
--------------------------------------------------------------------------------
101      11.11.11.51      0000.2222.1010 Static Remote  ----
```

```
101      11.11.11.1       3c2c.99d6.167a Static Local   ----
101      11.11.11.201     0000.5555.1010 Static Local   ----
Total number of entries are 3
Total number of entries are 1
TB2-VTEP5#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001

TB2-VTEP5#show ip route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
B            4.4.4.4/32 [0/0] is directly connected, tunvxlan2, 00:01:26
C            11.11.11.0/24 is directly connected, irb1001, 00:34:43
B            21.21.21.0/24 [200/0] via 4.4.4.4 (recursive is directly connected, tunvxlan2), 00:01:26
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:40:36

Gateway of last resort is not set
TB2-VTEP5#show ip route summary

-------------------------------------------------
IP routing table name is Default-IP-Routing-Table(0)
-------------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 12
Total number of IPv4 paths       : 12
Pending routes (due to route max reached): 0
Route Source    Networks
connected       3
ospf            9
Total           12
FIB             12

ECMP statistics (active in ASIC):
 Total number of IPv4 ECMP routes : 0
 Total number of IPv4 ECMP paths  : 0
TB2-VTEP5#show bgp l2vpn evpn
BGP table version is 13, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[41000:11]
*>i  [5]:[0]:[1000]:[24]:[21.21.21.0]:[0.0.0.0]:[1000]
                      4.4.4.4             0        100       0    ?  4.4.4.4         VXLAN

RD[1.1.1.1:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,1.1.1.1]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1             0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
```

```
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[1.1.1.1:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
*>i  [3]:[201]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN

RD[2.2.2.2:1]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [4]:[00:00:00:00:00:22:22:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[101]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:21]
*>i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[4.4.4.4:21]
*>i  [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [2]:[0]:[201]:[48,3c2c:99c7:077a]:[32,21.21.21.1]:[201]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [2]:[0]:[201]:[48,3c2c:99c7:077a]:[128,21:21::21:1]:[201]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i  [3]:[201]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[5.5.5.5:11] VRF[L2VRF1]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[101]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[101]:[48,0000:2222:1010]:[32,11.11.11.51]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[101]:[48,0000:5555:1010]:[32,11.11.11.201]:[101]
                      5.5.5.5              0        100       32768  i  ----------      VXLAN
*>   [2]:[0]:[101]:[48,3c2c:99d6:167a]:[32,11.11.11.1]:[101]
                      5.5.5.5              0        100       32768  i  ----------      VXLAN
*>   [2]:[0]:[101]:[48,3c2c:99d6:167a]:[128,11:11::11:1]:[101]
                      5.5.5.5              0        100       32768  i  ----------      VXLAN
* i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[101]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>   [3]:[101]:[32,5.5.5.5]
                      5.5.5.5              0        100       32768  i  ----------      VXLAN

RD[5.5.5.5:21] VRF[L2VRF2]:
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[201]:[201]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
```

```
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[00:00:00:00:00:22:22:00:00:00]:[201]:[48,0000:2222:1020]:[32,21.21.21.51]:[201]
                       1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i                    2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0000:4444:1020]:[32,21.21.21.101]:[201]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i  [2]:[0]:[201]:[48,3c2c:99c7:077a]:[32,21.21.21.1]:[201]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i  [2]:[0]:[201]:[48,3c2c:99c7:077a]:[128,21:21::21:1]:[201]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i  [3]:[201]:[32,1.1.1.1]
                       1.1.1.1              0        100       0    i  1.1.1.1         VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i  [3]:[201]:[32,4.4.4.4]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

Total number of prefixes 39
TB2-VTEP5#
```

VXLAN IRB ECMP

In multihoming, anycast-IP and the same subnet is configured on the multihomed devices within the same VPN on IRB interfaces connected to the multihomed CE. Both VTEP's will advertise same connected prefix route, remote VTEP need to understand this and treat the traffic destined to multihomed CE as ECMP traffic i.e Routed traffic should loadshare to both the VTEP's.

IRB ECMP Configuration

Configure from Base Configuration-L2 VXLAN section and perform commit after configuration, then configure below commands for ECMP approach.

VTEP1

Configure `max-path ibgp 2` on VTEP1 under BGP IPv4 VRF address family.

`#configure terminal`	Enter Configure mode.
`(config)#router bgp 5000`	Enter into BGP router mode
`(config-router)#address-family ipv4 vrf L3VRF1`	Enter into address-family mode for L3VRF1
`(config-router)#max-paths ibgp 2`	Configure BGP max-path
`(config-router-af)#redistribute connected`	Redistribute connected
`(config-router-af)#exit-address-family`	Exit from address-family
`(config-router-af)#commit`	Commit the transaction

VTEP1 IRB configuration

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#commit	Commit the transaction and save config and reload board
(config)#ip vrf L3VRF1	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
(config-vrf)#rd 11000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from VRF mode
(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast MAC address
(config)#commit	Commit the transaction
(config)# interface irb1001	Configure IRV interface 1001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 11.11.11.1/24 anycast	Configure IP address
(config-if)ipv6 address 11:11::11:1/48	Configure IPv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)exit	Exit from interface config mode
(config)# interface irb 2001	Configure IRB interface 2001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 21.21.21.1/24 anycast	Configure IP address
(config-if)ipv6 address 21:21::21:1/48	Configure IPv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)exit	Exit from interface config mode
(config)#commit	Commit the transaction
(config)router bgp 5000	Enter into BGP router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit form address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1001	Configure IRB1001 under VXLAN ID 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb2001	Configure irb2001 under VXLAN ID 201

`(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.
`(config)#commit`	Commit the transaction

VTEP2

Configure `max-path ibgp 2` on VTEP1 under BGP IPv4 VRF address family.

`#configure terminal`	Enter Configure mode.
`(config)#router bgp 5000`	Enter into BGP router mode
`(config-router)#address-family ipv4 vrf L3VRF1`	Enter into address-family mode for L3VRF1
`(config-router-af)#max-paths ibgp 2`	Configure BGP max-path
`(config-router-af)#redistribute connected`	Redistribute connected
`(config-router-af)#exit-address-family`	Exit from address-family
`(config-router-af)#commit`	Commit the transaction

VTEP2 IRB configuration

(config)#nvo vxlan irb	Enable VXLAN irb
(config)#commit	Commit the transaction and save config and reload board
(config)#ip vrf L3VRF1	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
(config-vrf)#rd 21000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from VRF mode
(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast MAC address
(config)#commit	Commit the transaction
(config)# interface irb 1001	Configure IRB interface 1001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 11.11.11.1/24 anycast	Configure IP address
(config-if)ipv6 address 11:11::11:1/48	Configure IPv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)exit	Exit from interface config mode
(config)# interface irb 2001	Configure IRB interface 2001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 21.21.21.1/24 anycast	Configure IP address
(config-if)ipv6 address 21:21::21:1/48	Configure IPv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast MAC address
(config-if)exit	Exit from interface config mode
(config)#commit	Commit the transaction
(config)router bgp 5000	Enter into BGP router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit form address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign VRF for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1001	Configure irb1001 under VXLAN ID 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign VRF for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb2001	Configure irb2001 under VXLAN id 201

`(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.
`(config)#commit`	Commit the transaction

VTEP5

Unconfigure - `evpn irb-forwarding anycast-gateway-mac` and assign different IP address and IPv6 address to IRB interfaces on VTEP1. Resolve the ARP on Traffic generator and verify the learnt MAC is same as IRB interface MAC not the anycast MAC (0000.0000.1111). Configure BGP max-path under BGP process.

Enable VXLAN Multihhoming on VTEP5 and reboot the node to apply the Multihoming configuration to hardware.

`#configure terminal`	Enter Configure mode.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#commit`	Commit the transaction and save config and reload board
`(config)#nvo vxlan irb`	Enable VXLAN IRB
`(config)#commit`	Commit the transaction
`(config)#ip vrf L3VRF1`	Create MAC routing/forwarding instance with L3VRF1 name and enter into VRF mode
`(config-vrf)#rd 51000:11`	Assign RD value
`(config-vrf)# route-target both 100:100`	Assign route-target value for same for import and export.
`(config-vrf)# l3vni 1000`	Configure L3VNI as 1000 for L3VRF1
`(config-vrf)#commit`	Commit the transaction
`(config)# no evpn irb-forwarding anycast-gateway-mac`	Delete `evpn irb-forwarding anycast-gateway-mac` address
`(config)#commit`	Commit the transaction
`(config)# interface irb1001`	Configure IRB interface 1001
`(config-irb-if)#ip vrf forwarding L3VRF1`	Configure L3VRF1
`(config-irb-if)#ip address 101.11.11.1/24`	Configure IP address
`(config-irb-if)#ipv6 address 101:11::11:1/48`	Configure IPv6 address
`(config-irb-if)#commit`	Commit the transaction
`(config)#router bgp 5000`	Enter into BGP router mode
`(config-router)#address-family ipv4 vrf L3VRF1`	Enter into address-family mode for L3VRF1
`(config-router-af)# max-paths ibgp 2`	Configure BGP max-path .
`(config-router-af)#redistribute connected`	Redistribute connected
`(config-router-af)#exit-address-family`	Exit from address-family
`(config-router-af)#commit`	Commit the transaction
`(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled`	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1`	Assign VRF for EVPN-BGP to carry EVPN route
`(config-nvo)# evpn irb1001`	Configure IRB under VXLAN ID 101

(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config-nvo)#commit	Commit the transaction

Validations

On VTEP5, verify that in the VRF routing table , ECMP path for the IRB address (11.11.11.1) is via VTEP1 - 1.1.1.1 and VTEP2 -2.2.2.2 . Send the Traffic from VTEP5 Single homed to Multihomed. Traffic should be forwarded via VTEP1 and VTEP2 and is loadshared between the Multihome VTEPs.

VTEP5

```
TB2-VTEP5#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status            Up/Down         Update
================================================================================
5.5.5.5           2.2.2.2          Installed         00:34:13        00:34:13
5.5.5.5           4.4.4.4          Installed         00:01:26        00:01:26
5.5.5.5           1.1.1.1          Installed         00:34:13        00:34:13

Total number of entries are 3


TB2-VTEP5#show nvo vxlan
VXLAN Information
================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________
___
101      VNI-101      L2       NW   ----      ----                         ---- ----      5.5.5.5          2.2.2.2
101      VNI-101      L2       NW   ----      ----                         ---- ----      5.5.5.5          1.1.1.1
101      VNI-101      --       AC   xe48      --- Single Homed Port ---    10   ----      ----             ----
1000     ----         L3       NW   ----      ----                         ---- ----      5.5.5.5          4.4.4.4

Total number of entries are 4

TB2-VTEP5#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
101      11.11.11.1       0000.0000.1111 Static       Remote      ----
101      101.11.11.1      3c2c.99d6.168a Static       Local       ----
101      11.11.11.201     0000.5555.1010 Static       Local       ----
Total number of entries are 3
Total number of entries are 1

TB2-VTEP5#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001

TB2-VTEP5#show ip route vrf L3VRF1
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C           101.11.11.0/24 is directly connected, irb1001, 00:34:43
B           11.11.11.0/24 [200/0] via 1.1.1.1 (recursive is directly connected, tunvxlan2), 00:01:26
                          [200/0] via 2.2.2.2 (recursive is directly connected, tunvxlan2), 00:01:26
C           127.0.0.0/8 is directly connected, lo.L3VRF1, 00:40:36

Gateway of last resort is not set
```

Send 10000 pps from VTEP5 (Traffic generator- SH5) and verify the counters on VTEP5, VTEP1, VTEP2 and Switch

```
TB2-VTEP5#show interface counter rate mbps
+------------------+-------------+------------+-------------+------------+
|     Interface    |    Rx mbps  |    Rx pps  |    Tx mbps  |    Tx pps  |
+------------------+-------------+------------+-------------+------------+
 xe48                100           10000          0.01          8
 xe40                0.00          0              106.76        10000
```

On VTEP1 and VTEP2, verify that traffic is load-balanced on ECMP path from VTEP5.

VTEP1

```
TB2-VTEP1#show interface counter rate mbps
+------------------+-------------+------------+-------------+------------+
|     Interface    |    Rx mbps  |    Rx pps  |    Tx mbps  |    Tx pps  |
+------------------+-------------+------------+-------------+------------+
 po2                 62.75         5000           0.01          8
 po1                 0.00          0              62.98         5000
 xe25                31.98         2500           0             0
 xe26                30.95         2501           0             0
 xe2                 0.00          0              31.53         2500
 xe3                 0.00          0              30.53         2500
```

VTEP2

```
TB2-VTEP2#show interface counter rate mbps
+------------------+-------------+------------+-------------+------------+
|     Interface    |    Rx mbps  |    Rx pps  |    Tx mbps  |    Tx pps  |
+------------------+-------------+------------+-------------+------------+
 po3                 62.75         5000           0.01          8
 po1                 0.00          0              62.98         5000
 xe27                31.98         2500           0             0
 xe28                30.95         2501           0             0
 xe8                 0.00          0              31.53         2500
 xe9                 0.00          0              30.53         2500
```

Verify the Traffic on Multihomed Switch:

SW1 (Multihomed)

```
TB2-SW1#show interface counter rate mbps
+------------------+-------------+------------+-------------+------------+
```

Interface	Rx mbps	Rx pps	Tx mbps	Tx pps
po1	100	10000	0.01	8
xe7	0.00	0	100	10000
xe8	25.01	2501	0	0
xe9	24.99	2499	0	0
xe2	24.98	2499	0	0
xe3	25.02	2501	0	0

CHAPTER 2 EVPN IRB - Anycast Support for Multiple Subnets

Overview

An EVPN-based Integrated Routing and Bridging (IRB) solution enables communication between two Layer-2 Virtual Network Identifiers (VNIDs) using IP-based Virtual Routing and Forwarding (IP-VRF). This enhancement provides Anycast Gateway Routing support for multiple subnets under the IRB interface per VNID. It allows for efficient Layer-3 termination at the VTEP for traffic originating from various subnets within a broadcast domain.

Feature Characteristics

- Connects primary or secondary subnets with either router MAC or anycast MAC address.
- Supports Anycast Gateway for multiple subnets under the layer-2 VNID's.
- Subnets A, B, and C can have Anycast Gateway support, while subnet D is reserved for BGP.
- Supports the Interfacefull model for ARP/ND requests and the interfaceless model using the kernel interface with a unique MAC per interface (Router MAC or Anycast MAC) for all subnets.

ARP/ND replies come from the ARP/ND cache table for host requests. Hosts can send Layer-3 packets with either Anycast MAC or Router MAC, and Layer-3 termination in the VTEP happens for both. By default, each subnet uses the Router MAC in ARP/ND cache. When the anycast argument is configured, it updates the ARP cache with the Anycast MAC. BGP withdraws the Route-Type 2 and update with Anycast/Router MAC to inform the configured gateway for the layer-2 VNID to non-default gateway nodes.

Benefits

EVPN IRB Anycast Support for Multiple Subnets offers benefits in networking environments by providing efficient traffic control, enhanced security, scalability, and improved performance:

- Distributed Default Gateway: Enables consistent Anycast gateway across all VTEPs.
- Operational Flexibility: Supports multiple subnets under a single VNID without additional interfaces.
- Traffic Optimization: Ensures local routing at ingress VTEP.
- Scalability: Supports multiple tenant VRFs and VNIs.
- Multihoming Support: Integrates with EVPN ESI multihoming for redundancy.
- Control-Plane Synchronization: Automatically updates MAC-IP advertisement upon gateway mode changes.

Topology

The topology illustrates Anycast Gateway Routing across multiple primary and secondary subnets (such as subnets A, B, and C) under the IRB interface for each VNID. This allows hosts to send Layer 3 packets using either an Anycast MAC or a Router MAC, with termination happening at the local VTEP to minimize traffic hairpining.

- VTEPs (VTEP1, VTEP2, VTEP4, VTEP5): These nodes function as the VXLAN Tunnel End Points where Layer-3 termination occurs for both Router MAC and Anycast MAC addresses.
- Multi-homed Group (MH 2): VTEP1 and VTEP2 form a redundant connection to Switch1 (SW1) to provide high availability for the connected tenant systems.

- Switch1 (SW1): Connects the multihomed Tenant Systems (TS1 and TS2) to VTEP1 and VTEP2 via a Port Channel (po1). It handles the VLAN trunking required to carry traffic for subnets 10 and 20.
- Switch2 (SW2): Provides connectivity for Tenant System TS2-21 to VTEP4. It participates in OSPF to ensure underlay reachability for the traffic it forwards.
- P3: Acts as the central underlay connectivity point where all VTEP interfaces (such as xe25, xe27, xe7, and xe40) are interconnected.

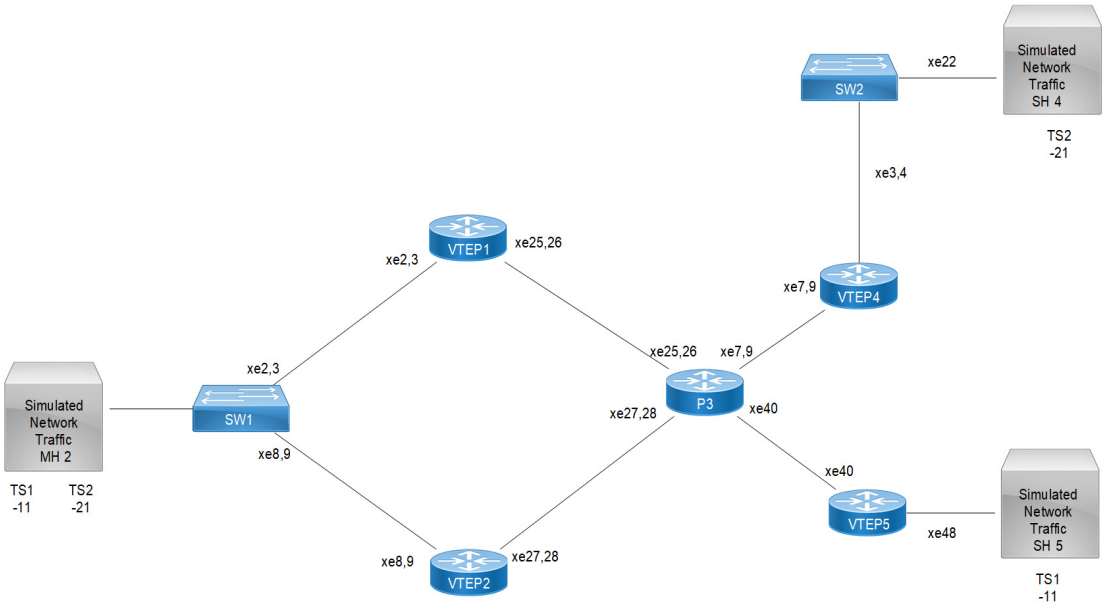

Figure 2-12: EVPN IRB for Multiple Subnets

Configuration

Configure the nodes within the topology to set up EVPN IRB with Anycast support for multiple subnets.

Prerequisites

Ensure the following base configurations are active before configuring Anycast support for multiple subnets:

- Establish Underlay or Overlay Connectivity: Configure OSPF and BGP (L2VPN EVPN address family) to establish reachability between VTEPs.
- Configure EVPN Multihoming: Enable multihoming and configure the ESI hold time. For configurations, refer the VXLAN Multi-homing Configuration section.
- Define VRFs: Create L2 and L3 VRFs with appropriate Route Distinguishers (RD) and Route Targets (RT).

Perform the following steps to configure the L3 VRF and the IRB interface for multi-subnet support:

1. Define the virtual MAC address that all VTEPs will share to act as a single logical default gateway for the subnets.
   ```
   (config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
   ```
2. Configure the IRB interface, associate it with an L3VRF, and enable anycast gateway functionality.
   ```
   (config)# interface irb1001
   (config-if)# ip vrf forwarding L3VRF1
   (config-if)# evpn irb-if-forwarding anycast-gateway-mac
   ```
3. Configure primary and secondary IPv4 and IPv6 addresses on the IRB interface using the anycast keyword..
   ```
   (config-irb-if)# ip address 11.11.11.1/24 anycast
   (config-irb-if)# ip address 12.11.11.1/24 secondary anycast
   (config-irb-if)# ip address 15.11.11.1/16 secondary anycast
   ```

```
(config-irb-if)# ipv6 address 11:11::11:1/48 anycast
(config-irb-if)# ipv6 address 13:11::11:1/48 anycast
```

4. Associate the specific VXLAN Network Identifier (VNID) with the IRB interface and define the reachability protocol.

```
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled
(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1
(config-nvo)# evpn irb1001
```

Note: Perform these configuration steps on all participating VTEPs within EVPN to ensure consistent Anycast Gateway and seamless host mobility. Ensure that the Anycast Gateway MAC address remains identical across all nodes.

Running Configurations

VTEP1

```
hostname VTEP1
!
# Prerequisites (Global Enablement & Multihoming)
nvo vxlan enable
nvo vxlan irb
evpn esi hold-time 90
evpn vxlan multihoming enable
evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
nvo vxlan vtep-ip-global 1.1.1.1
!
ip vrf L3VRF1
 rd 1.1.1.1:11
 route-target both 5000:11
 l3vni 1000
!
mac vrf L2VRF1
 description MAC VRF RED
 rd 1.1.1.1:1
 route-target both 5000:1
!
mac vrf L2VRF2
 description MAC VRF BLUE
 rd 1.1.1.1:2
 route-target both 5000:2
!
interface po1
 switchport
 evpn multi-homed system-mac 0000.0000.1212
!
interface xe2
 channel-group 1 mode active
!
interface xe25
 ip address 100.11.11.1/24
!
interface irb1001
 ip vrf forwarding L3VRF1
```

```
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 11.11.11.1/24 anycast
 # Secondary IPv4 Anycast Subnets
 ip address 12.11.11.1/24 secondary anycast
 ip address 15.11.11.1/16 secondary anycast
 ip address 16.11.11.1/8 secondary anycast
 ip address 17.11.11.1/24 secondary anycast
 ip address 18.11.11.1/16 secondary anycast
 ip address 20.11.11.1/24 secondary anycast
 ip address 21.11.11.1/8 secondary anycast
 ip address 22.11.11.1/24 secondary anycast
 ip address 25.11.11.1/16 secondary anycast
 # Non-Anycast Subnets (Router MAC)
 ip address 13.11.11.1/16 secondary
 ip address 14.11.11.1/24 secondary
 ip address 19.11.11.1/8 secondary
 # IPv6 Anycast Subnets
 ipv6 address 11:11::11:1/48 anycast
 ipv6 address 13:11::11:1/48 anycast
 ipv6 address 14:11::11:1/50 anycast
 ipv6 address 12:11::11:1/64
!
router ospf 100
 ospf router-id 1.1.1.1
 network 1.1.1.1/32 area 0.0.0.0
 network 100.11.11.0/24 area 0.0.0.0
 bfd all-interfaces
!
router bgp 5000
 bgp router-id 1.1.1.1
 neighbor 2.2.2.2 remote-as 5000
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 4.4.4.4 remote-as 5000
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 neighbor 5.5.5.5 remote-as 5000
 neighbor 5.5.5.5 update-source lo
 neighbor 5.5.5.5 advertisement-interval 0
 address-family l2vpn evpn
  neighbor 2.2.2.2 activate
  neighbor 4.4.4.4 activate
  neighbor 5.5.5.5 activate
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
```

```
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan po1 10
 map vnid 101
 mac 0000.2222.1010 ip 11.11.11.51
!
nvo vxlan access-if port-vlan po1 20
 map vnid 201
 mac 0000.2222.1020 ip 21.21.21.51
```

VTEP2

```
hostname VTEP2
!
nvo vxlan enable
nvo vxlan irb
evpn esi hold-time 90
evpn vxlan multihoming enable
evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
nvo vxlan vtep-ip-global 2.2.2.2
!
mac vrf L2VRF1
 description MAC VRF RED
 rd 2.2.2.2:1
 route-target both 5000:1
!
mac vrf L2VRF2
 description MAC VRF BLUE
 rd 2.2.2.2:2
 route-target both 5000:2
!
ip vrf L3VRF1
 rd 2.2.2.2:11
 route-target both 5000:11
 l3vni 1000
!
interface po1
 switchport
 evpn multi-homed system-mac 0000.0000.1212
!
interface xe8
 channel-group 1 mode active
!
interface xe27
 ip address 100.22.22.1/24
!
interface irb1001
 ip vrf forwarding L3VRF1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 11.11.11.1/24 anycast
```

```
 ip address 12.11.11.1/24 secondary anycast
 ip address 13.11.11.1/16 secondary
 ip address 14.11.11.1/24 secondary
 ip address 15.11.11.1/16 secondary anycast
 ip address 16.11.11.1/8 secondary anycast
 ip address 17.11.11.1/24 secondary anycast
 ip address 18.11.11.1/16 secondary anycast
 ip address 19.11.11.1/8 secondary
 ip address 20.11.11.1/24 secondary anycast
 ip address 21.11.11.1/8 secondary anycast
 ip address 22.11.11.1/24 secondary anycast
 ip address 25.11.11.1/16 secondary anycast
 ipv6 address 11:11::11:1/48 anycast
 ipv6 address 12:11::11:1/64
 ipv6 address 13:11::11:1/48 anycast
 ipv6 address 14:11::11:1/50 anycast
!
router ospf 100
 ospf router-id 2.2.2.2
 network 2.2.2.2/32 area 0.0.0.0
 network 100.22.22.0/24 area 0.0.0.0
 bfd all-interfaces
!
router bgp 5000
 bgp router-id 2.2.2.2
 neighbor 1.1.1.1 remote-as 5000
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 4.4.4.4 remote-as 5000
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 neighbor 5.5.5.5 remote-as 5000
 neighbor 5.5.5.5 update-source lo
 neighbor 5.5.5.5 advertisement-interval 0
 address-family l2vpn evpn
  neighbor 1.1.1.1 activate
  neighbor 4.4.4.4 activate
  neighbor 5.5.5.5 activate
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan po1 10
```

```
 map vnid 101
 mac 0000.2222.1010 ip 11.11.11.51
!
nvo vxlan access-if port-vlan po1 20
 map vnid 201
 mac 0000.2222.1020 ip 21.21.21.51
```

VTEP4

```
hostname VTEP4
!
nvo vxlan enable
nvo vxlan irb
evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
nvo vxlan vtep-ip-global 4.4.4.4
!
mac vrf L2VRF1
 description MAC VRF RED
 rd 4.4.4.4:1
 route-target both 5000:1
!
mac vrf L2VRF2
 description MAC VRF BLUE
 rd 4.4.4.4:2
 route-target both 5000:2
!
ip vrf L3VRF1
 rd 4.4.4.4:11
 route-target both 5000:11
 l3vni 1000
!
interface po1
 switchport
!
interface xe3
 channel-group 1 mode active
!
interface xe7
 ip address 100.44.44.1/24
!
interface irb1001
 ip vrf forwarding L3VRF1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 11.11.11.1/24 anycast
 ip address 12.11.11.1/24 secondary anycast
 ip address 13.11.11.1/16 secondary
 ip address 14.11.11.1/24 secondary
 ip address 15.11.11.1/16 secondary anycast
 ip address 16.11.11.1/8 secondary anycast
 ip address 17.11.11.1/24 secondary anycast
 ip address 18.11.11.1/16 secondary anycast
 ip address 19.11.11.1/8 secondary
```

```
 ip address 20.11.11.1/24 secondary anycast
 ip address 21.11.11.1/8 secondary anycast
 ip address 22.11.11.1/24 secondary anycast
 ip address 25.11.11.1/16 secondary anycast
 ipv6 address 11:11::11:1/48 anycast
 ipv6 address 12:11::11:1/64
 ipv6 address 13:11::11:1/48 anycast
 ipv6 address 14:11::11:1/50 anycast
!
router ospf 100
 ospf router-id 4.4.4.4
 network 4.4.4.4/32 area 0.0.0.0
 network 100.44.44.0/24 area 0.0.0.0
 bfd all-interfaces
!
router bgp 5000
 bgp router-id 4.4.4.4
 neighbor 2.2.2.2 remote-as 5000
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 1.1.1.1 remote-as 5000
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 5.5.5.5 remote-as 5000
 neighbor 5.5.5.5 update-source lo
 neighbor 5.5.5.5 advertisement-interval 0
 address-family l2vpn evpn
  neighbor 2.2.2.2 activate
  neighbor 1.1.1.1 activate
  neighbor 5.5.5.5 activate
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan po1 10
 map vnid 101
 mac 0000.2222.1010 ip 11.11.11.51
!
nvo vxlan access-if port-vlan po1 20
 map vnid 201
 mac 0000.2222.1020 ip 21.21.21.51
```

VTEP4

```
hostname VTEP3
```

```
!
nvo vxlan enable
nvo vxlan irb
evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
nvo vxlan vtep-ip-global 4.4.4.4
!
mac vrf L2VRF1
 description MAC VRF RED
 rd 4.4.4.4:1
 route-target both 5000:1
!
mac vrf L2VRF2
 description MAC VRF BLUE
 rd 4.4.4.4:2
 route-target both 5000:2
!
ip vrf L3VRF1
 rd 4.4.4.4:11
 route-target both 5000:11
 l3vni 1000
!
interface po1
 switchport
!
interface xe3
 channel-group 1 mode active
!
interface xe7
 ip address 100.44.44.1/24
!
interface irb1001
 ip vrf forwarding L3VRF1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 11.11.11.1/24 anycast
 ip address 12.11.11.1/24 secondary anycast
 ip address 13.11.11.1/16 secondary
 ip address 14.11.11.1/24 secondary
 ip address 15.11.11.1/16 secondary anycast
 ip address 16.11.11.1/8 secondary anycast
 ip address 17.11.11.1/24 secondary anycast
 ip address 18.11.11.1/16 secondary anycast
 ip address 19.11.11.1/8 secondary
 ip address 20.11.11.1/24 secondary anycast
 ip address 21.11.11.1/8 secondary anycast
 ip address 22.11.11.1/24 secondary anycast
 ip address 25.11.11.1/16 secondary anycast
 ipv6 address 11:11::11:1/48 anycast
 ipv6 address 12:11::11:1/64
 ipv6 address 13:11::11:1/48 anycast
 ipv6 address 14:11::11:1/50 anycast
```

```
!
router ospf 100
 ospf router-id 4.4.4.4
 network 4.4.4.4/32 area 0.0.0.0
 network 100.44.44.0/24 area 0.0.0.0
 bfd all-interfaces
!
router bgp 5000
 bgp router-id 4.4.4.4
 neighbor 2.2.2.2 remote-as 5000
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 1.1.1.1 remote-as 5000
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 5.5.5.5 remote-as 5000
 neighbor 5.5.5.5 update-source lo
 neighbor 5.5.5.5 advertisement-interval 0
 address-family l2vpn evpn
  neighbor 2.2.2.2 activate
  neighbor 1.1.1.1 activate
  neighbor 5.5.5.5 activate
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan po1 10
 map vnid 101
 mac 0000.2222.1010 ip 11.11.11.51
!
nvo vxlan access-if port-vlan po1 20
 map vnid 201
 mac 0000.2222.1020 ip 21.21.21.51
!
```

VTEP5

```
hostname VTEP5
!
nvo vxlan enable
nvo vxlan irb
evpn irb-forwarding anycast-gateway-mac 0000.0000.abcd
nvo vxlan vtep-ip-global 5.5.5.5
!
mac vrf L2VRF1
```

```
 description MAC VRF RED
 rd 5.5.5.5:1
 route-target both 5000:1
!
mac vrf L2VRF2
 description MAC VRF BLUE
 rd 5.5.5.5:2
 route-target both 5000:2
!
ip vrf L3VRF1
 rd 5.5.5.5:11
 route-target both 5.5.5.5:11
 l3vni 1000
!
interface xe48
 switchport
!
interface xe40
 ip address 100.55.55.1/24
!
interface irb1001
 ip vrf forwarding L3VRF1
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 11.11.11.1/24 anycast
 ip address 12.11.11.1/24 secondary anycast
 ip address 13.11.11.1/16 secondary
 ip address 14.11.11.1/24 secondary
 ip address 15.11.11.1/16 secondary anycast
 ip address 16.11.11.1/8 secondary anycast
 ip address 17.11.11.1/24 secondary anycast
 ip address 18.11.11.1/16 secondary anycast
 ip address 19.11.11.1/8 secondary
 ip address 20.11.11.1/24 secondary anycast
 ip address 21.11.11.1/8 secondary anycast
 ip address 22.11.11.1/24 secondary anycast
 ip address 25.11.11.1/16 secondary anycast
 ipv6 address 11:11::11:1/48 anycast
 ipv6 address 12:11::11:1/64
 ipv6 address 13:11::11:1/48 anycast
 ipv6 address 14:11::11:1/50 anycast
!
router ospf 100
 ospf router-id 5.5.5.5
 network 5.5.5.5/32 area 0.0.0.0
 network 100.55.55.0/24 area 0.0.0.0
 bfd all-interfaces
!
router bgp 5000
 bgp router-id 5.5.5.5
 neighbor 2.2.2.2 remote-as 5000
```

```
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 4.4.4.4 remote-as 5000
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 neighbor 1.1.1.1 remote-as 5000
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 address-family l2vpn evpn
  neighbor 2.2.2.2 activate
  neighbor 4.4.4.4 activate
  neighbor 1.1.1.1 activate
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan xe48 10
 map vnid 101
 mac 0000.2222.1010 ip 11.11.11.51
!
nvo vxlan access-if port-vlan xe48 20
 map vnid 201
 mac 0000.2222.1020 ip 21.21.21.51
```

Switch1

```
interface po1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan add 2,10,20
!
interface xe2
 channel-group 1 mode active
!
interface xe8
 channel-group 1 mode active
!
interface vlan1.20
 ip address 21.21.21.2/24
 ipv6 address 21:21::21:2/48
!
interface vlan1.2
 ip address 100.12.12.2/24
!
```

```
interface xe1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan add 2,10,20
```

Switch2

```
router ospf 100
 ospf router-id 6.6.6.6
 network 6.6.6.6/32 area 0.0.0.0
 network 100.14.14.0/24 area 0.0.0.0
 bfd all-interfaces
```

Validation

VTEP1:

```
VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.51      0000.2222.1010 Static Local   ----
101      11.11.11.201     0000.4444.1010 Static Remote  ----
101      12.11.11.1       0000.0000.1111 Static Local   ----
101      13.11.11.1       1444.8f10.d8c9 Static Local   ----
101      14.11.11.1       1444.8f10.d8c9 Static Local   ----
101      15.11.11.1       0000.0000.1111 Static Local   ----
101      16.11.11.1       0000.0000.1111 Static Local   ----
101      17.11.11.1       0000.0000.1111 Static Local   ----
101      18.11.11.1       0000.0000.1111 Static Local   ----
101      19.11.11.1       1444.8f10.d8c9 Static Local   ----
101      20.11.11.1       0000.0000.1111 Static Local   ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.51      0000.2222.1020 Static Local   ----
201      22.22.22.1       1444.8f10.d8c9 Static Local   ----
201      23.23.23.1       0000.0000.1111 Static Local   ----
201      24.24.24.1       0000.0000.1111 Static Local   ----
201      25.25.25.1       1444.8f10.d8c9 Static Local   ----
201      26.26.26.1       0000.0000.1111 Static Local   ----
201      27.27.27.1       1444.8f10.d8c9 Static Local   ----
201      28.28.28.1       0000.0000.1111 Static Local   ----
201      29.29.29.1       1444.8f10.d8c9 Static Local   ----
201      30.30.30.1       0000.0000.1111 Static Local   ----
201      192.85.1.2       0039.4400.0020 Dynamic Local  ----
Total number of entries are 24
------------------------------------------------------------
VTEP1#sh nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age
-Out    Retries-Left
_____________________________________________________________________________
_____________________
101      11:11::11:1                   0000.0000.1111 Static Local     ----
101      12:11::11:1                   1444.8f10.d8c9 Static Local     ----
101      13:11::11:1                   0000.0000.1111 Static Local     ----
101      14:11::11:1                   0000.0000.1111 Static Local     ----
101      15:11::11:1                   0000.0000.1111 Static Local     ----
101      16:11::11:1                   1444.8f10.d8c9 Static Local     ----
101      17:11::11:1                   0000.0000.1111 Static Local     ----
101      18:11::11:1                   1444.8f10.d8c9 Static Local     ----
101      19:11::11:1                   1444.8f10.d8c9 Static Local     ----
101      20:20::20:1                   1444.8f10.d8c9 Static Local     ----
201      21:21::21:1                   1444.8f10.d8c9 Static Local     ----
```

```
201      22:22::22:1                    0000.0000.1111 Static Local     ----
201      23:23::23:1                    1444.8f10.d8c9 Static Local     ----
201      24:24::24:1                    1444.8f10.d8c9 Static Local     ----
201      25:25::25:1                    0000.0000.1111 Static Local     ----
201      26:26::26:1                    0000.0000.1111 Static Local     ----
201      27:27::27:1                    1444.8f10.d8c9 Static Local     ----
201      28:28::28:1                    0000.0000.1111 Static Local     ----
201      29:29::29:1                    1444.8f10.d8c9 Static Local     ----
201      30:30::30:1                    0000.0000.1111 Static Local     ----
Total number of entries are 20
Leaf1#
```

VTEP2:

```
VTEP2#show int counters rate mbps
+------------------+-------------+-------------+--------------++
|    Interface|Rx mbps| Rx pps|   Tx mbps|  Tx pps|
+------------------+-------------+-------------+--------------++
ce49 1099.237719330.001
ce50 0.00 0    0.00 0
ce51 0.00 0    0.00 0
ce52 0.00 0    0.00 0
ce53 0.00 0    0.00 0
ce54 0.00 0    0.00 0
ce55 0.00 0    0.00 0
ce56 0.00 0    0.00 0
po3  1099.237719340.001
xe3  0.00 0    0.00 0
xe5  0.00 0    0.00 0



VTEP2#show nvo vxlan arp-cache VXLAN ARP-CACHE Information
============================
VNID Ip-AddrMac-AddrTypeAge-OutRetries-Left

101  11.11.11.10000.0000.1111StaticLocal----
101  11.11.11.510000.2222.1010StaticRemote----
101  11.11.11.2010000.4444.1010StaticRemote----
101  12.11.11.10000.0000.1111StaticLocal----
101  13.11.11.1b86a.979c.1669StaticLocal----
101  14.11.11.1b86a.979c.1669StaticLocal----
101  15.11.11.10000.0000.1111StaticLocal----
101  16.11.11.10000.0000.1111StaticLocal----
101  17.11.11.10000.0000.1111StaticLocal----
101  18.11.11.10000.0000.1111StaticLocal----
101  19.11.11.1b86a.979c.1669StaticLocal----
101  20.11.11.10000.0000.1111StaticLocal----
201  21.21.21.10000.0000.1111StaticLocal----
201  21.21.21.510000.2222.1020StaticRemote----
201  22.22.22.1b86a.979c.1669StaticLocal----
201  23.23.23.10000.0000.1111StaticLocal----
201  24.24.24.10000.0000.1111StaticLocal----
201  25.25.25.1b86a.979c.1669StaticLocal----
201  26.26.26.10000.0000.1111StaticLocal----
201  27.27.27.1b86a.979c.1669StaticLocal----
201  28.28.28.10000.0000.1111StaticLocal----
201  29.29.29.1b86a.979c.1669StaticLocal----
201  30.30.30.10000.0000.1111StaticLocal----
201  192.85.1.20039.4400.0020Dynamic Remote----
Total number of entries are 24

VTEP2# show nvo vxlan nd-cache VXLAN ND-CACHE Information
============================
VNID Ip-AddrMac-AddrTypeAge
-Out Retries-Left


101  11:11::11:10000.0000.1111StaticLocal----
101  12:11::11:1b86a.979c.1669StaticLocal----
101  13:11::11:10000.0000.1111StaticLocal----
101  14:11::11:10000.0000.1111StaticLocal----
```

```
101  15:11::11:10000.0000.1111StaticLocal----
101  16:11::11:1b86a.979c.1669StaticLocal----
101  17:11::11:10000.0000.1111StaticLocal----
101  18:11::11:1b86a.979c.1669StaticLocal----
101  19:11::11:1b86a.979c.1669StaticLocal----
101  20:20::20:1b86a.979c.1669StaticLocal----
201  21:21::21:1b86a.979c.1669StaticLocal----
201  22:22::22:10000.0000.1111StaticLocal----
201  23:23::23:1b86a.979c.1669StaticLocal----
201  24:24::24:1b86a.979c.1669StaticLocal----
201  25:25::25:10000.0000.1111StaticLocal----
201  26:26::26:10000.0000.1111StaticLocal----
201  27:27::27:1b86a.979c.1669StaticLocal----
201  28:28::28:10000.0000.1111StaticLocal----
201  29:29::29:1b86a.979c.1669StaticLocal----
201  30:30::30:10000.0000.1111StaticLocal----
Total number of entries are 20
Leaf4#
```

Implementation Examples

Here is an example scenario and a solution for implementing EVPN IRB Anycast Support for Multiple Subnets:

Multi-Tenant Data Center Monitoring

Scenario:

A cloud service provider (CSP) hosts thousands of virtual machines and containerized workloads for enterprises. Maintaining high-performance networking is critical to prevent latency-sensitive applications from failing. The network operations team requires a way to support multiple subnets per tenant without wasting IP addresses on unique per-switch gateway IPs.

Implementation:

Enable EVPN IRB and configure a global anycast MAC on all leaf switches.

Configure IRB interfaces with multiple secondary subnets using the anycast keyword.

This allows virtual machines in any subnet to use the same gateway IP regardless of which leaf switch they are connected to, facilitating seamless VM mobility.

Switch1 (SW1) Configuration

SW1 provides multihomed connectivity to VTEP1 and VTEP2 using a Port Channel (po1).

1. Configure VLAN Bridge and Trunking: Enable the VLAN bridge and configure the port channel as a trunk port to allow specific VLANs.

```
(config)# interface po1
(config-if)# switchport
(config-if)# bridge-group 1
(config-if)# switchport mode trunk
(config-if)# switchport trunk allowed vlan add 2,10,20
```

2. Map Physical Interfaces to Port Channel: Associate the physical interfaces connected to VTEP1 and VTEP2 with the port channel.

```
(config)# interface xe2
(config-if)# channel-group 1 mode active
(config)# interface xe8
(config-if)# channel-group 1 mode active
```

3. Configure Layer 3 VLAN Interfaces:Set up the SVI (Switch Virtual Interface) for the subnets to facilitate communication.

```
(config)# interface vlan1.20
(config)# ip address 21.21.21.2/24
```

```
(config)# ipv6 address 21:21::21:2/48
```

Switch2 (SW2) Configuration

SW2 connects Tenant System TS2-21 to VTEP4.

1. OSPF Routing Configuration:Configure OSPF to maintain reachability within the underlay network.

```
(config)# router ospf 100
(config-router)# ospf router-id 6.6.6.6
(config-router)# network 6.6.6.6/32 area 0.0.0.0
(config-router)# network 100.14.14.0/24 area 0.0.0.0
(config-router)# bfd all-interfaces
```

Note: Perform the Anycast Gateway steps on the VTEPs, the access switches (SW1 and SW2) must be correctly configured with matching VLANs and trunk settings to ensure that host traffic reaches the Anycast Gateways on the VTEPs.

CHAPTER 3 EVPN-VxLAN Symmetric IRB Support with Connected host

Overview

EVPN-IRB enables communication between two L2VNI's by employing Routing through IP-VRF. This functionality offers Host (/32) based Symmetric IRB support, directing inter-subnet traffic directly to the Host attached VTEP. To implement this, configure "evpn irb-advertise-host-route" within VNID (BGP type 2) settings, or "redistribute connected-host-routes" under BGP (BGP type 5) configurations.

Note:

- In EVPN-VxLAN Interface-less mode, only the 'redistribute connected-host-routes' command is supported. However, in Interface-full mode, both commands are supported.
- It is advisable to configure a route map in ESI configured Multi-Homed (MH) nodes to block Hosts from peer MH. This configuration is not required in non-ESI MH VTEPs.

Feature Characteristics

The preferred and recommended approach for AOS-CX VXLAN/EVPN Distributed L3 Gateways is Symmetric IRB. This implementation offers superior scalability by eliminating the need to manage MAC/ARP entries for both source and destination hosts, and it doesn't require configuring the same VLAN/VNI as in the case of Asymmetric IRB. These advantages facilitate simpler and more scalable deployments in both Data Center and Campus networks.

Benefits

The advantages of utilizing a EVPN-VxLAN Symmetric IRB Support:

- Routing is employed on both ingress and egress VTEPs.
- Bi-directional traffic follows a symmetric path, such as utilizing an L3 VNI per VRF.
- VTEPs are relieved from holding unnecessary ARP/MAC resources.
- Configuration of the destination VLAN/VNI on the source VTEP is unnecessary.

Configuration

Ensure that the VTEPs have a base configuration with Symmetric IRB settings. Then, initiate dynamic traffic from VTEP4 originating from the same subnet (53.1.X.XX/XXXX::XX) as the IRB interface. Typically, in EVPN, a single IP-VRF can accommodate multiple IRB interfaces. Each IRB interface corresponds to a VNI, and multiple VNIs can be associated with a MAC-VRF.

Topology

The procedures in this section use the topology in Figure 3-1

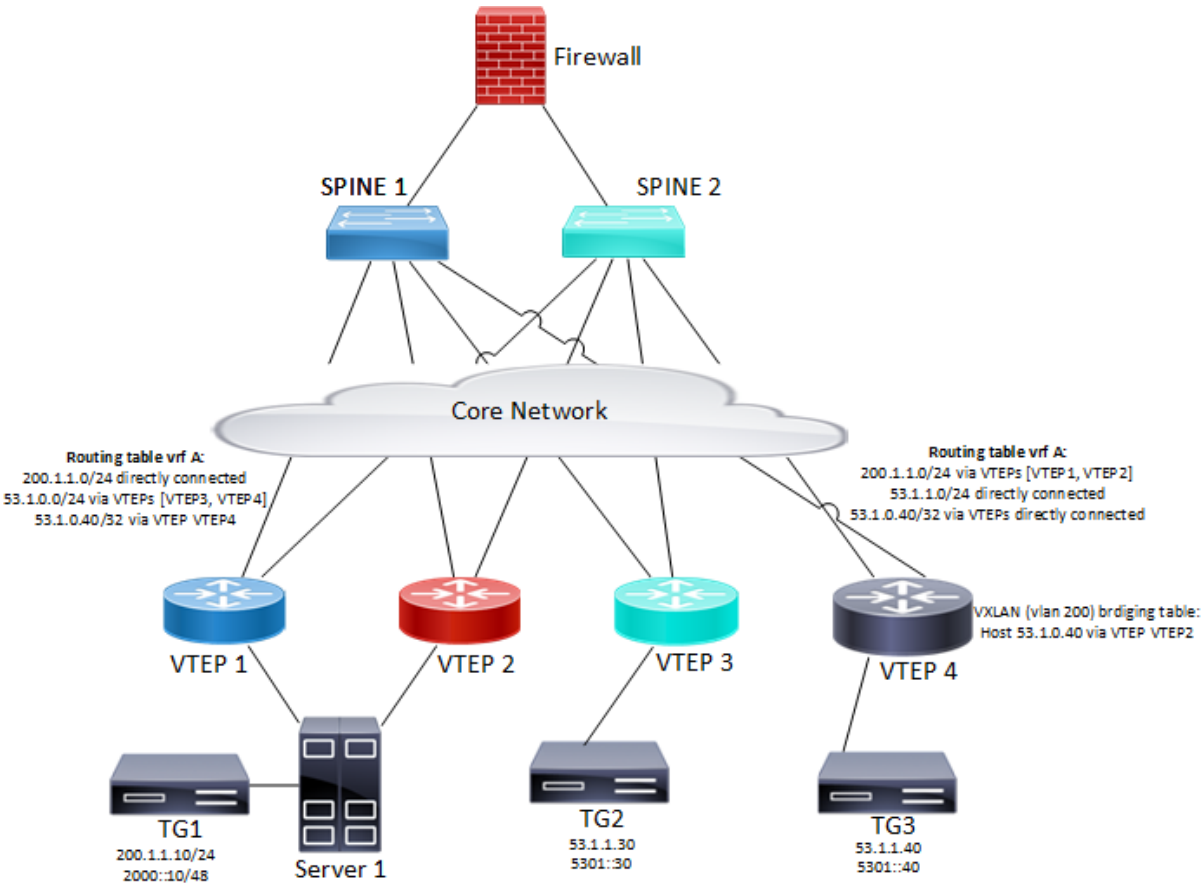

Figure 3-1: VxLAN EVPN IRB Connected host

Note: In the above topology TG1 is Multi homed Host and TG2 and TG3 are Single homed host with same subnet configured so there will be ECMP for 53 network in VTEP1 and VTEP2.

Base Configurations

Begin with a basic configuration that includes Symmetric IRB configurations on VTEPs, then initiate dynamic traffic transmission from VTEP4 within the same subnet (53.1.1.40/5301::40) as the IRB interface.

Validation

Verification before configuring evpn irb-advertise-host-route under VNID configurations or redistribute connected-host-routes under bgp.

In VTEP1:

```
VTEP1#show ip route vrf vxlan_l3_elan_mhsh
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
```

```
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "vxlan_l3_elan_mhsh"
B            2.2.2.2/32 [0/0] is directly connected, tunvxlan3, 00:21:33
B            5.5.5.5/32 [0/0] is directly connected, tunvxlan3, 00:21:33
B            6.6.6.6/32 [0/0] is directly connected, tunvxlan3, 00:21:33
B            53.1.1.0/24 [200/0] via 6.6.6.6 (recursive is directly connected,
tunvxlan3), 00:21:34
                                 [200/0] via 5.5.5.5 (recursive is directly connected,
tunvxlan3)
C            127.0.0.0/8 is directly connected, lo.vxlan_l3_elan_mhsh, 07:17:43
C            200.1.1.0/24 is directly connected, irb1604, 07:17:41

Gateway of last resort is not set
VTEP1#
VTEP1#show ipv6 route vrf vxlan_l3_elan_mhsh
IPv6 Routing Table
IP Route Table for VRF "vxlan_l3_elan_mhsh"
C       ::1/128 via ::, lo.vxlan_l3_elan_mhsh, 07:18:01
B       ::ffff:202:202/128 [0/0] via ::, tunvxlan3, 00:21:51
B       ::ffff:505:505/128 [0/0] via ::, tunvxlan3, 00:21:51
B       ::ffff:606:606/128 [0/0] via ::, tunvxlan3, 00:21:51
C       2000::/48 via ::, irb1604, 07:17:59
B       5301::/48 [200/0] via ::ffff:606:606 (recursive via ::, tunvxlan3), 00:21:52
                  [200/0] via ::ffff:505:505 (recursive via ::, tunvxlan3)
C       fe80::/64 via ::, irb1604, 07:17:59
VTEP1#
```

In VTEP2:

```
VTEP2#show ip route vrf vxlan_l3_elan_mhsh
IP Route Table for VRF "vxlan_l3_elan_mhsh"
B            1.1.1.1/32 [0/0] is directly connected, tunvxlan3, 00:22:50
B            5.5.5.5/32 [0/0] is directly connected, tunvxlan3, 00:22:50
B            6.6.6.6/32 [0/0] is directly connected, tunvxlan3, 00:22:50
B            53.1.1.0/24 [200/0] via 6.6.6.6 (recursive is directly connected,
tunvxlan3), 00:22:51
                                 [200/0] via 5.5.5.5 (recursive is directly connected,
tunvxlan3)
C            127.0.0.0/8 is directly connected, lo.vxlan_l3_elan_mhsh, 07:19:21
C            200.1.1.0/24 is directly connected, irb1604, 07:19:19

Gateway of last resort is not set
VTEP2#
VTEP2#
VTEP2#show ipv6 route vrf vxlan_l3_elan_mhsh
IPv6 Routing Table
IP Route Table for VRF "vxlan_l3_elan_mhsh"
C       ::1/128 via ::, lo.vxlan_l3_elan_mhsh, 07:19:22
```

```
B       ::ffff:101:101/128 [0/0] via ::, tunvxlan3, 00:22:51
B       ::ffff:505:505/128 [0/0] via ::, tunvxlan3, 00:22:51
B       ::ffff:606:606/128 [0/0] via ::, tunvxlan3, 00:22:51
C       2000::/48 via ::, irb1604, 07:19:20
B       5301::/48 [200/0] via ::ffff:606:606 (recursive via ::, tunvxlan3), 00:22:51
                  [200/0] via ::ffff:505:505 (recursive via ::, tunvxlan3)
C       fe80::/64 via ::, irb1604, 07:19:20
VTEP2#
VTEP2#show bgp l2vpn evpn mac-ip | grep 0000:0053:0040
0                         605       0000:0053:0040 --
605             0         6.6.6.6         --              VxLAN
0                          605        0000:0053:0040 53.1.1.40
605             0         6.6.6.6         --              VxLAN
0                          605        0000:0053:0040 5301::40
605             0         6.6.6.6         --              VxLAN
VTEP2#
```

In VTEP4:

```
VTEP4#show bgp l2vpn evpn mac-ip | grep 0000:0053:0040
0                         605       0000:0053:0040 --
605             0         6.6.6.6         --              VxLAN
0                          605        0000:0053:0040 53.1.1.40
605             0         6.6.6.6         --              VxLAN
0                          605        0000:0053:0040 5301::40
605             0         6.6.6.6         --              VxLAN
VTEP4#
```

Evpn irb-advertise-host-route configuration

1. To enable the EVPN irb-advertise-host-route, execute the following command in the config mode.
   ```
   (config)#nvo vxlan id 605 ingress-replication inner-vid-disabled
   (config-nvo)#vxlan host-reachability-protocol evpn-bgp vxlan_l2_elan_sh2
   (config-nvo)#evpn irb605
   (config-nvo)#evpn irb-advertise-host-route
   ```
2. To redistributed connected-host-routes, execute the following command.
   ```
   (config)#nvo vxlan id 605 ingress-replication inner-vid-disabled
   (config)#router bgp 1
   (config-router)#address-family ipv4 vrf vxlan_l3_elan_sh
   VTEP4(config-router-af)#redistribute connected-host-routes
   ```

Note: With static mac ip configured on vxlan access interface and when redistribute connected-host-routes is configured under bgp. Then routes will not be advertised as /32 or /128 because for static mac-ip Arp entry will not be present so only for dynamic routes.

Note: With redistribute connected-host-routes, show bgp l2vpn evpn mac-ip will not show the l3vnid.

Validation

Use this command to validate the EVPN-VxLAN Symmetric IRB.

In VTEP1:

```
VTEP1#show ip route vrf vxlan_l3_elan_mhsh
IP Route Table for VRF "vxlan_l3_elan_mhsh"
B            2.2.2.2/32 [0/0] is directly connected, tunvxlan3, 00:37:03
B            5.5.5.5/32 [0/0] is directly connected, tunvxlan3, 00:37:03
B            6.6.6.6/32 [0/0] is directly connected, tunvxlan3, 00:37:03
B            53.1.1.0/24 [200/0] via 6.6.6.6 (recursive is directly connected,
tunvxlan3), 00:37:04
                                 [200/0] via 5.5.5.5 (recursive is directly connected,
tunvxlan3)
B            53.1.1.40/32 [200/0] via 6.6.6.6 (recursive is directly connected,
tunvxlan3), 00:05:49
C            127.0.0.0/8 is directly connected, lo.vxlan_l3_elan_mhsh, 07:33:13
C            200.1.1.0/24 is directly connected, irb1604, 07:33:11
Gateway of last resort is not set
VTEP1#
VTEP1#show ipv6 route vrf vxlan_l3_elan_mhsh
IPv6 Routing Table
IP Route Table for VRF "vxlan_l3_elan_mhsh"
C       ::1/128 via ::, lo.vxlan_l3_elan_mhsh, 07:33:21
B       ::ffff:202:202/128 [0/0] via ::, tunvxlan3, 00:37:11
B       ::ffff:505:505/128 [0/0] via ::, tunvxlan3, 00:37:11
B       ::ffff:606:606/128 [0/0] via ::, tunvxlan3, 00:37:11
C       2000::/48 via ::, irb1604, 07:33:19
B       5301::/48 [200/0] via ::ffff:606:606 (recursive via ::, tunvxlan3), 00:37:12
                          [200/0] via ::ffff:505:505 (recursive via ::, tunvxlan3)
B       5301::40/128 [200/0] via ::ffff:606:606 (recursive via ::, tunvxlan3), 00:05:57
C       fe80::/64 via ::, irb1604, 07:33:19
VTEP1#
VTEP1#show bgp l2vpn evpn mac-ip | grep 0000:0053:0040
0                        605       0000:0053:0040 --
605              0         6.6.6.6         --              VxLAN
0                         605       0000:0053:0040 53.1.1.40
605            1604      6.6.6.6         --              VxLAN
0                         605       0000:0053:0040 5301::40
605            1604      6.6.6.6         --              VxLAN
VTEP1#
```

In VTEP2:

```
VTEP2#show ip route vrf vxlan_l3_elan_mhsh
IP Route Table for VRF "vxlan_l3_elan_mhsh"
B            1.1.1.1/32 [0/0] is directly connected, tunvxlan3, 00:31:16
B            5.5.5.5/32 [0/0] is directly connected, tunvxlan3, 00:31:16
B            6.6.6.6/32 [0/0] is directly connected, tunvxlan3, 00:31:16
B            53.1.1.0/24 [200/0] via 6.6.6.6 (recursive is directly connected,
tunvxlan3), 00:31:17
                                 [200/0] via 5.5.5.5 (recursive is directly connected,
tunvxlan3)
B            53.1.1.40/32 [200/0] via 6.6.6.6 (recursive is directly connected,
tunvxlan3), 00:00:03
C            127.0.0.0/8 is directly connected, lo.vxlan_l3_elan_mhsh, 07:27:47
```

```
C            200.1.1.0/24 is directly connected, irb1604, 07:27:45
Gateway of last resort is not set
VTEP2#
VTEP2#show ipv6 route vrf vxlan_l3_elan_mhsh
IPv6 Routing Table
IP Route Table for VRF "vxlan_l3_elan_mhsh"
C       ::1/128 via ::, lo.vxlan_l3_elan_mhsh, 07:27:54
B       ::ffff:101:101/128 [0/0] via ::, tunvxlan3, 00:31:23
B       ::ffff:505:505/128 [0/0] via ::, tunvxlan3, 00:31:23
B       ::ffff:606:606/128 [0/0] via ::, tunvxlan3, 00:31:23
C       2000::/48 via ::, irb1604, 07:27:52
B       5301::/48 [200/0] via ::ffff:606:606 (recursive via ::, tunvxlan3), 00:31:23
                  [200/0] via ::ffff:505:505 (recursive via ::, tunvxlan3)
B       5301::40/128 [200/0] via ::ffff:606:606 (recursive via ::, tunvxlan3), 00:00:10
C       fe80::/64 via ::, irb1604, 07:27:52
VTEP2#
VTEP2#show bgp l2vpn evpn mac-ip | grep 0000:0053:0040
0                         605       0000:0053:0040 --
605           0         6.6.6.6         --              VxLAN
0                          605        0000:0053:0040 53.1.1.40
605           1604      6.6.6.6         --              VxLAN
0                          605        0000:0053:0040 5301::40
605           1604      6.6.6.6         --              VxLAN
VTEP2#
```

In VTEP4:

```
VTEP4#show arp vrf vxlan_l3_elan_mhsh
Flags: D - Static Adjacencies attached to down interface
IP ARP Table for context vxlan_l3_elan_mhsh
Total number of entries: 1
Address          Age      MAC Address     Interface  State
1.1.1.1             -     e8c5.7aa3.2cb0  tunvxlan3  PERMANENT
2.2.2.2             -     e001.a657.ef01  tunvxlan3  PERMANENT
5.5.5.5             -     6cb9.c5b1.ab9c  tunvxlan3  PERMANENT
53.1.1.40        00:02:57 0000.0053.0040  irb604     STALE
VTEP4#

VTEP4#show bgp l2vpn evpn mac-ip | grep 0000:0053:0040
0                         605       0000:0053:0040 --
605           0         6.6.6.6         --              VxLAN
0                          605        0000:0053:0040 53.1.1.40
605           1604      6.6.6.6         --              VxLAN
0                          605        0000:0053:0040 5301::40
605           1604      6.6.6.6         --              VxLAN
VTEP4#
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
ECMP	ECMP stands for Equal-Cost Multi-Path. It's a routing technique used in computer networks, particularly in IP-based routing protocols like OSPF (Open Shortest Path First) and BGP (Border Gateway Protocol).
VTEPs	VTEP stands for VXLAN Tunnel Endpoint. It's a crucial component in network virtualization architectures, particularly in overlay networks using VXLAN (Virtual Extensible LAN) technology.
ARP/MAC	ARP (Address Resolution Protocol) and MAC (Media Access Control) address are both essential components of networking, particularly in Ethernet-based networks.
VLAN/VNI	VLAN (Virtual Local Area Network) and VNI (Virtual Network Identifier) are both technologies used in networking to segment and manage traffic within a larger network infrastructure.

CHAPTER 4 VXLAN-IRB-Inter-VRF Route Leaking

A VRF is a mechanism used to provide logical separation between routing tables on the same router. It is locally significant to the router. Each interface on a router can only be assigned to one VRF, but a VRF can have multiple interfaces. VRF route leaking can be done using route-target import/export.

The routes of VRF catering shared services shall be leaked to tenant VRFs. The leaking of routes shall be possible over one overlay VRF to another overlay VRF. By doing so shared services like Internet access through gateway routes can be made possible. Introduction of this feature shall cater various use cases of shared services like storage / Internet access etc.

Topology

The procedures in this section use the topology in Figure 4-2.

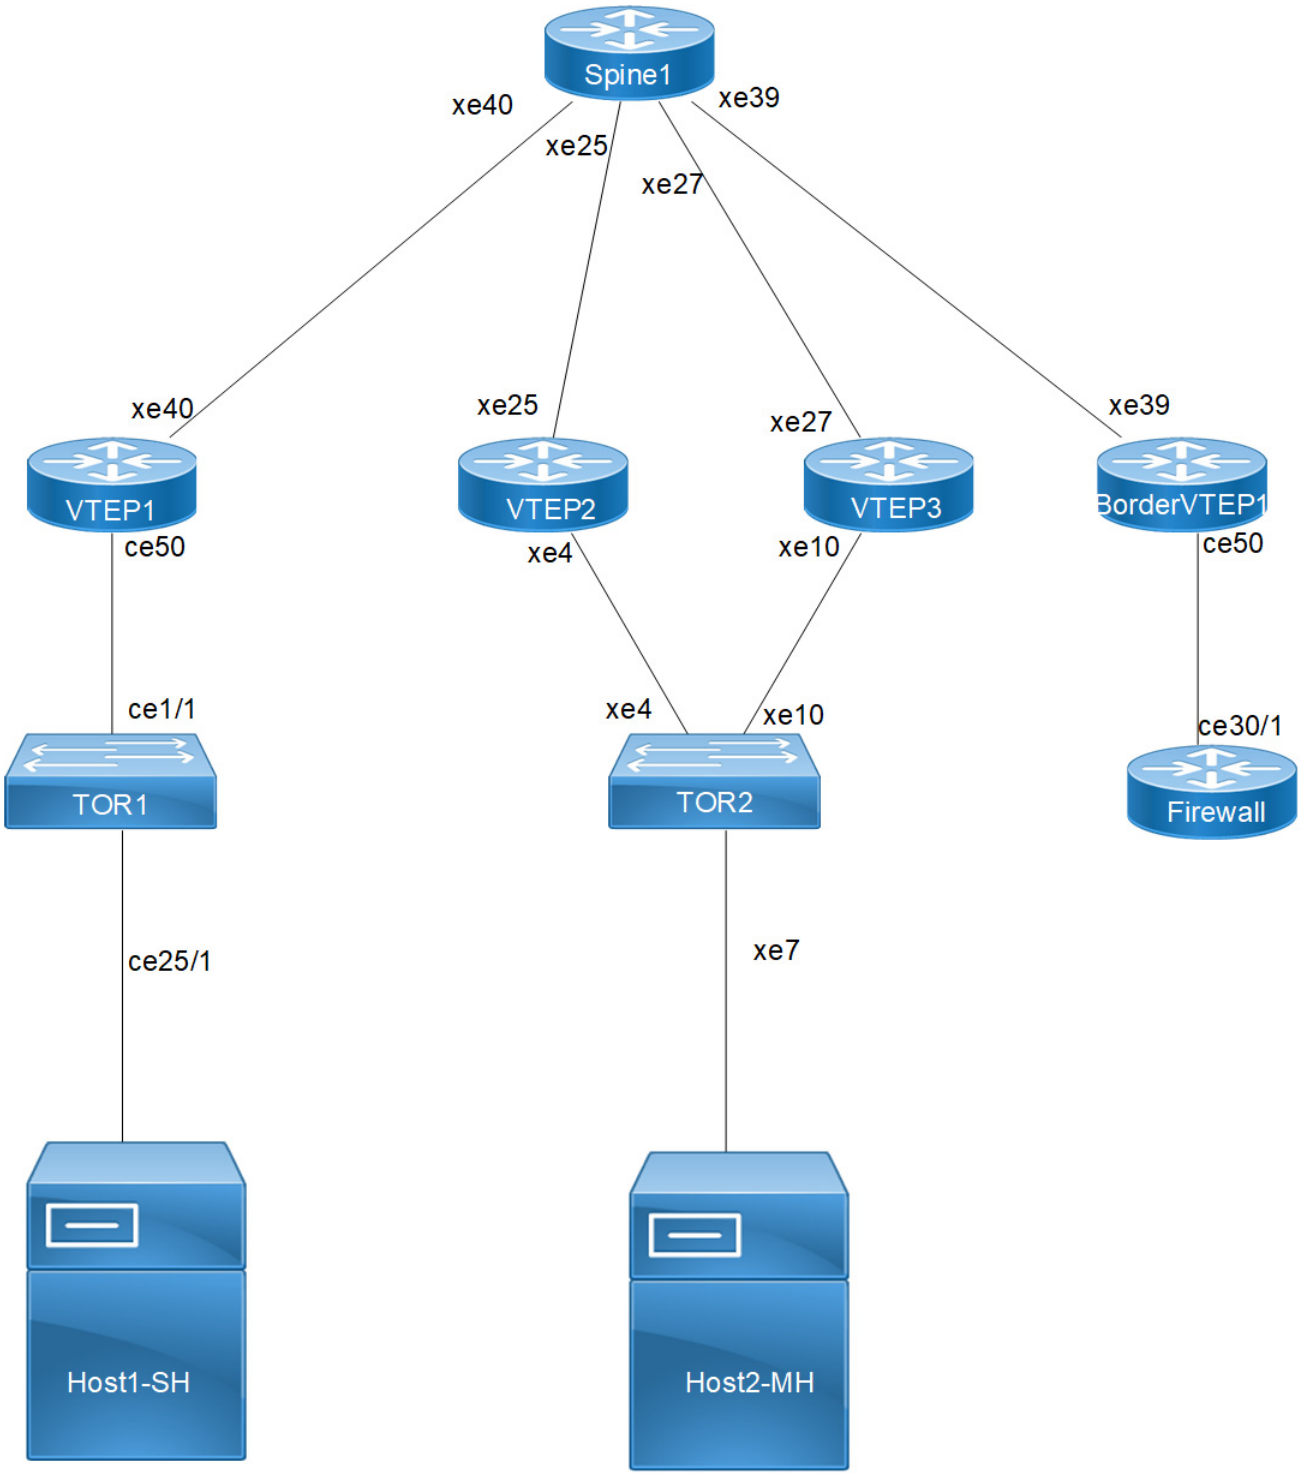

Figure 4-2: VxLAN_EVPN_IVRF

Note: SH means Single homing host and MH means Multihoming host.

VTEP1

Single Home -SH

Hardware profile and generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#nvo vxlan enable`	Enable VXLAN
`(config)#nvo vxlan irb`	Enable VXLAN irb
`(Config)#qos enable`	Enabling qos
`(Config)# bfd interval 3 minrx 3 multiplier 3`	Configure bfd
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

`(config)#interface ce50`	Enter Interface mode for ce50 (SH1)
`(config-if)# description ***Connected to TOR1***`	Interface description
`(config-if)#switchport`	Make it L2 interface
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 51.51.51.51/32 secondary`	Configure loopback ip address
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) interface xe40`	Enter interface mode
`(config-if)# description ***Connected to Spine2***`	Interface description
`(config-if)# ip address 10.10.10.1/31`	Configure ip address on network side of Spine1
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board
`(config-if)#exit`	Exit Interface mode and return to Configure mode.

OSPF configuration:

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 51.51.51.51`	Configure router-id as lo ip address
`(config-router)#network 51.51.51.51/32 area 0.0.0.0`	Add lo ip address network into area 0
`(config-router)#network 10.10.10.0/24 area 0.0.0.0`	Add Spine-connected network into area 0
`(config-router)#bfd all-interfaces`	Enabling bfd on all ospf interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

BGP configuration:

(config)#router bgp 500	Enter into Router BGP mode
(config-router)#bgp router-id 51.51.51.51	Configure router-id as lo ip address
(config-router)#neighbor 66.66.66.66 remote-as 500	Specify a BorderVTEP1 loopback ip address and remote-as defined
(config-router)#neighbor 66.66.66.66 update-source lo	Configure update as loopback for BorderVTEP1
(config-router)#neighbor 66.66.66.66 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for BorderVTEP1
(config-router)# address-family ipv4 unicast	Enter into IPV4 unicast address family mode
(config-router-af)# neighbor 66.66.66.66 activate	Activate BorderVTEP1 into ipv4 unicast family
(config-router-af)# exit-address-family	Exit from IPV4 unicast address family
(config-router)#address-family l2vpn evpn	Enter into l2vpn EVPN address family mode
(config-router-af)#neighbor 66.66.66.66 activate	Activate BorderVTEP1 into l2vpn evpn address family mode
(config-router-af)#exit-address-family	Exit from l2vpn address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L2 VRF Configuration:

(config)# mac vrf RED	Create mac routing/forwarding instance with RED name and enter into vrf mode
(config-vrf)# rd 2.3.4.5:1	Assign RD value
(config-vrf)# route-target both 6000:6000	Assign route-target value for same for import and export. Should be same on all node for RED
(config-vrf)#exit	Exit from vrf mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L3 VRF and BGP Configuration:

(config)# ip vrf FAX	Create mac routing/forwarding instance with FAX name and enter into vrf mode
(config-vrf)# rd 51.51.51.51:1050	Assign RD value
(config-vrf)# route-target both 1050:1050	Assign route-target value for same for import and export.
(config-vrf)# l3vni 10502	Configure L3VNI as 10502 for FAX vrf
(config-vrf)#exit	Exit from vrf mode
(config)# interface irb 1050	Configure irb interface 1050
(config-if)ip vrf forwarding FAX	Configure FAX
(config-if) ip address 10.12.32.1/24	Configure ip address
(config-if)exit	Exit from interface config mode

(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)router bgp 500	Enter into bgp router mode
(config-router)#address-family ipv4 vrf FAX	Enter into address-family mode for FAX
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit form address-family
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan vtep-ip-global 51.51.51.51	Configure Source vtep-ip-global configuration. Use loopback ip address
(config)#nvo vxlan id 1050 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp RED	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1050	Configure irb1050 under vxlan id 1050
(config)# nvo vxlan access-if port-vlan ce50 1050	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN)
(config-nvo-acc-if)# map vnid 1050	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.3333.1050 ip 10.12.32.10	Configure static mac-ip
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VTEP2

(Multi-homed group) - Part of both Multi-homed with po1000(MH).

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan irb	Enable VXLAN irb
(config)# evpn irb-forwarding anycast-gateway-mac 0000.2222.3333	Configure Anycast gateway mac
(Config)#qos enable	Enabling qos
(Config)# bfd interval 3 minrx 3 multiplier 3	Configure bfd
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

(config)#interface xe4	Enter Interface mode for xe4(MH)
(config-if)# description ***Connected to TOR2***	Interface description
(config-if)# channel-group 1000 mode active	Make it member of po1000
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)# interface po1000	Enter into po1000 mode
(config-if)# switchport	Configure L2 mode
(config-if)# evpn multi-homed system-mac 0000.4444.5555	Configure System mac
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 60.60.60.60/32	Configure loopback ip address
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config) interface xe25	Enter interface mode
(config-if)# description ***Connected to Spine1***	Interface description
(config-if)# ip address ip address 10.10.12.1/31	Configure ip address on network side of Spine1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 60.60.60.60	Configure router-id as lo ip address
(config-router)#network 60.60.60.60/32 area 0.0.0.0	Add lo ip address network into area 0
(config-router)#network 10.10.12.0/24 area 0.0.0.0	Add Spine-connected network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

BGP configuration:

(config)#router bgp 500	Enter into Router BGP mode
(config-router)#bgp router-id 60.60.60.60	Configure router-id as lo ip address
(config-router)#neighbor 66.66.66.66 remote-as 500	Specify a BorderVTEP1 loopback ip address and remote-as defined
(config-router)#neighbor 66.66.66.66 update-source lo	Configure update as loopback for BorderVTEP1
(config-router)#neighbor 66.66.66.66 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for BorderVTEP1

(config-router)#neighbor 76.76.76.76 remote-as 500	Specify a VTEP3 loopback ip address and remote-as defined
(config-router)#neighbor 76.76.76.76 update-source lo	Configure update as loopback for VTEP3
(config-router)#neighbor 76.76.76.76 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP3
(config-router)# address-family ipv4 unicast	Enter into IPV4 unicast address family mode
(config-router-af)# neighbor 66.66.66.66 activate	Activate BorderVTEP1 into ipv4 unicast family
(config-router-af)# neighbor 76.76.76.76 activate	Activate VTEP3 into ipv4 unicast family
(config-router-af)# exit-address-family	Exit from IPV4 unicast address family
(config-router)#address-family l2vpn evpn	Enter into l2vpn EVPN address family mode
(config-router-af)#neighbor 66.66.66.66 activate	Activate BorderVTEP1 into l2vpn evpn address family mode
(config-router-af)#neighbor 76.76.76.76 activate	Activate VTEP3 into l2vpn evpn address family mode
(config-router-af)#exit-address-family	Exit from l2vpn address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L2 VRF Configuration:

(config)# mac vrf RED	Create mac routing/forwarding instance with RED name and enter into vrf mode
(config-vrf)# rd 2.3.4.5:2	Assign RD value
(config-vrf)# route-target both 6000:6000	Assign route-target value for same for import and export. Should be same on all node for RED
(config-vrf)#exit	Exit from vrf mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L3 VRF and BGP Configuration:

(config)# ip vrf SMS	Create mac routing/forwarding instance with SMS name and enter into vrf mode
(config-vrf)# rd 60.60.60.60:1040	Assign RD value
(config-vrf)# route-target both 1040:1040	Assign route-target value for same for import and export.
(config-vrf)# l3vni 10402	Configure L3VNI as 10402 for SMS vrf
(config-vrf)#exit	Exit from vrf mode
(config)# interface irb 1060	Configure irb interface 1060
(config-if)ip vrf forwarding SMS	Configure ip vrf forwarding
(config-if) ip address 10.240.38.1/24	Configure ip address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Anycast mac configured
(config-if)exit	Exit from interface config mode

(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)router bgp 500	Enter into bgp router mode
(config-router)#address-family ipv4 vrf SMS	Enter into address-family mode for SMS
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit form address-family
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan vtep-ip-global 60.60.60.60	Configure Source vtep-ip-global configuration. Use loopback ip address
(config)#nvo vxlan id 1060 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp RED	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1060	Configure irb1060 under vxlan id 1060
(config)# nvo vxlan access-if port-vlan po1000 1060	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)# map vnid 1060	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1060 ip 10.240.38.10	Configure static mac-ip
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VTEP3

(Multi-homed group) - Part of both Multi-homed with po1000(MH).

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan irb	Enable VXLAN irb
(config)# evpn irb-forwarding anycast-gateway-mac 0000.2222.3333	Configure Anycast gateway mac
(Config)#qos enable	Enabling qos
(Config)# bfd interval 3 minrx 3 multiplier 3	Configure bfd
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

`(config)#interface xe10`	Enter Interface mode for xe10(MH)
`(config-if)# description ***Connected to TOR2***`	Interface description
`(config-if)# channel-group 1000 mode active`	Make it member of po1000
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)# interface po1000`	Enter into po1000 mode
`(config-if)# switchport`	Configure L2 mode
`(config-if)# evpn multi-homed system-mac 0000.4444.5555`	Configure System mac
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 76.76.76.76/32 secondary`	Configure loopback ip address
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) interface xe27`	Enter interface mode
`(config-if)# description ***Connected to Spine1***`	Interface description
`(config-if)# ip address ip address 10.10.24.1/31`	Configure ip address on network side of Spine1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 76.76.76.76`	Configure router-id as lo ip address
`(config-router)#network 76.76.76.76/32 area 0.0.0.0`	Add lo ip address network into area 0
`(config-router)#network 10.10.24.0/24 area 0.0.0.0`	Add Spine-connected network into area 0
`(config-router)#bfd all-interfaces`	Enabling bfd on all ospf interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

BGP configuration:

`(config)#router bgp 500`	Enter into Router BGP mode
`(config-router)#bgp router-id 76.76.76.76`	Configure router-id as lo ip address
`(config-router)#neighbor 66.66.66.66 remote-as 500`	Specify a BorderVTEP1 loopback ip address and remote-as defined
`(config-router)#neighbor 66.66.66.66 update-source lo`	Configure update as loopback for BorderVTEP1

(config-router)#neighbor 66.66.66.66 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for BorderVTEP1
(config-router)#neighbor 60.60.60.60 remote-as 500	Specify a VTEP2 loopback ip address and remote-as defined
(config-router)#neighbor 60.60.60.60 update-source lo	Configure update as loopback for VTEP2
(config-router)#neighbor 60.60.60.60 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP2
(config-router)# address-family ipv4 unicast	Enter into IPV4 unicast address family mode
(config-router-af)# neighbor 66.66.66.66 activate	Activate BorderVTEP1 into ipv4 unicast family
(config-router-af)# neighbor 60.60.60.60 activate	Activate VTEP2 into ipv4 unicast family
(config-router-af)# exit-address-family	Exit from IPV4 unicast address family
(config-router)#address-family l2vpn evpn	Enter into l2vpn EVPN address family mode
(config-router-af)#neighbor 66.66.66.66 activate	Activate BorderVTEP1 into l2vpn evpn address family mode
(config-router-af)#neighbor 60.60.60.60 activate	Activate VTEP2 into l2vpn evpn address family mode
(config-router-af)#exit-address-family	Exit from l2vpn address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L2 VRF Configuration:

(config)# mac vrf RED	Create mac routing/forwarding instance with RED name and enter into vrf mode
(config-vrf)# rd 2.3.4.6:2	Assign RD value
(config-vrf)# route-target both 6000:6000	Assign route-target value for same for import and export. Should be same on all node for RED
(config-vrf)#exit	Exit from vrf mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L3 VRF and BGP Configuration:

(config)# ip vrf SMS	Create mac routing/forwarding instance with SMS name and enter into vrf mode
(config-vrf)# rd 76.76.76.76:1040	Assign RD value
(config-vrf)# route-target both 1040:1040	Assign route-target value for same for import and export.
(config-vrf)# l3vni 10402	Configure L3VNI as 10402 for SMS vrf
(config-vrf)#exit	Exit from vrf mode
(config)# interface irb 1060	Configure irb interface 1060
(config-if)ip vrf forwarding SMS	Configure ip vrf forwarding
(config-if) ip address 10.240.38.1/24	Configure ip address

(config-if) evpn irb-if-forwarding anycast-gateway-mac	Anycast mac configured
(config-if)exit	Exit from interface config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)router bgp 500	Enter into bgp router mode
(config-router)#address-family ipv4 vrf SMS	Enter into address-family mode for SMS
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit form address-family
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VXLAN configuration:

(config)#nvo vxlan vtep-ip-global 76.76.76.76	Configure Source vtep-ip-global configuration. Use loopback ip address
(config)#nvo vxlan id 1060 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp RED	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1060	Configure irb1060 under vxlan id 1060
(config)# nvo vxlan access-if port-vlan po1000 1060	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)# map vnid 1060	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1060 ip 10.240.38.10	Configure static mac-ip
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

BorderVTEP1

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan irb	Enable VXLAN irb
(Config)#qos enable	Enabling qos
(Config)# bfd interval 3 minrx 3 multiplier 3	Configure bfd
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

`(config)#interface ce50`	Enter Interface mode
`(config-if)# description ***Connected to FW***`	Interface description
`(config-if)# switchport`	Configure L2 mode
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address  66.66.66.66/32`	Configure loopback ip address
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) interface xe39`	Enter interface mode
`(config-if)# description ***Connected to Spine1***`	Interface description
`(config-if)# ip address ip address 10.10.14.1/31`	Configure ip address on network side of Spine1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id  66.66.66.66`	Configure router-id as lo ip address
`(config-router)#network  66.66.66.66/32 area 0.0.0.0`	Add lo ip address network into area 0
`(config-router)#network 10.10.14.0/24 area 0.0.0.0`	Add Spine-connected network into area 0
`(config-router)#bfd all-interfaces`	Enabling bfd on all ospf interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

BGP configuration:

`(config)#router bgp 500`	Enter into Router BGP mode
`(config-router)#bgp router-id 66.66.66.66`	Configure router-id as lo ip address
`(config-router)#neighbor 51.51.51.51 remote-as 500`	Specify a VTEP1 loopback ip address and remote-as defined
`(config-router)#neighbor 51.51.51.51 update-source lo`	Configure update as loopback for VTEP1
`(config-router)#neighbor 51.51.51.51 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP1
`(config-router)#neighbor 60.60.60.60 remote-as 500`	Specify a VTEP2 loopback ip address and remote-as defined
`(config-router)#neighbor 60.60.60.60 update-source lo`	Configure update as loopback for VTEP2
`(config-router)#neighbor 60.60.60.60 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP2

(config-router)#neighbor 76.76.76.76 remote-as 500	Specify a VTEP3 loopback ip address and remote-as defined
(config-router)#neighbor 76.76.76.76 update-source lo	Configure update as loopback for VTEP3
(config-router)#neighbor 76.76.76.76 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP3
(config-router)# no bgp default ipv4-unicast	It will avoid default ipv4 unicast routing
(config-router)# address-family ipv4 unicast	Enter into IPV4 unicast address family mode
(config-router-af)# neighbor 51.51.51.51 activate	Activate VTEP1 into ipv4 unicast family
(config-router-af)# neighbor 60.60.60.60 activate	Activate VTEP2 into ipv4 unicast family
(config-router-af)# neighbor 76.76.76.76 activate	Activate VTEP3 into ipv4 unicast family
(config-router-af)# exit-address-family	Exit from IPV4 unicast address family
(config-router)#address-family l2vpn evpn	Enter into l2vpn EVPN address family mode
(config-router-af)#neighbor 51.51.51.51 activate	Activate VTEP1 into l2vpn evpn address family mode
(config-router-af)#neighbor 60.60.60.60 activate	Activate VTEP2 into l2vpn evpn address family mode
(config-router-af)#neighbor 76.76.76.76 activate	Activate VTEP3 into l2vpn evpn address family mode
(config-router-af)#exit-address-family	Exit from l2vpn address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L2 VRF Configuration:

(config)# mac vrf RED	Create mac routing/forwarding instance with RED name and enter into vrf mode
(config-vrf)# rd 2.2.4.4:4	Assign RD value
(config-vrf)# route-target both 6000:6000	Assign route-target value for same for import and export. Should be same on all node for RED
(config-vrf)#exit	Exit from vrf mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

L3 VRF and BGP Configuration:

(config)# ip vrf gvrf	Create mac routing/forwarding instance with gvrf name and enter into vrf mode
(config-vrf)# rd 4.5.6.8:6	Assign RD value
(config-vrf)# route-target import 100:100	Assign route-target value for import from FAX vrf
(config-vrf)# route-target import 300:300	Assign route-target value for import from SMS vrf
(config-vrf)# route-target export 1000:1000	Assign route-target value for export from gvrf
(config-vrf)# l3vni 500	Configure L3VNI as 500 for gvrf vrf

(config-vrf)#exit	Exit from vrf mode
(config)# interface irb 1067	Configure irb interface 1060
(config-if) ip vrf forwarding gvrf	Configure ip vrf forwarding
(config-if) ip address 10.10.18.1/24	
	Configure ip address
(config-if)exit	Exit from interface config mode
(config)# ip vrf FAX	Create mac routing/forwarding instance with FAX name and enter into vrf mode
(config-vrf)# rd 66.66.66.66:1050	Assign RD value
(config-vrf)# route-target both 1050:1050	Assign route-target value for same for import and export.
(config-vrf)# route-target export 100:100	Assign route-target value export from FAX
(config-vrf)# route-target import 1000:1000	Assign route-target value for import from gvrf
(config-vrf)# l3vni 10502	Configure L3VNI as 10502 for FAX vrf
(config-vrf)#exit	Exit from vrf mode
(config)# ip vrf SMS	Create mac routing/forwarding instance with SMS name and enter into vrf mode
(config-vrf)# rd 66.66.66.66:1060	Assign RD value
(config-vrf)# route-target both 1040:1040	Assign route-target value for same for import and export.
(config-vrf)# route-target export 300:300	Assign route-target value export from SMS
(config-vrf)# route-target import 1000:1000	Assign route-target value for import from gvrf
(config-vrf)# l3vni 10402	Configure L3VNI as 10402 for SMS vrf
(config-vrf)#exit	Exit from vrf mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board
(config)router bgp 500	Enter into bgp router mode
(config-router)#address-family ipv4 vrf gvrf	Enter into address-family mode for gvrf
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)# neighbor 10.10.18.2 remote-as 64603	Add Firewall as neighbor
(config-router-af)# neighbor 10.10.18.2 fall-over bfd	Configure bfd for better convergence
(config-router-af)# neighbor 10.10.18.2 activate	Activate the neighbor
(config-router-af)# neighbor 10.10.18.2 advertisement-interval 0	Configure interval 0 for better convergence
(config-router-af)#exit-address-family	Exit form address-family
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

VXLAN configuration:

(config)#nvo vxlan vtep-ip-global 66.66.66.66	Configure Source vtep-ip-global configuration. Use loopback ip address
(config)# nvo vxlan id 1067 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode

(config-nvo)#vxlan host-reachability-protocol evpn-bgp RED	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1067	Configure irb1067 under vxlan id 1067
(config)# nvo vxlan access-if port-vlan ce50 1067	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)# map vnid 1067	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Firewall

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(Config)#qos enable	Enabling qos
(Config)#bfd interval 3 minrx 3 multiplier 3	Configure bfd
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

#configure terminal	Enter Configure mode.
(config)# bridge 1 protocol rstp vlan-bridge	Configure rstp vlan bridge
(config)# vlan database	
(config)#vlan 1067 bridge 1 state enable	Configure vlans from 1067 and associate with bridge 1
(config)#interface ce30/1	Enter Interface mode for ce30/1
(config-if)# description ***Connected to BorderVTEP1***	Interface description
(config-if)#bridge-group 1	Associate to bridge 1
(config-if)# bridge-group 1 spanning-tree disable	Configure stp disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 1067	Trunk allowed vlan as 1067
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface vlan1.1067	Enter Interface mode
(config-if)# ip address 10.10.18.2/24	Configure ip address
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface ce1/1	Enter Interface mode
(config-if)#ip address 10.10.20.1/24	Configure ip address to advertise
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

BGP configuration:

(Config)#router bgp 64603	Enter into Router BGP mode
(config-router)# neighbor 10.10.18.1 remote-as 500	Specify a BorderVTEP1 gvrf ip address and remote-as defined
(config-router)# neighbor 10.10.18.1 fall-over bfd	Configure fall-over bfd for fast convergence
(config-router)#neighbor 10.10.18.1 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence
(config-router)# address-family ipv4 unicast	Enter into IPV4 unicast address family mode
(config-router-af)# network 10.10.20.0/24	Add lo adders as network for advertise
(config-router-af)# max-paths ebgp 8	Add max path
(config-router-af)# neighbor 10.10.18.1 activate	Activate the neighbor
(config-router-af)# neighbor 10.10.18.1 default-originate	Do default originate
(config-router-af)# exit-address-family	Exit from IPV4 unicast address family
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

TOR1 (SH)

#configure terminal	Enter Configure mode.
(config)# bridge 1 protocol rstp vlan-bridge	Configure rstp vlan bridge
(config)# vlan database	
(config)#vlan 1050 bridge 1 state enable	Configure vlans from 1050 and associate with bridge 1
(config)#interface ce1/1	Enter Interface mode for ce1/1
(config-if)#switchport	Make as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate to bridge 1
(config-if)# bridge-group 1 spanning-tree disable	Configure stp disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 1050	Trunk allowed vlan as 1050
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface ce25/1	Enter Interface mode for ce25/1
(config-if)#switchport	Make as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate to bridge 1
(config-if)# bridge-group 1 spanning-tree disable	Configure stp disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 1050	Trunk allowed vlan as 1050
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

TOR2 (MH)

Multihomed to 2-VTEPs (VTEP2 and VTEP3).

`#configure terminal`	Enter Configure mode.
`(config)# bridge 1 protocol rstp vlan-bridge`	Configure rstp vlan bridge
`(config)# vlan database`	
`(config)#vlan 1060 bridge 1 state enable`	Configure vlans from 1060 and associate with bridge 1
`(config)#interface po1000`	Enter Interface mode for po1000
`(config-if)#switchport`	Make as L2 port by configuring switchport
`(config-if)#bridge-group 1`	Associate to bridge 1
`(config-if)# bridge-group 1 spanning-tree disable`	Configure stp disable
`(config-if)# switchport mode trunk`	Mode as trunk
`(config-if)# switchport trunk allowed vlan add 1060`	Trunk allowed vlan as 1060
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe4`	Enter Interface mode for xe4
`(config-if)# channel-group 1000 mode active`	Make it member of po1000
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe7`	Enter Interface mode for xe7
`(config-if)# channel-group 1000 mode active`	Make it member of po1000
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe7`	Enter Interface mode for ce25/1
`(config-if)#switchport`	Make as L2 port by configuring switchport
`(config-if)#bridge-group 1`	Associate to bridge 1
`(config-if)# bridge-group 1 spanning-tree disable`	Configure stp disable
`(config-if)# switchport mode trunk`	Mode as trunk
`(config-if)# switchport trunk allowed vlan add 1060`	Trunk allowed vlan as 1060
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to running configuration and save config and reload board

Spine1

Spine node where all VTEPs are connected.

Generic configuration:

`#configure terminal`	Enter Configure mode.
`(Config)#qos enable`	Enabling qos

(Config)# bfd interval 3 minrx 3 multiplier 3	Configure bfd
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Interface and loopback configuration:

(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 62.62.62.62/32 secondary	Configure loopback ip address
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe40	Enter Interface mode
(config-if)# description ***Connected to VTEP1***	Description of interface
(config-if)#ip address ip address 10.10.10.2/31	Configure ip address on network side of VTEP1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe25	Enter Interface mode
(config-if)# description ***Connected to VTEP2***	Description of interface
(config-if)#ip address ip address 10.10.12.2/31	Configure ip address on network side of VTEP2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe27	Enter Interface mode
(config-if)# description ***Connected to VTEP3***	Description of interface
(config-if)#ip address ip address 10.10.24.2/31	Configure ip address on network side of VTEP3
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe39	Enter Interface mode
(config-if)# description ***Connected to BorderVTEP1***	Description of interface
(config-if)#ip address ip address 10.10.14.2/31	Configure ip address on network side of BorderVTEP1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 62.62.62.62	Configure router-id as lo ip address
(config-router)#network 62.62.62.62/32 area 0.0.0.0	Add lo ip address network into area 0
(config-router)#network 10.10.10.0/24 area 0.0.0.0	Add VTEP1 network into area 0

(config-router)#network 10.10.12.0/24 area 0.0.0.0	Add VTEP2 network into area 0
(config-router)#network 10.10.14.0/24 area 0.0.0.0	Add VTEP4 network into area 0
(config-router)#network 10.10.24.0/24 area 0.0.0.0	Add VTEP3 network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration and save config and reload board

Validations

Firewall

```
=========
Firewall#show ip bgp summary
BGP router identifier 10.10.19.2, local AS number 64603
BGP table version is 3
2 BGP AS-PATH entries
0 BGP community entries
8  Configured ebgp ECMP multipath: Currently set at 8

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Dow
n   State/PfxRcd
10.10.18.1               4   500   46         58       3      0      0  00:17:36
              3

Total number of neighbors 1

Total number of Established sessions 1
Firewall#show ip roy
Firewall#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.18.0/24 is directly connected, vlan1.1067, 00:19:40
C            10.10.20.0/24 is directly connected, ce1/1, 00:00:13
B            10.12.32.0/24 [20/0] via 10.10.18.1, vlan1.1067, 00:17:43
B            10.240.38.0/24 [20/0] via 10.10.18.1, vlan1.1067, 00:17:43
C            127.0.0.0/8 is directly connected, lo, 00:52:18
IP Route Table for VRF "management"
C            10.12.85.0/24 is directly connected, eth0, 00:52:07
```

```
C            127.0.0.0/8 is directly connected, lo.management, 00:52:18

Gateway of last resort is not set
FW#

BorderVTEP1
===========

BorderVTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID    VNI-Name    VNI-Type Type Interface ESI                           VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
1067    ----        --       AC   ce50     --- Single Homed Port ---    1067 ----
----             ----
10402   ----        L3       NW   ----     ----                         ---- ----
66.66.66.66      60.60.60.60
10402   ----        L3       NW   ----     ----                         ---- ----
66.66.66.66      76.76.76.76
10502   ----        L3       NW   ----     ----                         ---- ----
66.66.66.66      51.51.51.51

Total number of entries are 4
BorderVTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
66.66.66.66      51.51.51.51      Installed        00:22:36       00:22:36
66.66.66.66      60.60.60.60      Installed        00:22:36       00:22:36
66.66.66.66      76.76.76.76      Installed        00:22:36       00:22:36

Total number of entries are 3
BorderVTEP1#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            10.10.10.0/31 [110/2] via 10.10.14.0, xe39, 00:22:44
O            10.10.12.0/31 [110/2] via 10.10.14.0, xe39, 00:22:44
C            10.10.14.0/31 is directly connected, xe39, 00:23:29
```

```
O            10.10.24.0/31 [110/2] via 10.10.14.0, xe39, 00:22:44
O            51.51.51.51/32 [110/3] via 10.10.14.0, xe39, 00:22:44
O            60.60.60.60/32 [110/3] via 10.10.14.0, xe39, 00:22:44
C            66.66.66.66/32 is directly connected, lo, 00:23:32
O            76.76.76.76/32 [110/3] via 10.10.14.0, xe39, 00:22:44
C            127.0.0.0/8 is directly connected, lo, 00:24:12
IP Route Table for VRF "management"
C            10.12.86.0/24 is directly connected, eth0, 00:23:38
C            127.0.0.0/8 is directly connected, lo.management, 00:24:12
IP Route Table for VRF "gvrf"
Gateway of last resort is 10.10.18.2 to network 0.0.0.0

B*           0.0.0.0/0 [20/0] via 10.10.18.2, irb1067, 00:21:31
C            10.10.18.0/24 is directly connected, irb1067, 00:23:30
B            10.10.20.0/24 [20/0] via 10.10.18.2, irb1067, 00:04:01
B    v       10.12.32.0/24 [200/0] via 51.51.51.51 (recursive is directly connected,
tunvxlan3), 00:22:39
B    v       10.240.38.0/24 [200/0] via 60.60.60.60 (recursive is directly connected,
tunvxlan4), 00:22:40
C            127.0.0.0/8 is directly connected, lo.gvrf, 00:23:36
IP Route Table for VRF "SMS"
Gateway of last resort is 10.10.18.2 to network 0.0.0.0

B*   v       0.0.0.0/0 [20/0] via 10.10.18.2, irb1067, 00:21:31
B    v       10.10.18.0/24 [20/0] is directly connected, irb1067, 00:23:30
B    v       10.10.20.0/24 [20/0] via 10.10.18.2, irb1067, 00:04:01
B            10.240.38.0/24 [200/0] via 60.60.60.60 (recursive is directly connected,
tunvxlan4), 00:22:40
B            60.60.60.60/32 [0/0] is directly connected, tunvxlan4, 00:22:39
B            76.76.76.76/32 [0/0] is directly connected, tunvxlan4, 00:22:39
C            127.0.0.0/8 is directly connected, lo.SMS, 00:23:35
IP Route Table for VRF "FAX"
Gateway of last resort is 10.10.18.2 to network 0.0.0.0

B*   v       0.0.0.0/0 [20/0] via 10.10.18.2, irb1067, 00:21:31
B    v       10.10.18.0/24 [20/0] is directly connected, irb1067, 00:23:30
B    v       10.10.20.0/24 [20/0] via 10.10.18.2, irb1067, 00:04:01
B            10.12.32.0/24 [200/0] via 51.51.51.51 (recursive is directly connected,
tunvxlan3), 00:22:39
B            51.51.51.51/32 [0/0] is directly connected, tunvxlan3, 00:22:39
C            127.0.0.0/8 is directly connected, lo.FAX, 00:23:35
IP Route Table for VRF "SMM"
C            127.0.0.0/8 is directly connected, lo.SMM, 00:23:35

Gateway of last resort is not set
BorderVTEP1#                  show bgp l2vpn evpn summary
BGP router identifier 66.66.66.66, local AS number 500
BGP table version is 6
2 BGP AS-PATH entries
0 BGP community entries
```

```
Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
51.51.51.51               4   500   62         71       6      0      0  00:22:50
7       0       5       1       0       1
60.60.60.60               4   500   65         71       6      0      0  00:22:50
8       2       3       1       1       1
76.76.76.76               4   500   65         70       6      0      0  00:22:50
9       2       4       1       1       1

Total number of neighbors 3

Total number of Established sessions 3
BorderVTEP1#                    show bgp l2vpn evpn
BGP table version is 6, local router ID is 66.66.66.66
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network          Next Hop            Metric    LocPrf      Weight     Path  Peer
Encap

RD[2.2.4.4:4] VRF[RED]:
* i  [1]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i                  60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
* i  [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i                  60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
* i  [2]:[0]:[1050]:[48,0000:0e8d:561a]:[0]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
* i  [2]:[0]:[1050]:[48,0000:0e8d:561a]:[32,10.12.32.11]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
* i  [2]:[0]:[1050]:[48,0000:3333:1050]:[32,10.12.32.10]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
* i  [2]:[0]:[1050]:[48,3c2c:99d6:167a]:[32,10.12.32.1]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
* i  [2]:[0]:[1050]:[48,3c2c:99d6:167a]:[128,2401::1]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
* i  [2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:0e8d:5619]:[0]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i
[2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:2222:1060]:[32,10.240.38.10]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i                  60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
* i  [2]:[0]:[1060]:[48,0000:2222:3333]:[32,10.240.38.1]:[1060]
```

```
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i                   60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
* i  [2]:[0]:[1060]:[48,0000:2222:3333]:[128,1601::1]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i                   60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>   [2]:[0]:[1067]:[48,3c2c:991c:dc7a]:[32,10.10.18.1]:[1067]
                      66.66.66.66           0        100       32768  i  ----------
VXLAN
*>   [2]:[0]:[1067]:[48,a82b:b5cf:f806]:[32,10.10.18.2]:[1067]
                      66.66.66.66           0        100       32768  i  ----------
VXLAN
* i  [3]:[1050]:[32,51.51.51.51]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
* i  [3]:[1060]:[32,60.60.60.60]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
* i  [3]:[1060]:[32,76.76.76.76]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>   [3]:[1067]:[32,66.66.66.66]
                      66.66.66.66           0        100       32768  i  ----------
VXLAN

RD[2.3.4.5:1]
*>i  [2]:[0]:[1050]:[48,0000:0e8d:561a]:[0]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
*>i  [2]:[0]:[1050]:[48,0000:0e8d:561a]:[32,10.12.32.11]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
*>i  [2]:[0]:[1050]:[48,0000:3333:1050]:[32,10.12.32.10]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
*>i  [2]:[0]:[1050]:[48,3c2c:99d6:167a]:[32,10.12.32.1]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
*>i  [2]:[0]:[1050]:[48,3c2c:99d6:167a]:[128,2401::1]:[1050]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN
*>i  [3]:[1050]:[32,51.51.51.51]
                     51.51.51.51          0        100       0    i  51.51.51.51     VXLAN

RD[2.3.4.5:2]
*>i  [1]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i
[2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:2222:1060]:[32,10.240.38.10]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i  [2]:[0]:[1060]:[48,0000:2222:3333]:[32,10.240.38.1]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i  [2]:[0]:[1060]:[48,0000:2222:3333]:[128,1601::1]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i  [3]:[1060]:[32,60.60.60.60]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN

RD[2.3.4.6:2]
*>i  [1]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
```

```
*>i  [2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:0e8d:5619]:[0]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i
[2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:2222:1060]:[32,10.240.38.10]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [2]:[0]:[1060]:[48,0000:2222:3333]:[32,10.240.38.1]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [2]:[0]:[1060]:[48,0000:2222:3333]:[128,1601::1]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [3]:[1060]:[32,76.76.76.76]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN

RD[51.51.51.51:1050]
*>i  [5]:[0]:[10502]:[24]:[10.12.32.0]:[0.0.0.0]:[10502]
                     51.51.51.51          0        100       0    ?  51.51.51.51     VXLAN

RD[60.60.60.60:1]
*>i  [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i  [4]:[00:00:00:44:44:55:55:00:00:00]:[32,60.60.60.60]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN

RD[60.60.60.60:1040]
*>i  [5]:[0]:[10402]:[24]:[10.240.38.0]:[0.0.0.0]:[10402]
                     60.60.60.60          0        100       0    ?  60.60.60.60     VXLAN

RD[76.76.76.76:1]
*>i  [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [4]:[00:00:00:44:44:55:55:00:00:00]:[32,76.76.76.76]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN

RD[76.76.76.76:1040]
*>i  [5]:[0]:[10402]:[24]:[10.240.38.0]:[0.0.0.0]:[10402]
                     76.76.76.76          0        100       0    ?  76.76.76.76     VXLAN

Total number of prefixes 41
BorderVTEP1#

BorderVTEP1#show bgp l2vpn evpn prefix-route

RD[51.51.51.51:1050]
ESI                  Eth-Tag Prefix-Length    IP-Address                            GW-
IPAddress                            L3VNID    Nexthop         Encap     Router-Mac
0            10502  24          10.12.32.0                0.0.0.0
10502     51.51.51.51     VXLAN     3c2c:99d6:167a

RD[60.60.60.60:1040]
ESI                  Eth-Tag Prefix-Length    IP-Address                            GW-
IPAddress                            L3VNID    Nexthop         Encap     Router-Mac
```

```
0            10402  24          10.240.38.0                 0.0.0.0
10402     60.60.60.60     VXLAN     3c2c:99d1:117a

RD[76.76.76.76:1040]
ESI                  Eth-Tag Prefix-Length    IP-Address                              GW-
IPAddress                            L3VNID    Nexthop         Encap     Router-Mac
0            10402  24          10.240.38.0                 0.0.0.0
10402     76.76.76.76     VXLAN     3c2c:99de:1e7a
```

VTEP3

```
======


VTEP3#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
1060     ----         L2       NW   ----      ----                         ---- ----
76.76.76.76      60.60.60.60
1060     ----         --       AC   po1000    00:00:00:44:44:55:55:00:00:00 1060 NON-DF
----             ----
10402    ----         L3       NW   ----      ----                         ---- ----
76.76.76.76      66.66.66.66

Total number of entries are 3
VTEP3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
76.76.76.76      66.66.66.66      Installed        00:24:35       00:24:35
76.76.76.76      60.60.60.60      Installed        00:54:40       00:54:40

Total number of entries are 2
VTEP3#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
```

```
O            10.10.10.0/31 [110/2] via 10.10.24.1, xe27, 00:54:56
O            10.10.12.0/31 [110/2] via 10.10.24.1, xe27, 00:54:56
O            10.10.14.0/31 [110/2] via 10.10.24.1, xe27, 00:25:31
C            10.10.24.0/31 is directly connected, xe27, 00:55:37
O            51.51.51.51/32 [110/3] via 10.10.24.1, xe27, 00:54:47
O            60.60.60.60/32 [110/3] via 10.10.24.1, xe27, 00:54:45
O            66.66.66.66/32 [110/3] via 10.10.24.1, xe27, 00:24:46
C            76.76.76.76/32 is directly connected, lo, 00:55:38
C            127.0.0.0/8 is directly connected, lo, 00:55:39
IP Route Table for VRF "management"
C            10.12.20.0/24 is directly connected, eth0, 00:55:10
C            127.0.0.0/8 is directly connected, lo.management, 00:55:39
IP Route Table for VRF "SMS"
Gateway of last resort is 66.66.66.66 to network 0.0.0.0

B*           0.0.0.0/0 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:23:33
B            10.10.18.0/24 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:24:41
B            10.10.20.0/24 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:06:03
C            10.240.38.0/24 is directly connected, irb1060, 00:55:38
B            66.66.66.66/32 [0/0] is directly connected, tunvxlan2, 00:24:41
C            127.0.0.0/8 is directly connected, lo.SMS, 00:55:39
VTEP3#        show bgp l2vpn evpn summary
BGP router identifier 76.76.76.76, local AS number 500
BGP table version is 8
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
60.60.60.60             4   500  144       140      7     0     0  00:54:55
8      2      3      1      1      1
66.66.66.66             4   500  127       125      7     0     0  00:24:51
12      0      2      1      0      9

Total number of neighbors 2

Total number of Established sessions 2
VTEP3#        show bgp l2vpn evpn
BGP table version is 8, local router ID is 76.76.76.76
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
```

```
   Network          Next Hop            Metric    LocPrf      Weight     Path  Peer
Encap

RD[2.2.4.4:4]
*>i  [2]:[0]:[1067]:[48,3c2c:991c:dc7a]:[32,10.10.18.1]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
*>i  [2]:[0]:[1067]:[48,a82b:b5cf:f806]:[32,10.10.18.2]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
*>i  [3]:[1067]:[32,66.66.66.66]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN

RD[2.3.4.5:2]
*>i  [1]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i
[2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:2222:1060]:[32,10.240.38.10]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i  [2]:[0]:[1060]:[48,0000:2222:3333]:[32,10.240.38.1]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i  [2]:[0]:[1060]:[48,0000:2222:3333]:[128,1601::1]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i  [3]:[1060]:[32,60.60.60.60]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN

RD[2.3.4.6:2] VRF[RED]:
*>   [1]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[1060]
                      76.76.76.76           0        100       32768  i  ----------
VXLAN
* i                   60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
* i  [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>   [2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:0e8d:5619]:[0]:[1060]
                      76.76.76.76           0        100       32768  i  ----------
VXLAN
*>
[2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:2222:1060]:[32,10.240.38.10]:[1060]
                      76.76.76.76           0        100       32768  i  ----------
VXLAN
* i                   60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
* i  [2]:[0]:[1060]:[48,0000:2222:3333]:[32,10.240.38.1]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>                    76.76.76.76           0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[1060]:[48,0000:2222:3333]:[128,1601::1]:[1060]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>                    76.76.76.76           0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[1067]:[48,3c2c:991c:dc7a]:[32,10.10.18.1]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
* i  [2]:[0]:[1067]:[48,a82b:b5cf:f806]:[32,10.10.18.2]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
```

```
* i  [3]:[1060]:[32,60.60.60.60]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>   [3]:[1060]:[32,76.76.76.76]
                      76.76.76.76          0        100       32768  i  ----------
VXLAN
* i  [3]:[1067]:[32,66.66.66.66]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN

RD[4.5.6.8:6]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

RD[60.60.60.60:1]
*>i  [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>i  [4]:[00:00:00:44:44:55:55:00:00:00]:[32,60.60.60.60]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN

RD[60.60.60.60:1040]
*>i  [5]:[0]:[10402]:[24]:[10.240.38.0]:[0.0.0.0]:[10402]
                     60.60.60.60          0        100       0    ?  60.60.60.60     VXLAN

RD[66.66.66.66:1050]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

RD[66.66.66.66:1060]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

RD[76.76.76.76:1] VRF[evpn-gvrf-1]:
*>   [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                      76.76.76.76          0        100       32768  i  ----------
VXLAN
```

```
* i  [4]:[00:00:00:44:44:55:55:00:00:00]:[32,60.60.60.60]
                     60.60.60.60          0        100       0    i  60.60.60.60     VXLAN
*>   [4]:[00:00:00:44:44:55:55:00:00:00]:[32,76.76.76.76]
                      76.76.76.76           0        100       32768  i  ----------
VXLAN

Total number of prefixes 34
VTEP3#
VTEP3#show  bgp l2vpn evpn prefix-route

RD[4.5.6.8:6]
ESI                   Eth-Tag Prefix-Length    IP-Address                            GW-
IPAddress                              L3VNID    Nexthop         Encap     Router-Mac
0              500    0           0.0.0.0                      0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24          10.10.18.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24          10.10.20.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a

RD[60.60.60.60:1040]
ESI                   Eth-Tag Prefix-Length    IP-Address                            GW-
IPAddress                              L3VNID    Nexthop         Encap     Router-Mac
0             10402  24          10.240.38.0                   0.0.0.0
10402      60.60.60.60     VXLAN     3c2c:99d1:117a

RD[66.66.66.66:1050]
ESI                   Eth-Tag Prefix-Length    IP-Address                            GW-
IPAddress                              L3VNID    Nexthop         Encap     Router-Mac
0              500    0           0.0.0.0                      0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24          10.10.18.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24          10.10.20.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a

RD[66.66.66.66:1060]
ESI                   Eth-Tag Prefix-Length    IP-Address                            GW-
IPAddress                              L3VNID    Nexthop         Encap     Router-Mac
0              500    0           0.0.0.0                      0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24          10.10.18.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24          10.10.20.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
```

VTEP2

```
======

VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
```

```
          AC - Access Port
         (u) - Untagged

VNID    VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________________
1060    ----         L2       NW   ----      ----                           ---- ----
60.60.60.60      76.76.76.76
1060    ----         --       AC   po1000    00:00:00:44:44:55:55:00:00:00 1060 DF
----             ----
10402   ----         L3       NW   ----      ----                           ---- ----
60.60.60.60      66.66.66.66

Total number of entries are 3
VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
60.60.60.60      66.66.66.66      Installed        00:26:50       00:26:50
60.60.60.60      76.76.76.76      Installed        00:56:51       00:56:51

Total number of entries are 2
VTEP2#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            10.10.10.0/31 [110/2] via 10.10.12.0, xe25, 00:57:13
C            10.10.12.0/31 is directly connected, xe25, 00:57:57
O            10.10.14.0/31 [110/2] via 10.10.12.0, xe25, 00:27:47
O            10.10.16.0/31 [110/2] via 10.10.12.0, xe25, 00:57:13
O            10.10.24.0/31 [110/2] via 10.10.12.0, xe25, 00:57:13
O            51.51.51.51/32 [110/3] via 10.10.12.0, xe25, 00:57:03
C            60.60.60.60/32 is directly connected, lo, 00:57:59
O            66.66.66.66/32 [110/3] via 10.10.12.0, xe25, 00:27:02
O            76.76.76.76/32 [110/3] via 10.10.12.0, xe25, 00:57:13
C            127.0.0.0/8 is directly connected, lo, 00:58:00
IP Route Table for VRF "management"
C            10.12.20.0/24 is directly connected, eth0, 00:57:29
C            127.0.0.0/8 is directly connected, lo.management, 00:58:00
IP Route Table for VRF "SMS"
Gateway of last resort is 66.66.66.66 to network 0.0.0.0

B*           0.0.0.0/0 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:25:49
```

```
B            10.10.18.0/24 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:26:58
B            10.10.20.0/24 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:08:19
C            10.240.38.0/24 is directly connected, irb1060, 00:57:58
B            66.66.66.66/32 [0/0] is directly connected, tunvxlan2, 00:26:58
C            127.0.0.0/8 is directly connected, lo.SMS, 00:58:00

VTEP2#show bgp l2vpn evpn sum
BGP router identifier 60.60.60.60, local AS number 500
BGP table version is 12
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
66.66.66.66             4   500  135       133      11      0      0  00:27:29
12      0      2      1      0      9
76.76.76.76             4   500  146       150      11      0      0  00:57:30
9      2      4      1      1      1

Total number of neighbors 2

Total number of Established sessions 2
VTEP2#show bgp l2vpn evpn
BGP table version is 12, local router ID is 60.60.60.60
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[2.2.4.4:4]
*>i  [2]:[0]:[1067]:[48,3c2c:991c:dc7a]:[32,10.10.18.1]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
*>i  [2]:[0]:[1067]:[48,a82b:b5cf:f806]:[32,10.10.18.2]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
*>i  [3]:[1067]:[32,66.66.66.66]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN

RD[2.3.4.5:2] VRF[RED]:
* i  [1]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[1060]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
```

```
*>                    60.60.60.60          0        100       32768  i  ----------
VXLAN
* i  [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i  [2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:0e8d:5619]:[0]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i
[2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:2222:1060]:[32,10.240.38.10]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>                    60.60.60.60          0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[1060]:[48,0000:2222:3333]:[32,10.240.38.1]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>                    60.60.60.60          0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[1060]:[48,0000:2222:3333]:[128,1601::1]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>                    60.60.60.60          0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[1067]:[48,3c2c:991c:dc7a]:[32,10.10.18.1]:[1067]
                      66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
* i  [2]:[0]:[1067]:[48,a82b:b5cf:f806]:[32,10.10.18.2]:[1067]
                      66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
*>   [3]:[1060]:[32,60.60.60.60]
                       60.60.60.60         0        100       32768  i  ----------
VXLAN
* i  [3]:[1060]:[32,76.76.76.76]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
* i  [3]:[1067]:[32,66.66.66.66]
                      66.66.66.66          0        100       0    i  66.66.66.66     VXLAN

RD[2.3.4.6:2]
*>i  [1]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:0e8d:5619]:[0]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i
[2]:[00:00:00:44:44:55:55:00:00:00]:[1060]:[48,0000:2222:1060]:[32,10.240.38.10]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [2]:[0]:[1060]:[48,0000:2222:3333]:[32,10.240.38.1]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [2]:[0]:[1060]:[48,0000:2222:3333]:[128,1601::1]:[1060]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [3]:[1060]:[32,76.76.76.76]
                      76.76.76.76          0        100       0    i  76.76.76.76     VXLAN

RD[4.5.6.8:6]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                      66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                      66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
```

```
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

RD[60.60.60.60:1] VRF[evpn-gvrf-1]:
*>   [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                      60.60.60.60          0        100       32768  i  ----------
VXLAN
*>   [4]:[00:00:00:44:44:55:55:00:00:00]:[32,60.60.60.60]
                      60.60.60.60          0        100       32768  i  ----------
VXLAN
* i  [4]:[00:00:00:44:44:55:55:00:00:00]:[32,76.76.76.76]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN

RD[66.66.66.66:1050]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

RD[66.66.66.66:1060]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

RD[76.76.76.76:1]
*>i  [1]:[00:00:00:44:44:55:55:00:00:00]:[4294967295]:[0]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN
*>i  [4]:[00:00:00:44:44:55:55:00:00:00]:[32,76.76.76.76]
                     76.76.76.76          0        100       0    i  76.76.76.76     VXLAN

RD[76.76.76.76:1040]
*>i  [5]:[0]:[10402]:[24]:[10.240.38.0]:[0.0.0.0]:[10402]
                     76.76.76.76          0        100       0    ?  76.76.76.76     VXLAN

Total number of prefixes 35
VTEP2#
VTEP22#show  bgp l2vpn evpn prefix-route

RD[4.5.6.8:6]
ESI                  Eth-Tag Prefix-Length    IP-Address                          GW-
IPAddress                          L3VNID    Nexthop         Encap     Router-Mac
```

```
0             500   0          0.0.0.0                    0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a
0             500   24         10.10.18.0                 0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a
0             500   24         10.10.20.0                 0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a

RD[66.66.66.66:1050]
ESI                 Eth-Tag Prefix-Length    IP-Address                          GW-
IPAddress                           L3VNID    Nexthop         Encap     Router-Mac
0             500   0          0.0.0.0                    0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a
0             500   24         10.10.18.0                 0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a
0             500   24         10.10.20.0                 0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a

RD[66.66.66.66:1060]
ESI                 Eth-Tag Prefix-Length    IP-Address                          GW-
IPAddress                           L3VNID    Nexthop         Encap     Router-Mac
0             500   0          0.0.0.0                    0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a
0             500   24         10.10.18.0                 0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a
0             500   24         10.10.20.0                 0.0.0.0
500       66.66.66.66     VXLAN     3c2c:991c:dc7a

RD[76.76.76.76:1040]
ESI                 Eth-Tag Prefix-Length    IP-Address                          GW-
IPAddress                           L3VNID    Nexthop         Encap     Router-Mac
0           10402   24         10.240.38.0                0.0.0.0
10402     76.76.76.76     VXLAN     3c2c:99de:1e7a
VTEP2#
```

VTEP1

```
======
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID    VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
1050    ----         --       AC   ce50     --- Single Homed Port ---     1050 ----
----             ----
10502   ----         L3       NW   ----     ----                          ---- ----
51.51.51.51      66.66.66.66

Total number of entries are 2
```

```
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
51.51.51.51      66.66.66.66      Installed        00:28:13       00:28:13

Total number of entries are 1
VTEP1#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/31 is directly connected, xe40, 00:59:46
O            10.10.12.0/31 [110/2] via 10.10.10.0, xe40, 00:59:01
O            10.10.14.0/31 [110/2] via 10.10.10.0, xe40, 00:29:45
O            10.10.24.0/31 [110/2] via 10.10.10.0, xe40, 00:59:01
C            51.51.51.51/32 is directly connected, lo, 00:59:47
O            60.60.60.60/32 [110/3] via 10.10.10.0, xe40, 00:59:01
O            66.66.66.66/32 [110/3] via 10.10.10.0, xe40, 00:29:00
O            76.76.76.76/32 [110/3] via 10.10.10.0, xe40, 00:59:01
C            127.0.0.0/8 is directly connected, lo, 00:59:49
IP Route Table for VRF "management"
C            10.12.20.0/24 is directly connected, eth0, 00:59:22
C            127.0.0.0/8 is directly connected, lo.management, 00:59:49
IP Route Table for VRF "FAX"
Gateway of last resort is 66.66.66.66 to network 0.0.0.0

B*           0.0.0.0/0 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:27:47
B            10.10.18.0/24 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:28:55
B            10.10.20.0/24 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:10:17
C            10.12.32.0/24 is directly connected, irb1050, 00:59:47
B            66.66.66.66/32 [0/0] is directly connected, tunvxlan2, 00:28:55
C            127.0.0.0/8 is directly connected, lo.FAX, 00:59:48
VTEP1#               show  bgp l2vpn evpn summary
BGP router identifier 51.51.51.51, local AS number 500
BGP table version is 9
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
```

```
66.66.66.66             4   500  138       132      8      0      0  00:29:07
12      0      2      1      0      9

Total number of neighbors 1

Total number of Established sessions 1
VTEP1#                show  bgp l2vpn evpn
BGP table version is 9, local router ID is 51.51.51.51
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[2.2.4.4:4]
*>i  [2]:[0]:[1067]:[48,3c2c:991c:dc7a]:[32,10.10.18.1]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
*>i  [2]:[0]:[1067]:[48,a82b:b5cf:f806]:[32,10.10.18.2]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
*>i  [3]:[1067]:[32,66.66.66.66]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN

RD[2.3.4.5:1] VRF[RED]:
*>   [2]:[0]:[1050]:[48,0000:0e8d:561a]:[0]:[1050]
                      51.51.51.51           0        100       32768  i  ----------
VXLAN
*>   [2]:[0]:[1050]:[48,0000:0e8d:561a]:[32,10.12.32.11]:[1050]
                      51.51.51.51           0        100       32768  i  ----------
VXLAN
*>   [2]:[0]:[1050]:[48,0000:3333:1050]:[32,10.12.32.10]:[1050]
                      51.51.51.51           0        100       32768  i  ----------
VXLAN
*>   [2]:[0]:[1050]:[48,3c2c:99d6:167a]:[32,10.12.32.1]:[1050]
                      51.51.51.51           0        100       32768  i  ----------
VXLAN
*>   [2]:[0]:[1050]:[48,3c2c:99d6:167a]:[128,2401::1]:[1050]
                      51.51.51.51           0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[1067]:[48,3c2c:991c:dc7a]:[32,10.10.18.1]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
* i  [2]:[0]:[1067]:[48,a82b:b5cf:f806]:[32,10.10.18.2]:[1067]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN
*>   [3]:[1050]:[32,51.51.51.51]
```

```
                      51.51.51.51          0        100       32768  i  ----------
VXLAN
* i  [3]:[1067]:[32,66.66.66.66]
                     66.66.66.66          0        100       0    i  66.66.66.66     VXLAN

RD[4.5.6.8:6]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

RD[66.66.66.66:1050]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

RD[66.66.66.66:1060]
*>i  [5]:[0]:[500]:[0]:[0.0.0.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.18.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0    ?  66.66.66.66     VXLAN
*>i  [5]:[0]:[500]:[24]:[10.10.20.0]:[0.0.0.0]:[500]
                     66.66.66.66          0        100       0   64603 i  66.66.66.66
VXLAN

Total number of prefixes 21
VTEP1#


VTEP1#show ip route vrf FAX
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "FAX"
Gateway of last resort is 66.66.66.66 to network 0.0.0.0
```

```
B*           0.0.0.0/0 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:29:26
B            10.10.18.0/24 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:30:34
B            10.10.20.0/24 [200/0] via 66.66.66.66 (recursive is directly connected,
tunvxlan2), 00:11:56
C            10.12.32.0/24 is directly connected, irb1050, 01:01:26
B            66.66.66.66/32 [0/0] is directly connected, tunvxlan2, 00:30:34
C            127.0.0.0/8 is directly connected, lo.FAX, 01:01:27


VTEP1#show bgp l2vpn evpn prefix-route

RD[4.5.6.8:6]
ESI                  Eth-Tag Prefix-Length    IP-Address                          GW-
IPAddress                            L3VNID    Nexthop         Encap     Router-Mac
0              500    0          0.0.0.0                      0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24         10.10.18.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24         10.10.20.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a

RD[66.66.66.66:1050]
ESI                  Eth-Tag Prefix-Length    IP-Address                          GW-
IPAddress                            L3VNID    Nexthop         Encap     Router-Mac
0              500    0          0.0.0.0                      0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24         10.10.18.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24         10.10.20.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a

RD[66.66.66.66:1060]
ESI                  Eth-Tag Prefix-Length    IP-Address                          GW-
IPAddress                            L3VNID    Nexthop         Encap     Router-Mac
0              500    0          0.0.0.0                      0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24         10.10.18.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
0              500    24         10.10.20.0                    0.0.0.0
500        66.66.66.66     VXLAN     3c2c:991c:dc7a
```

Ping to 10.10.20.1 network which is advertised by Firewall from VTEP1 FAX vrf

```
VTEP1#       ping 10.10.20.1 vrf FAX
Press CTRL+C to exit
PING 10.10.20.1 (10.10.20.1) 56(84) bytes of data.
64 bytes from 10.10.20.1: icmp_seq=1 ttl=63 time=0.446 ms
64 bytes from 10.10.20.1: icmp_seq=2 ttl=63 time=0.413 ms
64 bytes from 10.10.20.1: icmp_seq=3 ttl=63 time=0.373 ms

--- 10.10.20.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 73ms
```

```
rtt min/avg/max/mdev = 0.373/0.410/0.446/0.037 ms
VTEP1#
```

CHAPTER 5 DHCP Relay Over IRB Interface

This chapter contains configurations for DHCP relay over IRB interface.

Overview

Dynamic Host Configuration Protocol (DHCP) is a protocol that allows a DHCP server to dynamically allocate IP addresses to DHCP clients. The DHCP relay agent forwards DHCP messages between DHCP clients and DHCP servers when they are on different networks.

For DHCP relay to function, uplink interface (server facing) and downlink interface (client facing) are to be configured along with DHCP server address. These configurations are done in the interface mode.

In the IRB deployment, IRB interface acts as gateway and hence it might need to provide the service of DHCP for the hosts connected to IRB network per L2 VPN. Since the DHCP server will not be present in the VTEP, it can forward the DHCP requests to the DHCP server acting as relay agent.

Topology

The procedures in this section use the topology in Figure 5-3.

Figure 5-3: DHCP Relay over IRB

ROUTER-1

`#configure terminal`	Enter Configure mode.
`OcNOS(config)#interface lo`	Enter Interface mode for loopback.
`OcNOS(config-if)#ip address 1.1.1.1/32 secondary`	Assign secondary IP address.
`OcNOS(config-if)#exit`	Exit Interface mode and return to Configure mode.
`OcNOS(config)#nvo vxlan enable`	Enable VXLAN
`OcNOS(config)#nvo vxlan irb`	Enable VXLAN IRB
`OcNOS(config)#ip vrf vrf1`	Create routing/forwarding instance with VRF1 name and enter into VRF mode
`OcNOS(config-vrf)#rd 200:1`	Assign RD value
`OcNOS(config-vrf)#route-target both 200:1`	Assign route target value
`OcNOS(config-vrf)#ip dhcp relay address 40.40.40.1`	The relay address configured should be server interface address connected to DUT machine
`OcNOS(config-vrf)#ip dhcp relay uplink evpn`	Configure the uplink interface as L3 VNI interface for specific VRF
`OcNOS(config-vrf)#l3vni 45001`	Configure L3VNI as 45001 for VRF1
`OcNOS(config-vrf)#exit`	Exit IP VRF mode

`OcNOS(config)#mac vrf vrfred`	Create MAC VRF instance with vrfred name and enter into VRF mode
`OcNOS(config-vrf)#rd 1.1.1.1:1`	Assign RD value
`OcNOS(config-vrf)#route-target both 1.1.1.1:1`	Assign route target value
`OcNOS(config-vrf)#exit`	Exit MAC VRF mode
`OcNOS(config)#interface irb 1`	Configure IRB interface
`OcNOS(config-irb-if)#ip vrf forwarding vrf1`	Configure IP VRF forwarding
`OcNOS(config-irb-if)#ip address 11.1.1.1/24`	Assign IP address on IRB interface.
`OcNOS(config-irb-if)#ip dhcp relay`	Relay should be configured on the interface connecting to the relay
`OcNOS(config-irb-if)#exit`	Exit IRB interface mode
`OcNOS(config)#interface irb 2`	Configure irb interface
`OcNOS(config-irb-if)#ip vrf forwarding vrf1`	Configure IP VRF forwarding
`OcNOS(config-irb-if)#ip address 70.70.70.1/24`	Assign IP address on IRB interface.
`OcNOS(config-irb-if)#exit`	Exit IRB interface mode
`OcNOS(config)#interface ce49`	Enter Interface mode for ce49.
`OcNOS(config-if)#ip address 10.1.1.2/24`	Assign IP address on ce49 interface.
`OcNOS(config-if)#exit`	Exit Interface mode and return to Configure mode.
`OcNOS(config)#interface xe5`	Enter Interface mode for xe5.
`OcNOS(config-if)#switchport`	Configure interface as L2 interface
`OcNOS(config-if)#exit`	Exit Interface mode and return to Configure mode.
`OcNOS(config)#router ospf`	Enter the Router OSPF mode
`OcNOS(config-router)#network 1.1.1.1/32 area 0.0.0.0`	Advertise loopback address in OSPF
`OcNOS(config-router)#network 10.1.1.0/24 area 0.0.0.0`	Advertise network address in OSPF
`OcNOS(config-router)#exit`	Exit from Router OSPF mode and enter into config mode
`OcNOS(config)#router bgp 1`	Enter into BGP router mode
`OcNOS(config-router)#neighbor 2.2.2.2 remote-as 1`	Specify a VTEP2 loopback IP address and remote-as defined
`OcNOS(config-router)#neighbor 2.2.2.2 update-source 1.1.1.1`	Configure update as loopback for VTEP2
`OcNOS(config-router)#address-family l2vpn evpn`	Enter into L2VPN EVPN address family mode
`OcNOS(config-router-af)#neighbor 2.2.2.2 activate`	Activate neighbor in L2VPN mode
`OcNOS(config-router-af)#exit-address-family`	Exit from Address family mode
`OcNOS(config-router)#address-family ipv4 vrf vrf1`	Enter into address-family mode for VRF1
`OcNOS(config-router-af)#redistribute connected`	Configure Redistribute connected
`OcNOS(config-router-af)#exit-address-family`	Exit from Address family mode
`OcNOS(config-router)#exit`	Exit from router BGP mode and enter into config mode

`OcNOS(config)#nvo vxlan vtep-ip-global 1.1.1.1`	Configure Source VTEP-IP-global configuration. Use loopback IP address
`OcNOS(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled`	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
`OcNOS(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrfred`	Assign VRF for EVPN-BGP to carry EVPN route
`OcNOS(config-nvo)#evpn irb1`	Configure IRB1 under VXLAN ID 10
`OcNOS(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.
`OcNOS(config)#nvo vxlan id 30 ingress-replication inner-vid-disabled`	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
`OcNOS(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrfred`	Assign VRF for EVPN-BGP to carry EVPN route
`OcNOS(config-nvo)#evpn irb2`	Configure IRB2 under VXLAN ID 30
`OcNOS(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.
`OcNOS(config)#nvo vxlan access-if port-vlan xe5 2`	Enable port-VLAN mapping i.e. access port to outer-VLAN (SVLAN) - Multihomed access port
`OcNOS(config-nvo-acc-if)#map vnid 10`	Map VXLAN Identified to access-port for VXLAN
`OcNOS(config-nvo-acc-if)#exit`	Exit from VXLAN access-interface mode and enter into configuration mode
`OcNOS(config)#commit`	Commit the candidate configuration to the running configuration
`OcNOS(config)#exit`	Exit from configuration mode

ROUTER-2

`OcNOS#configure terminal`	Enter Configure mode.
`OcNOS(config)#interface ce0`	Enter Interface mode for ce0.
`OcNOS(config-if)#ip address 10.1.1.1/24`	Assign IP address on ce0 interface.
`OcNOS(config-if)#exit`	Exit Interface mode and return to Configure mode.
`OcNOS(config)#interface xe13`	Enter Interface mode for xe13.
`OcNOS(config-if)#ip address 30.1.1.1/24`	Assign IP address on xe13 interface.
`OcNOS(config-if)#exit`	Exit Interface mode and return to Configure mode.
`OcNOS(config)#router ospf`	Enter the Router OSPF mode
`OcNOS(config-router)#network 10.1.1.0/24 area 0.0.0.0`	Advertise network address in OSPF
`OcNOS(config-router)#network 30.1.1.10/24 area 0.0.0.0`	Advertise network address in OSPF
`OcNOS(config-router)#exit`	Exit from Router OSPF mode and enter into config mode
`OcNOS(config)#commit`	commit the candidate configuration to the running configuration
`OcNOS(config)#exit`	Exit from configuration mode

ROUTER-3

`OcNOS#configure terminal`	Enter Configure mode.
`OcNOS(config)#interface lo`	Enter Interface mode for loopback.
`OcNOS(config-if)#ip address 2.2.2.2/32 secondary`	Assign secondary IP address.
`OcNOS(config-if)#exit`	Exit Interface mode and return to Configure mode.
`OcNOS(config)#nvo vxlan enable`	Enable VXLAN
`OcNOS(config)#nvo vxlan irb`	Enable VXLAN IRB
`OcNOS(config)#ip vrf vrf1`	Create routing/forwarding instance with VRF1 name and enter into VRF mode
`OcNOS(config-vrf)#rd 300:1`	Assign RD value
`OcNOS(config-vrf)#route-target both 200:1`	Assign route target value
`OcNOS(config-vrf)#ip dhcp relay uplink evpn`	Configure the uplink interface as L3 VNI interface for specific VRF
`OcNOS(config-vrf)#l3vni 45001`	Configure L3VNI as 45001 for VRF1
`OcNOS(config-vrf)#exit`	Exit IP VRF mode
`OcNOS(config)#mac vrf vrfred`	Create MAC VRF instance with vrfred name and enter into VRF mode
`OcNOS(config-vrf)#rd 2.2.2.1:1`	Assign RD value
`OcNOS(config-vrf)#route-target both 1.1.1.1:1`	Assign route target value
`OcNOS(config-vrf)#exit`	Exit MAC VRF mode
`OcNOS(config)#interface irb 2`	Configure IRB interface
`OcNOS(config-irb-if)#ip vrf forwarding vrf1`	Configure IP VRF forwarding
`OcNOS(config-irb-if)#ip address 40.40.40.2/24`	Assign IP address on IRB interface.
`OcNOS(config-irb-if)#exit`	Exit IRB interface mode
`OcNOS(config)#interface xe13`	Enter Interface mode for xe13.
`OcNOS(config-if)#ip address 30.1.1.2/24`	Assign IP address on xe13 interface.
`OcNOS(config-if)#exit`	Exit Interface mode and return to Configure mode.
`OcNOS(config)#interface xe19`	Enter Interface mode for xe19.
`OcNOS(config-if)#switchport`	Configure interface as L2 interface
`OcNOS(config-if)#exit`	Exit Interface mode and return to Configure mode.
`OcNOS(config)#router ospf`	Enter the Router OSPF mode
`OcNOS(config-router)#network 2.2.2.2/32 area 0.0.0.0`	Advertise loopback address in OSPF
`OcNOS(config-router)#network 30.1.1.0/24 area 0.0.0.0`	Advertise network address in OSPF
`OcNOS(config-router)#network 40.1.1.0/24 area 0.0.0.0`	Advertise network address in OSPF
`OcNOS(config-router)#exit`	Exit from Router OSPF mode and enter into config mode
`OcNOS(config)#router bgp 1`	Enter into BGP router mode
`OcNOS(config-router)#neighbor 1.1.1.1 remote-as 1`	Specify a VTEP1 loopback IP address and remote-as defined

`OcNOS(config-router)#neighbor 1.1.1.1 update-source 2.2.2.2`	Configure update as loopback for VTEP1
`OcNOS(config-router)#address-family l2vpn evpn`	Enter into L2VPN EVPN address family mode
`OcNOS(config-router-af)#neighbor 1.1.1.1 activate`	Activate neighbor in L2VPN mode
`OcNOS(config-router-af)#exit-address-family`	Exit from Address family mode
`OcNOS(config-router)#address-family ipv4 vrf vrf1`	Enter into address-family mode for VRF1
`OcNOS(config-router-af)#redistribute connected`	Configure Redistribute connected
`OcNOS(config-router-af)#exit-address-family`	Exit from Address family mode
`OcNOS(config-router)#exit`	Exit from router BGP mode and enter into config mode
`OcNOS(config)#nvo vxlan vtep-ip-global 2.2.2.2`	Configure Source VTEP-IP-global configuration. Use loopback IP address
`OcNOS(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled`	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
`OcNOS(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrfred`	Assign VRF for EVPN-BGP to carry EVPN route
`OcNOS(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.
`OcNOS(config)#nvo vxlan id 20 ingress-replication inner-vid-disabled`	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
`OcNOS(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrfred`	Assign VRF for EVPN-BGP to carry EVPN route
`OcNOS(config-nvo)#evpn irb2`	Configure IRB2 under VXLAN ID 20
`OcNOS(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.
`OcNOS(config)#nvo vxlan access-if port xe19`	Enable port mapping i.e. access port
`OcNOS(config-nvo-acc-if)#map vnid 20`	Map VXLAN Identified to access-port for VXLAN
`OcNOS(config-nvo-acc-if)#exit`	Exit from VXLAN access-interface mode and enter into configuration mode
`OcNOS(config)#commit`	Commit the candidate configuration to the running configuration
`OcNOS(config)#exit`	Exit from configuration mode

Validation

ROUTER-1

```
VTEP1#
!
nvo vxlan enable
!
nvo vxlan irb
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan id 10 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrfred
```

```
 evpn irb1
  !
 nvo vxlan id 30 ingress-replication inner-vid-disabled
   vxlan host-reachability-protocol evpn-bgp vrfred
   evpn irb2
!
nvo vxlan access-if port xe2
 map vnid 10
!
nvo vxlan access-if port-vlan xe5 2
 map vnid 10
!
VTEP1#show ip dhcp relay
DHCP relay service is Enabled.
VRF Name: vrf1
  Option 82: Disabled
  DHCP Servers configured: 40.40.40.1

  Interface                Uplink/Downlink
  ---------                -------------
  irb1                     Downlink
  evpn                     uplink
Incoming DHCPv4 packets which already contain relay agent option are FORWARDED
unchanged.
```

```
VTEP1#show nvo vxlan mac-table
================================================================================================================================================
==============
                                                    VXLAN MAC Entries
================================================================================================================================================
==============
VNID       Interface VlanId Vlan-RangeId Inner-VlanId Mac-Addr      VTEP-Ip/ESI                   Type           Status
AccessPortDesc
------------------------------------------------------------------------------------------------------------------------------------------------
------------

10         xe5       2      ----         ----         0000.2837.ddf5 1.1.1.1                     Dynamic Local   -------
-------
10         irb1                          b86a.97f9.85be 1.1.1.1                     Static Local    -------    -------
30         irb2                          b86a.97f9.85be 1.1.1.1                     Static Local    -------    -------

Total number of entries are : 3
```

```
VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID      Ip-Addr           Mac-Addr       Type         Age-Out    Retries-Left
------------------------------------------------------------------------------
30        70.70.70.1        b86a.97f9.85be Static Local   ----
10        11.1.1.1          b86a.97f9.85be Static Local   ----
10        11.1.1.30         0000.2837.ddf5 Dynamic Local  ----
Total number of entries are 3

VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
```

```
================================================================================
1.1.1.1          2.2.2.2          Installed        01:51:11       01:51:11

Total number of entries are 1
```

ROUTER-2

```
VTEP2#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan irb
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 10 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrfred
!
nvo vxlan id 20 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrfred
 evpn irb2
!
nvo vxlan access-if port xe19
 map vnid 20
!
!

VTEP2#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left

20       40.40.40.2       b86a.9761.ea3d Static      Local      ----
20       40.40.40.1       0002.a54f.1577 Dynamic     Local      ----
20       40.40.40.101     0000.2837.ddf3 Dynamic     Local      ----
10       11.1.1.1         b86a.97f9.85be Static      Remote     ----
10       11.1.1.30        0000.2837.ddf6 Dynamic     Remote     ----
Total number of entries are 5

VTEP2#show running-config dhcp
interface eth0
 ip address dhcp
!
!

ip vrf vrf1
 ip dhcp relay uplink evpn
```

CHAPTER 6 EVPN-VxLAN with IRB QoS

Overview

An EVPN-based Integrated Routing and Bridging solution used for forwarding of intra-subnets and inter-subnets traffic. Here QoS is applied on IRB solution for L3 packets.

VXLAN quality of service (QoS) provides differentiated service in VXLAN applications. A device implements mapping between QoS priorities in original packets, internal priorities (local precedence assigned by the device to differentiate service classes of packets), and priorities of encapsulated packets. In this way, the switch provides the differentiated QoS service based on original packets.

Topology

The procedures in this section use the topology in Figure 6-4

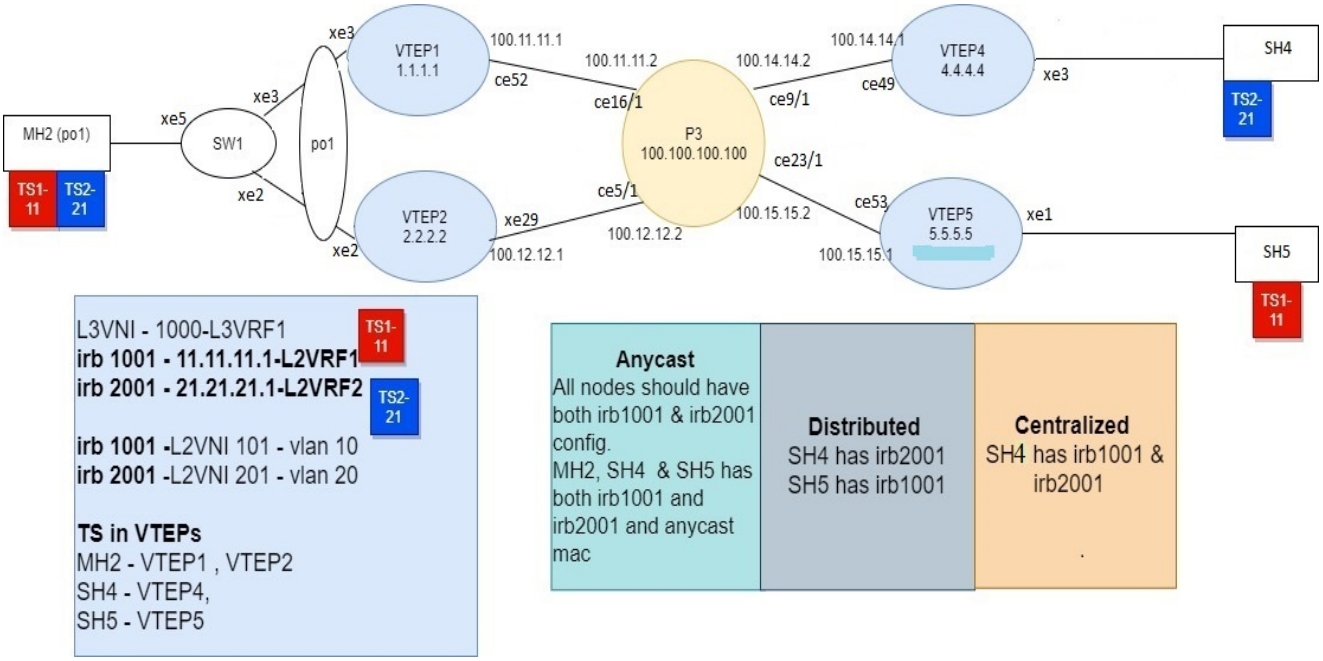

Figure 6-4: VxLAN_EVPN_IRB

Note: In the above topology TS1, TS2 are the tenant systems. The blue and red color denotes different subnets in the Tenant systems.

Base Configuration - L2 VXLAN

VTEP1

(Multi-homed group1) - Part of both Multi-homed with po1(MH2).

Generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#qos enable`	Enabling qos

Interface and loopback configuration:

`(config)#interface po1`	Enter Interface mode for po1 (MH2)
`(config-if)#switchport`	Make it L2 interface
`(config-if)# evpn multi-homed system-mac 0000.0000.2222`	Configure system mac as ESI value for Lag (po1) interface. VTEP1 and VTEP2 should have same ESI value
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface xe3`	Enter Interface mode for xe3
`(config-if)#channel-group 1 mode active`	Make it member port of po1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure loopback ip address as 1.1.1.1 for VTEP1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface ce52`	Enter Interface mode for network side port
`(config-if)#ip address 100.11.11.1/24`	Configure ip address as 100.11.11.1 on network side of Spine-P3
`(config-if)#exit`	Exit Interface mode and return to Configure mode.

OSPF configuration:

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 1.1.1.1`	Configure router-id as 1.1.1.1 (lo ip address)
`(config-router)#network 1.1.1.1/32 area 0.0.0.0`	Add 1.1.1.1 (lo ip address) network into area 0
`(config-router)#network 100.11.11.0/24 area 0.0.0.0`	Add 100.11.11.0(Spine-P3) network into area 0
`(config-router)#bfd all-interfaces`	Enabling bfd on all ospf interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.

BGP configuration:

`(config)#router bgp 5000`	Enter into Router BGP mode
`(config-router)#bgp router-id 1.1.1.1`	Configure router-id as 1.1.1.1 (lo ip address)
`(config-router)#neighbor 2.2.2.2 remote-as 5000`	Specify a VTEP2 loopback ip address and remote-as defined
`(config-router)#neighbor 2.2.2.2 update-source lo`	Configure update as loopback for VTEP2

(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP2
(config-router)#neighbor 4.4.4.4 remote-as 5000	Specify a VTEP4 loopback ip address and remote-as defined
(config-router)#neighbor 4.4.4.4 update-source lo	Configure update as loopback for VTEP4
(config-router)#neighbor 4.4.4.4 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP4
(config-router)#neighbor 5.5.5.5 remote-as 5000	Specify a VTEP5 loopback ip address and remote-as defined
(config-router)#neighbor 5.5.5.5 update-source lo	Configure update as loopback for VTEP5
(config-router)#neighbor 5.5.5.5 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP5
(config-router)#address-family l2vpn evpn	Enter into l2vpn EVPN address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2(VTEP2) into l2vpn evpn address family mode
(config-router-af)#neighbor 4.4.4.4 activate	Activate 3.3.3.3(VTEP4) into l2vpn evpn address family mode
(config-router-af)#neighbor 5.5.5.5 activate	Activate 5.5.5.5(VTEP5) into l2vpn evpn address family mode
(config-router-af)#exit-address-family	Exit from l2vpn address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode

L2 MAC VRF Configuration:

(config)#mac vrf L2VRF1	Create mac routing/forwarding instance with L2VRF1 name and enter into vrf mode
(config-vrf)#rd 1.1.1.1:11	Assign RD value
(config-vrf)#description MAC VRF RED	Give description to L2VRF1 as RED
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
(config-vrf)#exit	Exit from vrf mode
(config)#mac vrf L2VRF2	Create mac routing/forwarding instance with L2VRF2 name and enter into vrf mode
(config-vrf)#rd 1.1.1.1:21	Assign RD value
(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#description MAC VRF BLUE	Give description to L2VRF2 as BLUE
(config-vrf)#exit	Exit from vrf mode

L2 VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#evpn esi hold-time 90	Configure ESI hold time to allow tunnel to come up at the time of vxlan initialization before making the ESI up. It should be same on both VTEP1 and VTEP2
(config)#nvo vxlan vtep-ip-global 1.1.1.1	Configure Source vtep-ip-global configuration - Use loopback ip address

(config)#nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# vni-name VNI-101	Configure VNI name as VNI-101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# vni-name VNI-201	Configure VNI name as VNI-201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan access-if port-vlan po1 10	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 101	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1010 ip 11.11.11.51	Configure static mac-ip
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan po1 20	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 201	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1020 ip 21.21.21.51	Configure static mac-ip
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#exit	Exit from configuration mode

VTEP2

(Multi-homed group1) - Part of both Multi-homed with po1(MH1).

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#evpn vxlan multihoming enable	Enable Multihoming, save configs and reboot the board for multihoming to be effective
(config)#qos enable	Enabling qos

Interface and loopback configuration:

(config)#interface po1	Enter Interface mode for po1 (MH2)
(config-if)#switchport	Make it L2 interface
(config-if)# evpn multi-homed system-mac 0000.0000.2222	Configure system mac as ESI value for Lag (po1) interface. VTEP1 and VTEP2 should have same ESI value

(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe2	Enter Interface mode for xe2
(config-if)#channel-group 1 mode active	Make it member port of po1
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 2.2.2.2/32 secondary	Configure loopback ip address as 2.2.2.2 for VTEP2
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config) interface xe29	Enter into network side interface
(config-if)#ip address 100.12.12.1/24	Configure ip address as 100.12.12.1 on network side of Spine-P3
(config-if)#exit	Exit Interface mode and return to Configure mode.

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 2.2.2.2	Configure router-id as 2.2.2.2 (lo ip address)
(config-router)#network 2.2.2.2/32 area 0.0.0.0	Add 2.2.2.2 (lo ip address) network into area 0
(config-router)#network 100.12.12.0/24 area 0.0.0.0	Add 100.12.12.0(Spine-P3) network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.

BGP configuration:

(config)#router bgp 5000	Enter into Router BGP mode
(config-router)#bgp router-id 2.2.2.2	Configure router-id as 2.2.2.2 (lo ip address)
(config-router)#neighbor 1.1.1.1 remote-as 5000	Specify a VTEP1 loopback ip address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Configure update as loopback for VTEP1
(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP1
(config-router)#neighbor 4.4.4.4 remote-as 5000	Specify a VTEP4 loopback ip address and remote-as defined
(config-router)#neighbor 4.4.4.4 update-source lo	Configure update as loopback for VTEP4
(config-router)#neighbor 4.4.4.4 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP4
(config-router)#neighbor 5.5.5.5 remote-as 5000	Specify a VTEP5 loopback ip address and remote-as defined
(config-router)#neighbor 5.5.5.5 update-source lo	Configure update as loopback for VTEP5
(config-router)#neighbor 5.5.5.5 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP5
(config-router)#address-family l2vpn evpn	Enter into l2vpn EVPN address family mode

(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1(VTEP1) into l2vpn evpn address family mode
(config-router-af)#neighbor 4.4.4.4 activate	Activate 4.4.4.4(VTEP4) into l2vpn evpn address family mode
(config-router-af)#neighbor 5.5.5.5 activate	Activate 5.5.5.5(VTEP5) into l2vpn evpn address family mode
(config-router-af)#exit-address-family	Exit from l2vpn address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode

VRF Configuration:

(config)#mac vrf L2VRF1	Create mac routing/forwarding instance with L2VRF1 name and enter into vrf mode
(config-vrf)#rd 2.2.2.2:11	Assign RD value
(config-vrf)#description MAC VRF RED	Give description to L2VRF1 as RED
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
(config-vrf)#exit	Exit from vrf mode
(config)#mac vrf L2VRF2	Create mac routing/forwarding instance with L2VRF2 name and enter into vrf mode
(config-vrf)#rd 2.2.2.2:21	Assign RD value
(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#description MAC VRF BLUE	Give description to L2VRF2 as BLUE
(config-vrf)#exit	Exit from vrf mode

VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#evpn esi hold-time 90	Configure ESI hold time to allow tunnel to come up at the time of vxlan initialization before making the ESI up.It should be same on both VTEP1 and VTEP2
(config)#nvo vxlan vtep-ip-global 2.2.2.2	Configure Source vtep-ip-global configuration - Use loopback ip address
(config)#nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# vni-name VNI-101	Configure VNI name as VNI-101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# vni-name VNI-201	Configure VNI name as VNI-201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan access-if port-vlan po1 10	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port

(config-nvo-acc-if)#map vnid 101	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1010 ip 11.11.11.51	Configure static mac-ip
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#nvo vxlan access-if port-vlan po1 20	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) - Multihomed access port
(config-nvo-acc-if)#map vnid 201	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.2222.1020 ip 21.21.21.51	Configure static mac-ip
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#exit	Exit from configuration mode

VTEP4

Single Home -SH5.

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#evpn vxlan multihoming enable	Enable Multihoming, save configs and reboot the board for multihoming to be effective
(config)#qos enable	Enabling qos

Interface and loopback configuration:

(config)#interface xe3	Enter Interface mode for xe3
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface lo	Enter Interface mode for lo
(config-if)#ip address 4.4.4.4/32 secondary	Configure loopback ip address as 4.4.4.4 for VTEP4
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config) interface ce49	Enter interface towards network side
(config-if)#ip address 100.14.14.1/24	Configure ip address as 100.14.14.1 on network side of Spine-P3
(config-if)#exit	Exit Interface mode and return to Configure mode.

OSPF configuration:

(config)#router ospf 100	Enter into router OSPF mode
(config-router)#ospf router-id 4.4.4.4	Configure router-id as 4.4.4.4 (lo ip address)
(config-router)#network 4.4.4.4/32 area 0.0.0.0	Add 4.4.4.4 (lo ip address) network into area 0

(config-router)#network 100.14.14.0/24 area 0.0.0.0	Add 100.14.14.0(Spine-P3) network into area 0
(config-router)#bfd all-interfaces	Enabling bfd on all ospf interface for fast convergence
(config-if)#exit	Exit Interface mode and return to Configure mode.

BGP configuration:

(config)#router bgp 5000	Enter into Router BGP mode
(config-router)#bgp router-id 4.4.4.4	Configure router-id as 4.4.4.4 (lo ip address)
(config-router)#neighbor 1.1.1.1 remote-as 5000	Specify a VTEP1 loopback ip address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Configure update as loopback for VTEP1
(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP1
(config-router)#neighbor 2.2.2.2 remote-as 5000	Specify a VTEP2 loopback ip address and remote-as defined
(config-router)#neighbor 2.2.2.2 update-source lo	Configure update as loopback for VTEP2
(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP2
(config-router)#neighbor 5.5.5.5 remote-as 5000	Specify a VTEP5 loopback ip address and remote-as defined
(config-router)#neighbor 5.5.5.5 update-source lo	Configure update as loopback for VTEP5
(config-router)#neighbor 5.5.5.5 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP5
(config-router)#address-family l2vpn evpn	Enter into l2vpn EVPN address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1(VTEP1) into l2vpn evpn address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2(VTEP2) into l2vpn evpn address family mode
(config-router-af)#neighbor 5.5.5.5 activate	Activate 5.5.5.5(VTEP5) into l2vpn evpn address family mode
(config-router-af)#exit-address-family	Exit from l2vpn address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode

VRF Configuration:

(config)#mac vrf L2VRF1	Create mac routing/forwarding instance with L2VRF1 name and enter into vrf mode
(config-vrf)#rd 4.4.4.4:11	Assign RD value
(config-vrf)#description MAC VRF RED	Give description to L2VRF1 as RED
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
(config-vrf)#exit	Exit from vrf mode
(config)#mac vrf L2VRF2	Create mac routing/forwarding instance with L2VRF2 name and enter into vrf mode
(config-vrf)#rd 4.4.4.4:21	Assign RD value

(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#description MAC VRF BLUE	Give description to L2VRF2 as BLUE
(config-vrf)#exit	Exit from vrf mode

VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan vtep-ip-global 4.4.4.4	Configure Source vtep-ip-global configuration. Use loopback ip address
(config)#nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# vni-name VNI-101	Configure VNI name as VNI-101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# vni-name VNI-201	Configure VNI name as VNI-201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)# nvo vxlan access-if port-vlan xe3 20	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN)
(config-nvo-acc-if)#map vnid 201	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)# mac 0000.5555.1020 ip 21.21.21.101	Configure static mac-ip
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#exit	Exit from configuration mode

VTEP5

Single Home -SH3

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#evpn vxlan multihoming enable	Enable Multihoming, save configs and reboot the board for multihoming to be effective
(config)#qos enable	Enabling qos

Interface and loopback configuration:

`(config)#interface xe1`	Enter Interface mode for xe1 (SH5)
`(config-if)#switchport`	Make it L2 interface
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 5.5.5.5/32 secondary`	Configure loopback ip address as 5.5.5.5 for VTEP5
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) interface ce53`	Enter interface mode
`(config-if)#ip address 100.15.15.1/24`	Configure ip address as 100.15.15.1 on network side of Spine-P3
`(config-if)#exit`	Exit Interface mode and return to Configure mode.

OSPF configuration:

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 5.5.5.5`	Configure router-id as 5.5.5.5 (lo ip address)
`(config-router)#network 5.5.5.5/32 area 0.0.0.0`	Add 5.5.5.5 (lo ip address) network into area 0
`(config-router)#network 100.15.15.0/24 area 0.0.0.0`	Add 100.15.15.0(Spine-P3) network into area 0
`(config-router)#bfd all-interfaces`	Enabling bfd on all ospf interface for fast convergence
`(config-if)#exit`	Exit Interface mode and return to Configure mode.

BGP configuration:

`(config)#router bgp 5000`	Enter into Router BGP mode
`(config-router)#bgp router-id 5.5.5.5`	Configure router-id as 5.5.5.5(lo ip address)
`(config-router)#neighbor 1.1.1.1 remote-as 5000`	Specify a VTEP1 loopback ip address and remote-as defined
`(config-router)#neighbor 1.1.1.1 update-source lo`	Configure update as loopback for VTEP1
`(config-router)#neighbor 1.1.1.1 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP1
`(config-router)#neighbor 2.2.2.2 remote-as 5000`	Specify a VTEP2 loopback ip address and remote-as defined
`(config-router)#neighbor 2.2.2.2 update-source lo`	Configure update as loopback for VTEP2
`(config-router)#neighbor 2.2.2.2 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP2
`(config-router)#neighbor 4.4.4.4 remote-as 5000`	Specify a VTEP4 loopback ip address and remote-as defined
`(config-router)#neighbor 4.4.4.4 update-source lo`	Configure update as loopback for VTEP4
`(config-router)#neighbor 4.4.4.4 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP4

(config-router)#address-family l2vpn evpn	Enter into l2vpn EVPN address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate 1.1.1.1(VTEP1) into l2vpn evpn address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate 2.2.2.2(VTEP2) into l2vpn evpn address family mode
(config-router-af)#neighbor 4.4.4.4 activate	Activate 4.4.4.4(VTEP4) into l2vpn evpn address family mode
(config-router-af)#exit-address-family	Exit from l2vpn address family mode
(config-router)#exit	Exit from Router BGP mode and enter into config mode

VRF Configuration:

(config)#mac vrf L2VRF1	Create mac routing/forwarding instance with L2VRF1 name and enter into vrf mode
(config-vrf)#rd 5.5.5.5:11	Assign RD value
(config-vrf)#description MAC VRF RED	Give description to L2VRF1 as RED
(config-vrf)#route-target both 9.9.9.9:100	Assign route-target value for same for import and export. Should be same on all node for L2VRF1
(config-vrf)#exit	Exit from vrf mode
(config)#mac vrf L2VRF2	Create mac routing/forwarding instance with L2VRF2 name and enter into vrf mode
(config-vrf)#rd 5.5.5.5:21	Assign RD value
(config-vrf)#route-target both 90.90.90.90:100	Assign route-target value for same for import and export
(config-vrf)#description MAC VRF BLUE	Give description to L2VRF2 as BLUE
(config-vrf)#exit	Exit from vrf mode

VXLAN configuration:

(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan vtep-ip-global 5.5.5.5	Configure Source vtep-ip-global configuration. Use loopback ip address
(config)#nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# vni-name VNI-101	Configure VNI name as VNI-101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# vni-name VNI-201	Configure VNI name as VNI-201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)# nvo vxlan access-if port-vlan xe1 10	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN)
(config-nvo-acc-if)#map vnid 101	Map VXLAN Identified to access-port for VXLAN

(config-nvo-acc-if)# mac 0000.4444.1010 ip 11.11.11.201	Configure static mac-ip
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#exit	Exit from VXLAN access-interface mode and enter into configuration mode
(config)#exit	Exit from configuration mode

Switch1 (MH2)

Multihomed to 2-VTEPs (VTEP1 and VTEP2). It acts as Tenant system for vlan1.20.

#configure terminal	Enter Configure mode.
(config)# bridge 1 protocol rstp vlan-bridge	Configure rstp vlan bridge
(config)# vlan database	Enter vlan database config mode
(config)#vlan 2-20 bridge 1 state enable	Configure vlans from 2-20 and associate with bridge 1
(config)#interface xe5	Enter Interface mode for xe5 which is connected to TG
(config-if)#switchport	Make as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate bridge 1 into interface
(config-if)# bridge-group 1 spanning-tree disable	Configure interface as stp disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 10,20	Trunk allowed vlan as 10.20
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface po1	Enter Interface mode for po1
(config-if)#switchport	Make po1 as L2 port by configuring switchport
(config-if)#bridge-group 1	Associate po1 to bridge 1
(config-if)# bridge-group 1 spanning-tree disable	Configure po1 as stp disable
(config-if)# switchport mode trunk	Mode as trunk
(config-if)# switchport trunk allowed vlan add 10,20	Trunk allowed vlan as 2.10.20
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe3	Enter Interface mode for xe3
(config-if)#channel-group 1 mode active	Make it member port of po1
(config)#interface xe2	Enter Interface mode for xe2
(config-if)#channel-group 1 mode active	Make it member port of po1
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit from configuration mode

Spine-P3

Spine node where all VTEPs are connected.

Generic configuration:

`#configure terminal`	Enter Configure mode.
`(Config)#qos enable`	Enabling qos

Interface and loopback configuration:

`(config)#interface lo`	Enter Interface mode for lo
`(config-if)#ip address 100.100.100.100/32 secondary`	Configure loopback ip address as 100.100.100.100 for Spine-P3
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#interface ce16/1`	Enter Interface mode for ce16/1
`(config-if)#ip address 100.11.11.2/24`	Configure ip address as 100.11.11.2 on network side of VTEP1
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) interface ce5/1`	Enter into ce5/1 interface mode
`(config-if)#ip address 100.12.12.2/24`	Configure ip address as 100.12.12.2 on network side of VTEP2
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) interface ce9/1`	Enter ce9/1 interface mode
`(config-if)#ip address 100.14.14.2/24`	Configure ip address as 100.14.14.12 on network side of VTEP4
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config) interface ce23/1`	Enter interface mode
`(config-if)#ip address 100.15.15.1/24`	Configure ip address as 100.15.15.1 on network side of VTEP5
`(config-if)#exit`	Exit Interface mode and return to Configure mode.

OSPF configuration:

`(config)#router ospf 100`	Enter into router OSPF mode
`(config-router)#ospf router-id 100.100.100.100`	Configure router-id as 100.100.100.100 (lo ip address)
`(config-router)#network 100.100.100.100/32 area 0.0.0.0`	Add 100.100.100.100 (lo ip address) network into area 0
`(config-router)#network 100.11.11.0/24 area 0.0.0.0`	Add 100.11.11.0 (VTEP1) network into area 0
`(config-router)#network 100.12.12.0/24 area 0.0.0.0`	Add 100.12.12.0 (VTEP2) network into area 0
`(config-router)#network 100.14.14.0/24 area 0.0.0.0`	Add 100.14.14.0 (VTEP4) network into area 0
`(config-router)#network 100.15.15.0/24 area 0.0.0.0`	Add 100.15.15.0 (VTEP5) network into area 0
`(config-router)#bfd all-interfaces`	Enabling bfd on all ospf interface for fast convergence

(config-router)#commit	Commit the candidate configuration to the running configuration
(config-router)#exit	Exit Interface mode and return to Configure mode.

Centralized Gateway

In Centralized gateway approach, when two TS belonging to two different subnets connected to the same/different VTEP node, wanted to communicate with each other, their traffic needed to be back hauled from the VTEP node to the centralized gateway node where inter- subnet switching is performed and then back to the VTEP node.

IRB Configuration for Centralized Gateway

Configure from Base Configuration-L2 VXLAN section, then configure below commands for centralized gateway approach. Here VTEP4 is the centralized GW. In VTEP4, dscp-to-queue qos profile should be applied on the particular incoming L2VNID IRB interface.

Note: For L3 traffic, when L2VNID is sent in the traffic, then dscp-to-queue qos profile mapped at IRB interface of that particular L2VNID takes effect.

VTEP1

(config)#qos profile cos-to-queue COS_QUEUE	Create QoS profile for mapping traffic towards tunnel from access-if
(config-ingress-cos-map)#cos 1 queue 5	Configure particular COS value to queue value for the profile
(config-ingress-cos-map)#exit	Exit from qos profile mode
(config)#qos profile queue-color-to-dscp QUEUE_DSCP	Create QoS profile for attaching in vxlan tunnel egress
(config-egress-dscp-map)#queue 5 dscp 34	Configure queue value to DSCP value for the profile
(config)#nvo vxlan tunnel qos-map-mode cos-dscp egress QUEUE_DSCP	Map the configured QoS profile to vxlan tunnel egress
(config)#nvo vxlan access-if port-vlan po1 20	Enter into vxlan access port mode
(config-nvo-acc-if)#map qos-profile cos-to-queue COS_QUEUE	Map the qos profile in vxlan access-if
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#end	Exit from vxlan access port

VTEP2

(config)#qos profile cos-to-queue COS_QUEUE	Create QoS profile for mapping traffic towards tunnel from access-if
(config-ingress-cos-map)#cos 1 queue 5	Configure particular COS value to queue value for the profile
(config-ingress-cos-map)#exit	Exit from qos profile mode
(config)#qos profile queue-color-to-dscp QUEUE_DSCP	Create QoS profile for attaching in vxlan tunnel egress
(config-egress-dscp-map)#queue 5 dscp 34	Configure queue value to DSCP value for the profile

(config)#nvo vxlan tunnel qos-map-mode cos-dscp egress QUEUE_DSCP	Map the configured QoS profile to vxlan tunnel egress
(config)#nvo vxlan access-if port-vlan po1 20	Enter into vxlan access port mode
(config-nvo-acc-if)#map qos-profile cos-to-queue COS_QUEUE	Map the qos profile in vxlan access-if
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#end	Exit from vxlan access port

VTEP4

(config)#nvo vxlan irb	Enable VXLAN irb
(config)#ip vrf L3VRF1	Create mac routing/forwarding instance with L3VRF1 name and enter into vrf mode
(config-vrf)#rd 51000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from vrf mode
(config)# interface irb1001	Configure IRB interface 1001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 11.11.11.1/24	Configure ip address
(config-if)ipv6 address 1111::1/64	Configure ipv6 address
(config-if)exit	Exit from interface config mode
(config)# interface irb2001	Configure IRB interface 2001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 21.21.21.1/24	Configure ip address
(config-if)ipv6 address 2121::1/64	Configure ipv6 address
(config-if)exit	Exit from interface config mode
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1001	Configure irb1001 under vxlan id 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb2001	Configure irb2001 under vxlan id 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#qos profile dscp-to-queue DSCP_QUEUE	Configure the QoS profile to match the inner dscp value of the traffic at the IRB interface

(config-ingress-dscp-map)#dscp 20 queue 1	Configure particular dscp to a queue value. Configure particular dscp to a queue value. Here classification at the IRB L3 interface is based on customer dscp value.
(config-ingress-dscp-map)#exit	Exit from qos profile config mode
config)#int irb2001	Enter IRB L3 interface mode
(config-irb-if)#qos map-profile dscp-to-queue DSCP_QUEUE	Map the qos profile in the IRB interface
(config-irb-if)#exit	Exit from interface mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#end	Exit from global config mode

VTEP5

(config)#qos profile queue-color-to-cos QUEUE_COS	Create QoS profile for remark the queue value to COS value
(config-egress-cos-map)#queue 2 cos 5	Configure particular queue value to COS value for the profile
(config-ingress-cos-map)#exit	Exit from qos profile mode
(config)# qos profile dscp-to-queue DSCP_QUEUE	Create QoS profile for attaching in vxlan tunnel ingress
(config-ingress-dscp-map)#dscp 56 queue 2	Configure DSCP value to queue value for the profile
(config)#nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP_QUEUE	Map the configured QoS profile to vxlan tunnel ingress
(config)#nvo vxlan access-if port-vlan xe1 10	Enter into vxlan access port mode
(config-nvo-acc-if)#map qos-profile queue-color-to-cos QUEUE_COS	Map the qos profile in vxlan access-if
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#end	Exit from vxlan access port

Validation

Send traffic from TS2-21 to MH2 access-if with dscp value 20 and COS value 1(vlan20) and verify traffic reaceived at TS1-11 with dscp value 32 and COS value 5(vlan10) at the VTEP5 access-if.

VTEP1

```
VTEP1#show running-config qos
qos enable
!
qos profile cos-to-queue COS_QUEUE
 cos 1 queue 5
!
qos profile queue-color-to-dscp QUEUE_DSCP
 queue 5 color all dscp 34
!
!
!
VTEP1#show running-config nvo vxlan
!
nvo vxlan enable
!
```

```
evpn vxlan multihoming enable
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan tunnel qos-map-mode cos-dscp egress QUEUE_DSCP
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 vni-name VNI-201
!
nvo vxlan access-if port-vlan po1 10
 map vnid 101
 mac 0000.2222.1010 ip 11.11.11.51
!
nvo vxlan access-if port-vlan po1 20
 map vnid 201
 mac 0000.2222.1020 ip 21.21.21.51
 map qos-profile cos-to-queue COS_QUEUE
!
!
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
1.1.1.1          5.5.5.5          Installed        01:15:13       01:15:13
1.1.1.1          4.4.4.4          Installed        01:15:28       01:15:28
1.1.1.1          2.2.2.2          Installed        01:11:40       01:11:40

Total number of entries are 3
VTEP1#show interface ce52 counters queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes      |  Dropped pkts   |   Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------+-------------------+
q0                 (E) 125304832 0                 0                   0                 0
q1                 (E) 125304832 0                 0                   0                 0
q2                 (E) 125304832 0                 0                   0                 0
q3                 (E) 125304832 0                 0                   0                 0
q4                 (E) 125304832 0                 0                   0                 0
q5                 (E) 125304832 1316880           1316879000          0                 0
q6                 (E) 125304832 0                 0                   0                 0
q7                 (E) 125304832 0                 0                   0                 0
VTEP1#show qos-profile COS_QUEUE
profile name: COS_QUEUE
profile type: cos-to-queue
profile attached to 1 instances
configured mapping:
 cos 1 queue 5
Detailed mapping:
 ---------------+----------------- | ---------------+-----------------
       INPUT     |     OUTPUT       |       INPUT     |     OUTPUT
 ---------------+----------------- | ---------------+-----------------
    COS  |  DEI  | Queue |  Color   |   COS  |  DEI  | Queue |  Color
 -------+-------+-------+--------- | -------+-------+-------+---------
   0       0       0       green    |  0       1       0       yellow
   1       0       5       green    |  1       1       5       yellow
   2       0       2       green    |  2       1       2       yellow
   3       0       3       green    |  3       1       3       yellow
   4       0       4       green    |  4       1       4       yellow
   5       0       5       green    |  5       1       5       yellow
   6       0       6       green    |  6       1       6       yellow
   7       0       7       green    |  7       1       7       yellow

VTEP1#show qos-profile QUEUE_DSCP
profile name: QUEUE_DSCP
profile type: queue-color-to-dscp
profile attached to 1 instances
configured mapping:
 queue 5 color all dscp 34
```

```
Detailed mapping:
 ----------------+-------- | ----------------+-------- | ----------------+--------
      INPUT      | OUTPUT  |      INPUT      | OUTPUT  |      INPUT      | OUTPUT
 ----------------+-------- | ----------------+-------- | ----------------+--------
  Queue | Color  | DSCP    |  Queue | Color  | DSCP    |  Queue | Color  | DSCP
 -------+--------+-------- | -------+--------+-------- | -------+--------+--------
  0       green    0       |  0       yellow   0       |  0       red      0
  1       green    10      |  1       yellow   12      |  1       red      14
  2       green    18      |  2       yellow   20      |  2       red      22
  3       green    26      |  3       yellow   28      |  3       red      30
  4       green    34      |  4       yellow   36      |  4       red      38
  5       green    34      |  5       yellow   34      |  5       red      34
  6       green    48      |  6       yellow   48      |  6       red      48
  7       green    56      |  7       yellow   56      |  7       red      56
```

VTEP4

```
VTEP4#show running-config qos
qos enable
!
qos profile dscp-to-queue DSCP_QUEUE
 dscp 20 queue 1
!
!
!
interface irb2001
 qos map-profile dscp-to-queue DSCP_QUEUE
!
VTEP4#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
nvo vxlan vtep-ip-global 4.4.4.4
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan xe3 20
 map vnid 201
 mac 0000.5555.1020 ip 21.21.21.101
!
!
VTEP4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status           Up/Down        Update
================================================================================
4.4.4.4           2.2.2.2          Installed        00:08:40       00:08:40
4.4.4.4           1.1.1.1          Installed        00:12:28       00:12:28
4.4.4.4           5.5.5.5          Installed        00:12:13       00:12:13

Total number of entries are 3
VTEP4#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
```

```
101      VNI-101      L2       NW   ----      ----                          ---- ----      4.4.4.4          2.2.2.2
101      VNI-101      L2       NW   ----      ----                          ---- ----      4.4.4.4          1.1.1.1
101      VNI-101      L2       NW   ----      ----                          ---- ----      4.4.4.4          5.5.5.5
201      VNI-201      L2       NW   ----      ----                          ---- ----      4.4.4.4          2.2.2.2
201      VNI-201      L2       NW   ----      ----                          ---- ----      4.4.4.4          1.1.1.1
201      VNI-201      L2       NW   ----      ----                          ---- ----      4.4.4.4          5.5.5.5
201      VNI-201      --       AC   xe3       --- Single Homed Port ---     20   ----      ----             ----

Total number of entries are 7
VTEP4#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       3c2c.991a.da7a Static Local   ----
201      21.21.21.101     0000.5555.1020 Static Local   ----
101      11.11.11.51      0000.2222.1010 Static Remote  ----
101      11.11.11.1       3c2c.991a.da7a Static Local   ----
101      11.11.11.201     0000.4444.1010 Static Remote  ----
Total number of entries are 6
VTEP4#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                 Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________________________________
201      2121::1                        3c2c.991a.da7a Static Local     ----
101      1111::1                        3c2c.991a.da7a Static Local     ----
Total number of entries are 2
VTEP4#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001
 1000         201          irb2001

VTEP4#show ipv4 route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:01:35
C            21.21.21.0/24 is directly connected, irb2001, 00:01:16
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:06:12

Gateway of last resort is not set
VTEP4#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C       ::1/128 via ::, lo.L3VRF1, 00:06:29
C       1111::/64 via ::, irb1001, 00:01:52
C       2121::/64 via ::, irb2001, 00:01:33
C       fe80::/64 via ::, irb2001, 00:01:33
VTEP4#show interface ce49 counters queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes      |  Dropped pkts   |   Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------+-------------------+
q0                (E) 125304832 0                 0                   0                 0
q1                (E) 125304832 1422755           1488201730          0                 0
q2                (E) 125304832 0                 0                   0                 0
```

```
q3                (E) 125304832 0                0                0                0
q4                (E) 125304832 0                0                0                0
q5                (E) 125304832 0                0                0                0
q6                (E) 125304832 0                0                0                0
q7                (E) 125304832 0                0                0                0
VTEP4#show qos-profile DSCP_QUEUE
profile name: DSCP_QUEUE
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
 dscp 20 queue 1
Detailed mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--------------------------- | -------+-
--------------------------
  INPUT |        OUTPUT              |  INPUT |        OUTPUT              |  INPUT |        OUTPUT              |  INPUT |
OUTPUT
 -------+--------------------------- | -------+--------------------------- | -------+--------------------------- | -------+-
--------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+-------+--------+---------- | -------+-
------+--------+----------
  0      0      green    0         | 16     2      green    16        | 32     4      green    32        | 48     6
green    48
  1      0      green    1         | 17     2      green    17        | 33     4      green    33        | 49     6
green    49
  2      0      green    2         | 18     2      green    18        | 34     4      green    34        | 50     6
green    50
  3      0      green    3         | 19     2      green    19        | 35     4      green    35        | 51     6
green    51
  4      0      green    4         | 20     1      yellow   20        | 36     4      yellow   36        | 52     6
green    52
  5      0      green    5         | 21     2      green    21        | 37     4      green    37        | 53     6
green    53
  6      0      green    6         | 22     2      yellow   22        | 38     4      yellow   38        | 54     6
green    54
  7      0      green    7         | 23     2      green    23        | 39     4      green    39        | 55     6
green    55
  8      1      green    8         | 24     3      green    24        | 40     5      green    40        | 56     7
green    56
  9      1      green    9         | 25     3      green    25        | 41     5      green    41        | 57     7
green    57
  10     1      green    10        | 26     3      green    26        | 42     5      green    42        | 58     7
green    58
  11     1      green    11        | 27     3      green    27        | 43     5      green    43        | 59     7
green    59
  12     1      yellow   12        | 28     3      yellow   28        | 44     5      green    44        | 60     7
green    60
  13     1      green    13        | 29     3      green    29        | 45     5      green    45        | 61     7
green    61
  14     1      yellow   14        | 30     3      yellow   30        | 46     5      green    46        | 62     7
green    62
  15     1      green    15        | 31     3      green    31        | 47     5      green    47        | 63     7
green    63

VTEP4#show qos-profile interface irb2001
profile name: DSCP_QUEUE
 profile type: dscp-to-queue (Ingress)
 mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--------------------------- | -------+-
--------------------------
  INPUT |        OUTPUT              |  INPUT |        OUTPUT              |  INPUT |        OUTPUT              |  INPUT |
OUTPUT
 -------+--------------------------- | -------+--------------------------- | -------+--------------------------- | -------+-
--------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+-------+--------+---------- | -------+-
------+--------+----------
  0      0      green    0         | 16     2      green    16        | 32     4      green    32        | 48     6
green    48
  1      0      green    1         | 17     2      green    17        | 33     4      green    33        | 49     6
green    49
  2      0      green    2         | 18     2      green    18        | 34     4      green    34        | 50     6
green    50
  3      0      green    3         | 19     2      green    19        | 35     4      green    35        | 51     6
green    51
  4      0      green    4         | 20     1      yellow   20        | 36     4      yellow   36        | 52     6
green    52
```

```
  5       0       green    5         |  21      2       green    21        |  37      4       green    37        |  53      6
green     53
  6       0       green    6         |  22      2       yellow   22        |  38      4       yellow   38        |  54      6
green     54
  7       0       green    7         |  23      2       green    23        |  39      4       green    39        |  55      6
green     55
  8       1       green    8         |  24      3       green    24        |  40      5       green    40        |  56      7
green     56
  9       1       green    9         |  25      3       green    25        |  41      5       green    41        |  57      7
green     57
  10      1       green    10        |  26      3       green    26        |  42      5       green    42        |  58      7
green     58
  11      1       green    11        |  27      3       green    27        |  43      5       green    43        |  59      7
green     59
  12      1       yellow   12        |  28      3       yellow   28        |  44      5       green    44        |  60      7
green     60
  13      1       green    13        |  29      3       green    29        |  45      5       green    45        |  61      7
green     61
  14      1       yellow   14        |  30      3       yellow   30        |  46      5       green    46        |  62      7
green     62
  15      1       green    15        |  31      3       green    31        |  47      5       green    47        |  63      7
green     63
```

VTEP5

```
VTEP5#show running-config qos
qos enable
!
qos profile queue-color-to-cos QUEUE_COS
 queue 2 color all cos 5
!
qos profile dscp-to-queue DSCP_QUEUE
 dscp 56 queue 2
!
!
!
VTEP5#show running-config nvo vxlan
!
nvo vxlan enable
!
evpn vxlan multihoming enable
!
nvo vxlan vtep-ip-global 5.5.5.5
!
nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP_QUEUE
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 vni-name VNI-201
!
nvo vxlan access-if port-vlan xe1 10
 map vnid 101
 mac 0000.4444.1010 ip 11.11.11.201
 map qos-profile queue-color-to-cos QUEUE_COS
!
!
VTEP5#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status           Up/Down        Update
================================================================================
5.5.5.5           2.2.2.2          Installed        01:11:17       01:11:17
5.5.5.5           4.4.4.4          Installed        01:14:50       01:14:50
5.5.5.5           1.1.1.1          Installed        01:14:50       01:14:50

Total number of entries are 3
VTEP5#show interface xe1 counters queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |      Tx bytes     |  Dropped pkts   |   Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------+-------------------+
q0                (E) 12517376 0                 0                   0                 0
```

```
q1                (E) 12517376 0                 0                   0                 0
q2                (E) 12517376 1005800           1052066800          0                 0
q3                (E) 12517376 0                 0                   0                 0
q4                (E) 12517376 0                 0                   0                 0
q5                (E) 12517376 0                 0                   0                 0
q6                (E) 12517376 0                 0                   0                 0
q7                (E) 12517376 0                 0                   0                 0
VTEP5#sh qos-profile QUEUE_COS
profile name: QUEUE_COS
profile type: queue-color-to-cos
profile attached to 1 instances
configured mapping:
 queue 2 color all cos 5
Detailed mapping:
 ----------------+-------- | ----------------+-------- | ----------------+--------
      INPUT      | OUTPUT  |      INPUT      | OUTPUT  |      INPUT      | OUTPUT
 ----------------+-------- | ----------------+-------- | ----------------+--------
  Queue | Color  | COS     |  Queue | Color  | COS     |  Queue | Color  | COS
 -------+--------+-------- | -------+--------+-------- | -------+--------+--------
  0       green    0       |  0       yellow   0       |  0       red      0
  1       green    1       |  1       yellow   1       |  1       red      1
  2       green    5       |  2       yellow   5       |  2       red      5
  3       green    3       |  3       yellow   3       |  3       red      3
  4       green    4       |  4       yellow   4       |  4       red      4
  5       green    5       |  5       yellow   5       |  5       red      5
  6       green    6       |  6       yellow   6       |  6       red      6
  7       green    7       |  7       yellow   7       |  7       red      7

VTEP5#show qos-profile DSCP_QUEUE
profile name: DSCP_QUEUE
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
 dscp 56 queue 2
Detailed mapping:
 -------+---------------------------- | -------+---------------------------- | -------+---------------------------- | -------+-
--------------------------
  INPUT |        OUTPUT             |  INPUT |        OUTPUT             |  INPUT |        OUTPUT             |  INPUT |
OUTPUT
 -------+---------------------------- | -------+---------------------------- | -------+---------------------------- | -------+-
--------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+-------+--------+---------- | -------+-
------+--------+----------
  0       0       green    0         |  16      2       green    16        |  32      4       green    32        |  48      6
green    48
  1       0       green    1         |  17      2       green    17        |  33      4       green    33        |  49      6
green    49
  2       0       green    2         |  18      2       green    18        |  34      4       green    34        |  50      6
green    50
  3       0       green    3         |  19      2       green    19        |  35      4       green    35        |  51      6
green    51
  4       0       green    4         |  20      2       yellow   20        |  36      4       yellow   36        |  52      6
green    52
  5       0       green    5         |  21      2       green    21        |  37      4       green    37        |  53      6
green    53
  6       0       green    6         |  22      2       yellow   22        |  38      4       yellow   38        |  54      6
green    54
  7       0       green    7         |  23      2       green    23        |  39      4       green    39        |  55      6
green    55
  8       1       green    8         |  24      3       green    24        |  40      5       green    40        |  56      2
green    56
  9       1       green    9         |  25      3       green    25        |  41      5       green    41        |  57      7
green    57
  10      1       green    10        |  26      3       green    26        |  42      5       green    42        |  58      7
green    58
  11      1       green    11        |  27      3       green    27        |  43      5       green    43        |  59      7
green    59
  12      1       yellow   12        |  28      3       yellow   28        |  44      5       green    44        |  60      7
green    60
  13      1       green    13        |  29      3       green    29        |  45      5       green    45        |  61      7
green    61
  14      1       yellow   14        |  30      3       yellow   30        |  46      5       green    46        |  62      7
green    62
  15      1       green    15        |  31      3       green    31        |  47      5       green    47        |  63      7
green    63
```

Anycast Gateway

For today's large multi-tenant data center, centralized L3 gateway scheme is very inefficient and sometimes impractical. In order to overcome the drawback of centralized L3GW approach, anycast mode is used.

In Anycast gateway approach, all the VTEPs acts as default gateway for all the VNIDs. We will configure same anycast mac in all VTEPs.

IRB Configuration for Anycast

Configure from Base Configuration-L2 VXLAN section, then configure below commands for Anycast gateway approach.

Note: For L2 traffic, always dscp-to-queue qos profile at tunnel ingress takes effect.

Note: For L3 traffic in the local VTEP, routing is done at IRB level and also QoS applied at the IRB interface and it sends with l2vnid.

VTEP1

`(config)#nvo vxlan irb`	Enable VXLAN irb
`(config)#ip vrf L3VRF1`	Create mac routing/forwarding instance with L3VRF1 name and enter into vrf mode
`(config-vrf)#rd 11000:11`	Assign RD value
`(config-vrf)# route-target both 100:100`	Assign route-target value for same for import and export.
`(config-vrf)# l3vni 1000`	Configure L3VNI as 1000 for L3VRF1
`(config-vrf)#exit`	Exit from vrf mode
`(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111`	Configure anycast mac address
`(config)# interface irb1001`	Configure IRB interface 1001
`(config-if)ip vrf forwarding L3VRF1`	Configure L3VRF1
`(config-if)ip address 11.11.11.1/24`	Configure ip address
`(config-if)ipv6 address 1111::1/64`	Configure ipv6 address
`(config-if) evpn irb-if-forwarding anycast-gateway-mac`	Configure anycast mac address
`(config-if)exit`	Exit from interface config mode
`(config)# interface irb2001`	Configure IRB interface 2001
`(config-if)ip vrf forwarding L3VRF1`	Configure L3VRF1
`(config-if)ip address 21.21.21.1/24`	Configure ip address
`(config-if)ipv6 address 2121::1/64`	Configure ipv6 address
`(config-if) evpn irb-if-forwarding anycast-gateway-mac`	Configure anycast mac address
`(config-if)exit`	Exit from interface config mode
`(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled`	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode

(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1001	Configure irb1001 under vxlan id 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb2001	Configure irb2001 under vxlan id 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#qos profile dscp-to-queue DSCP_QUEUE	Configure QoS profile to remark the traffic received from access-if
(config-ingress-dscp-map)#dscp 20 queue 4	Configure particular dscp value to queue value in the profile
(config-ingress-dscp-map)#exit	Exit from qos profile mode
(config)#int irb2001	Enter IRB L3 interface
(config-irb-if)#qos map-profile dscp-to-queue DSCP_QUEUE	Map the qos profile
(config-irb-if)#exit	Exit from interface mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#end	Exit from global config mode

VTEP2

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#ip vrf L3VRF1	Create mac routing/forwarding instance with L3VRF1 name and enter into vrf mode
(config-vrf)#rd 21000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from vrf mode
(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast mac address
(config)# interface irb1001	Configure IRB interface 1001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 11.11.11.1/24	Configure ip address
(config-if)ipv6 address 1111::1/64	Configure ipv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast mac address
(config-if)exit	Exit from interface config mode
(config)# interface irb2001	Configure IRB interface 2001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 21.21.21.1/24	Configure ip address

(config-if)ipv6 address 2121::1/64	Configure ipv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast mac address
(config-if)exit	Exit from interface config mode
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1001	Configure irb1001 under vxlan id 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb2001	Configure irb2001 under vxlan id 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#qos profile dscp-to-queue DSCP_QUEUE	Configure QoS profile to remark the traffic at remote vtep at the ingress tunnel.
(config-ingress-dscp-map)#dscp 20 queue 4	Configure particular dscp value to queue value in the profile
(config-ingress-dscp-map)#exit	Exit from qos profile mode
(config)#int irb2001	Enter IRB L3 interface
(config-irb-if)#qos map-profile dscp-to-queue DSCP_QUEUE	Map the qos profile
(config-irb-if)#exit	Exit from interface mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#end	Exit from global config mode

VTEP4

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#ip vrf L3VRF1	Create mac routing/forwarding instance with L3VRF1 name and enter into vrf mode
(config-vrf)#rd 41000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from vrf mode
(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast mac address
(config)# interface irb1001	Configure IRB interface 1001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 11.11.11.1/24	Configure ip address
(config-if)ipv6 address 1111::1/64	Configure ipv6 address

(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast mac address
(config-if)exit	Exit from interface config mode
(config)# interface irb2001	Configure IRB interface 2001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 21.21.21.1/24	Configure ip address
(config-if)ipv6 address 2121::1/64	Configure ipv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast mac address
(config-if)exit	Exit from interface config mode
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1001	Configure irb1001 under vxlan id 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb2001	Configure irb2001 under vxlan id 201
(config-nvo)#commit	Commit the candidate configuration to the running configuration
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.

VTEP5

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#ip vrf L3VRF1	Create mac routing/forwarding instance with L3VRF1 name and enter into vrf mode
(config-vrf)#rd 51000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from vrf mode
(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111	Configure anycast mac address
(config)# interface irb1001	Configure IRB interface 1001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 11.11.11.1/24	Configure ip address
(config-if)ipv6 address 1111::1/64	Configure ipv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast mac address
(config-if)exit	Exit from interface config mode
(config)# interface irb2001	Configure IRB interface 2001

(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 21.21.21.1/24	Configure ip address
(config-if)ipv6 address 2121::1/64	Configure ipv6 address
(config-if) evpn irb-if-forwarding anycast-gateway-mac	Configure anycast mac address
(config-if)exit	Exit from interface config mode
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1001	Configure irb1001 under vxlan id 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb2001	Configure irb2001 under vxlan id 201
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#qos profile dscp-to-queue DSCP_QUEUE	Configure the qos profile to remark outer dscp to queue of the the ingress tunnel traffic. Here classification at the ingress vtep is based on overlay dscp value.
(config-ingress-dscp-map)#dscp 34 queue 2	Configure particular dscp value to queue value in the profile. Here dscp to <queue, dscp> is not applicable. I.e ingress remarking of dscp at the ingress tunnel is not applicable.
(config-ingress-dscp-map)#exit	Exit from qos profile mode
(config)#qos profile queue-color-to-cos QUEUE_COS	Configure qos profile to remark the queue value to COS value
(config-egress-cos-map)#queue 2 cos 6	Configure particular queue value to COS value
(config-egress-cos-map)#exit	Exit from qos profile config mode
(config)#nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP_QUEUE	Map the qos profile in tunnel ingress
(config)#nvo vxlan access-if port-vlan xe1 10	Enter to vxlan access port config mode
(config-nvo-acc-if)#map qos-profile queue-color-to-cos QUEUE_COS	Map the qos profile in vxlan access port
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#end	Exit from config mode

Validations

Send traffic from TS2-21 to MH2 access-if with dscp value 20 and COS value 1(vlan20) and verify traffic received at TS1-11 with dscp value 20 and COS value 6(vlan10) at the VTEP5 access-if.

VTEP1

```
VTEP1#show running-config qos
```

```
qos enable
!
qos profile dscp-to-queue DSCP_QUEUE
 dscp 20 queue 4
!
!
!
!
interface irb2001
 qos map-profile dscp-to-queue DSCP_QUEUE
!
VTEP1#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
evpn irb-forwarding anycast-gateway-mac 0000.0000.1111
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan po1 10
 map vnid 101
 mac 0000.2222.1010 ip 11.11.11.51
!
nvo vxlan access-if port-vlan po1 20
 map vnid 201
 mac 0000.2222.1020 ip 21.21.21.51
!
!
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status            Up/Down         Update
================================================================================
1.1.1.1           5.5.5.5          Installed         01:15:13        01:15:13
1.1.1.1           4.4.4.4          Installed         01:15:28        01:15:28
1.1.1.1           2.2.2.2          Installed         01:11:40        01:11:40
```

```
Total number of entries are 3
VTEP1#show nvo vxlan
VXLAN Information
================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID    VNI-Name    VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
101     VNI-101     L2       NW   ----      ----                         ---- ----
1.1.1.1          5.5.5.5
101     VNI-101     L2       NW   ----      ----                         ---- ----
1.1.1.1          4.4.4.4
101     VNI-101     L2       NW   ----      ----                         ---- ----
1.1.1.1          2.2.2.2
101     VNI-101     --       AC   po1       00:00:00:00:00:22:22:00:00:00 10   DF
----             ----
201     VNI-201     L2       NW   ----      ----                         ---- ----
1.1.1.1          5.5.5.5
201     VNI-201     L2       NW   ----      ----                         ---- ----
1.1.1.1          4.4.4.4
201     VNI-201     L2       NW   ----      ----                         ---- ----
1.1.1.1          2.2.2.2
201     VNI-201     --       AC   po1       00:00:00:00:00:22:22:00:00:00 20   DF
----             ----

Total number of entries are 8
VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Local   ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.101     0000.5555.1020 Static Remote  ----
101      11.11.11.51      0000.2222.1010 Static Local   ----
101      11.11.11.10      0010.9400.0002 Dynamic Remote ----
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.201     0000.4444.1010 Static Remote  ----
Total number of entries are 7
VTEP1#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________
_____________
201      2121::1                          0000.0000.1111 Static Local     ----
101      1111::10                         0010.9400.0002 Dynamic Remote   ----
101      1111::1                          0000.0000.1111 Static Local     ----
```

```
Total number of entries are 3
VTEP1#show ipv4 route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:13:19
C            21.21.21.0/24 is directly connected, irb2001, 00:12:56
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:17:13

Gateway of last resort is not set
VTEP1#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C       ::1/128 via ::, lo.L3VRF1, 00:17:23
C       1111::/64 via ::, irb1001, 00:13:29
C       2121::/64 via ::, irb2001, 00:13:06
C       fe80::/64 via ::, irb2001, 00:13:06
VTEP1#show interface ce52 counters queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes      |  Dropped pkts   |
Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
q0                (E) 125304832 0                 0                   0                 0
q1                (E) 125304832 0                 0                   0                 0
q2                (E) 125304832 0                 0                   0                 0
q3                (E) 125304832 0                 0                   0                 0
q4                (E) 125304832 0                 0                   0                 0
q5                (E) 125304832 0                 0                   0                 0
q6                (E) 125304832 0                 0                   0                 0
q7                (E) 125304832 0                 0                   0                 0

VTEP1#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
```

```
 1000         101          irb1001
 1000         201          irb2001

VTEP1#show qos-profile DSCP_QUEUE
profile name: DSCP_QUEUE
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
 dscp 20 queue 4
Detailed mapping:
 -------+---------------------------- | -------+---------------------------- | -------+--
------------------------ | -------+----------------------------
  INPUT |        OUTPUT             |  INPUT |        OUTPUT             |  INPUT |
OUTPUT              |  INPUT |         OUTPUT
 -------+---------------------------- | -------+---------------------------- | -------+--
------------------------ | -------+----------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0      0      green    0         |  16     2      green    16        |  32     4
green    32        |  48      6       green    48
  1      0      green    1         |  17     2      green    17        |  33     4
green    33        |  49      6       green    49
  2      0      green    2         |  18     2      green    18        |  34     4
green    34        |  50      6       green    50
  3      0      green    3         |  19     2      green    19        |  35     4
green    35        |  51      6       green    51
  4      0      green    4         |  20     4      yellow   20        |  36     4
yellow   36        |  52      6       green    52
  5      0      green    5         |  21     2      green    21        |  37     4
green    37        |  53      6       green    53
  6      0      green    6         |  22     2      yellow   22        |  38     4
yellow   38        |  54      6       green    54
  7      0      green    7         |  23     2      green    23        |  39     4
green    39        |  55      6       green    55
  8      1      green    8         |  24     3      green    24        |  40     5
green    40        |  56      7       green    56
  9      1      green    9         |  25     3      green    25        |  41     5
green    41        |  57      7       green    57
  10     1      green    10        |  26     3      green    26        |  42     5
green    42        |  58      7       green    58
  11     1      green    11        |  27     3      green    27        |  43     5
green    43        |  59      7       green    59
  12     1      yellow   12        |  28     3      yellow   28        |  44     5
green    44        |  60      7       green    60
  13     1      green    13        |  29     3      green    29        |  45     5
green    45        |  61      7       green    61
  14     1      yellow   14        |  30     3      yellow   30        |  46     5
green    46        |  62      7       green    62
  15     1      green    15        |  31     3      green    31        |  47     5
green    47        |  63      7       green    63

VTEP1#show qos-profile interface irb2001
profile name: DSCP_QUEUE
```

```
 profile type: dscp-to-queue (Ingress)
 mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------- | -------+---------------------------
  INPUT |        OUTPUT            |  INPUT |        OUTPUT            |  INPUT |
OUTPUT            |  INPUT |         OUTPUT
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------- | -------+---------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0      0      green    0         |  16     2      green    16        |  32     4
green    32        |  48     6       green    48
  1      0      green    1         |  17     2      green    17        |  33     4
green    33        |  49     6       green    49
  2      0      green    2         |  18     2      green    18        |  34     4
green    34        |  50     6       green    50
  3      0      green    3         |  19     2      green    19        |  35     4
green    35        |  51     6       green    51
  4      0      green    4         |  20     4      yellow   20        |  36     4
yellow   36        |  52     6       green    52
  5      0      green    5         |  21     2      green    21        |  37     4
green    37        |  53     6       green    53
  6      0      green    6         |  22     2      yellow   22        |  38     4
yellow   38        |  54     6       green    54
  7      0      green    7         |  23     2      green    23        |  39     4
green    39        |  55     6       green    55
  8      1      green    8         |  24     3      green    24        |  40     5
green    40        |  56     7       green    56
  9      1      green    9         |  25     3      green    25        |  41     5
green    41        |  57     7       green    57
  10     1      green    10        |  26     3      green    26        |  42     5
green    42        |  58     7       green    58
  11     1      green    11        |  27     3      green    27        |  43     5
green    43        |  59     7       green    59
  12     1      yellow   12        |  28     3      yellow   28        |  44     5
green    44        |  60     7       green    60
  13     1      green    13        |  29     3      green    29        |  45     5
green    45        |  61     7       green    61
  14     1      yellow   14        |  30     3      yellow   30        |  46     5
green    46        |  62     7       green    62
  15     1      green    15        |  31     3      green    31        |  47     5
green    47        |  63     7       green    63

VTEP2

VTEP2#show running-config qos
qos enable
!
qos profile dscp-to-queue DSCP_QUEUE
 dscp 20 queue 4
!
!
!
```

```
!
interface irb2001
 qos map-profile dscp-to-queue DSCP_QUEUE
!
VTEP2#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
evpn irb-forwarding anycast-gateway-mac 0000.0000.1111
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan po1 10
 map vnid 101
 mac 0000.2222.1010 ip 11.11.11.51
!
nvo vxlan access-if port-vlan po1 20
 map vnid 201
 mac 0000.2222.1020 ip 21.21.21.51
!
!
VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
2.2.2.2          4.4.4.4          Installed        01:13:43       01:13:43
2.2.2.2          1.1.1.1          Installed        01:13:43       01:13:43
2.2.2.2          5.5.5.5          Installed        01:13:43       01:13:43

Total number of entries are 3
VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
```

```
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                            VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________________
101      VNI-101      L2       NW   ----      ----                           ---- ----
2.2.2.2          4.4.4.4
101      VNI-101      L2       NW   ----      ----                           ---- ----
2.2.2.2          1.1.1.1
101      VNI-101      L2       NW   ----      ----                           ---- ----
2.2.2.2          5.5.5.5
101      VNI-101      --       AC   po1       00:00:00:00:00:22:22:00:00:00 10   NON-DF
----             ----
201      VNI-201      L2       NW   ----      ----                           ---- ----
2.2.2.2          4.4.4.4
201      VNI-201      L2       NW   ----      ----                           ---- ----
2.2.2.2          1.1.1.1
201      VNI-201      L2       NW   ----      ----                           ---- ----
2.2.2.2          5.5.5.5
201      VNI-201      --       AC   po1       00:00:00:00:00:22:22:00:00:00 20   NON-DF
----             ----

Total number of entries are 8
VTEP2#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Local   ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.101     0000.5555.1020 Static Remote  ----
101      11.11.11.51      0000.2222.1010 Static Local   ----
101      11.11.11.10      0010.9400.0002 Dynamic Remote ----
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.201     0000.4444.1010 Static Remote  ----
Total number of entries are 7
VTEP2#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr                                 Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________
_____________
201      2121::1                         0000.0000.1111 Static Local     ----
101      1111::10                        0010.9400.0002 Dynamic Remote   ----
101      1111::1                         0000.0000.1111 Static Local     ----
Total number of entries are 3
VTEP2#show ipv4 route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:10:20
C            21.21.21.0/24 is directly connected, irb2001, 00:09:55
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:13:30

Gateway of last resort is not set
VTEP2#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C       ::1/128 via ::, lo.L3VRF1, 00:13:36
C       1111::/64 via ::, irb1001, 00:10:26
C       2121::/64 via ::, irb2001, 00:10:01
C       fe80::/64 via ::, irb2001, 00:10:01
VTEP2#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001
 1000         201          irb2001

VTEP2#show interface xe29 counters queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes      |  Dropped pkts   |
Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
q0                (E) 12517376 0                 0                   0                 0
q1                (E) 12517376 0                 0                   0                 0
q2                (E) 12517376 0                 0                   0                 0
q3                (E) 12517376 0                 0                   0                 0
q4                (E) 12517376 2998022           2998045000          0                 0
q5                (E) 12517376 0                 0                   0                 0
q6                (E) 12517376 0                 0                   0                 0
q7                (E) 12517376 0                 0                   0                 0

VTEP2#show qos-profile DSCP_QUEUE
profile name: DSCP_QUEUE
```

```
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
 dscp 20 queue 4
Detailed mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------ | -------+---------------------------
  INPUT |        OUTPUT            |  INPUT |        OUTPUT            |  INPUT |
OUTPUT             |  INPUT |         OUTPUT
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------ | -------+---------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0       0       green    0         |  16      2       green    16        |  32      4
green    32         |  48       6        green    48
  1       0       green    1         |  17      2       green    17        |  33      4
green    33         |  49       6        green    49
  2       0       green    2         |  18      2       green    18        |  34      4
green    34         |  50       6        green    50
  3       0       green    3         |  19      2       green    19        |  35      4
green    35         |  51       6        green    51
  4       0       green    4         |  20      4       yellow   20        |  36      4
yellow   36         |  52       6        green    52
  5       0       green    5         |  21      2       green    21        |  37      4
green    37         |  53       6        green    53
  6       0       green    6         |  22      2       yellow   22        |  38      4
yellow   38         |  54       6        green    54
  7       0       green    7         |  23      2       green    23        |  39      4
green    39         |  55       6        green    55
  8       1       green    8         |  24      3       green    24        |  40      5
green    40         |  56       7        green    56
  9       1       green    9         |  25      3       green    25        |  41      5
green    41         |  57       7        green    57
  10      1       green    10        |  26      3       green    26        |  42      5
green    42         |  58       7        green    58
  11      1       green    11        |  27      3       green    27        |  43      5
green    43         |  59       7        green    59
  12      1       yellow   12        |  28      3       yellow   28        |  44      5
green    44         |  60       7        green    60
  13      1       green    13        |  29      3       green    29        |  45      5
green    45         |  61       7        green    61
  14      1       yellow   14        |  30      3       yellow   30        |  46      5
green    46         |  62       7        green    62
  15      1       green    15        |  31      3       green    31        |  47      5
green    47         |  63       7        green    63

VTEP2#show qos-profile interface irb2001
profile name: DSCP_QUEUE
 profile type: dscp-to-queue (Ingress)
 mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------ | -------+---------------------------
  INPUT |        OUTPUT            |  INPUT |        OUTPUT            |  INPUT |
OUTPUT             |  INPUT |         OUTPUT
```

```
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------ | -------+---------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0      0      green    0         |  16     2      green    16        |  32     4
green    32        |  48      6       green    48
  1      0      green    1         |  17     2      green    17        |  33     4
green    33        |  49      6       green    49
  2      0      green    2         |  18     2      green    18        |  34     4
green    34        |  50      6       green    50
  3      0      green    3         |  19     2      green    19        |  35     4
green    35        |  51      6       green    51
  4      0      green    4         |  20     4      yellow   20        |  36     4
yellow   36        |  52      6       green    52
  5      0      green    5         |  21     2      green    21        |  37     4
green    37        |  53      6       green    53
  6      0      green    6         |  22     2      yellow   22        |  38     4
yellow   38        |  54      6       green    54
  7      0      green    7         |  23     2      green    23        |  39     4
green    39        |  55      6       green    55
  8      1      green    8         |  24     3      green    24        |  40     5
green    40        |  56      7       green    56
  9      1      green    9         |  25     3      green    25        |  41     5
green    41        |  57      7       green    57
  10     1      green    10        |  26     3      green    26        |  42     5
green    42        |  58      7       green    58
  11     1      green    11        |  27     3      green    27        |  43     5
green    43        |  59      7       green    59
  12     1      yellow   12        |  28     3      yellow   28        |  44     5
green    44        |  60      7       green    60
  13     1      green    13        |  29     3      green    29        |  45     5
green    45        |  61      7       green    61
  14     1      yellow   14        |  30     3      yellow   30        |  46     5
green    46        |  62      7       green    62
  15     1      green    15        |  31     3      green    31        |  47     5
green    47        |  63      7       green    63
```

VTEP4

```
VTEP4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source             Destination       Status            Up/Down         Update
================================================================================
4.4.4.4            2.2.2.2           Installed         01:14:47        01:14:47
4.4.4.4            1.1.1.1           Installed         01:18:35        01:18:35
4.4.4.4            5.5.5.5           Installed         01:18:20        01:18:20

Total number of entries are 3
VTEP4#show nvo vxlan
VXLAN Information
================
   Codes: NW - Network Port
           AC - Access Port
          (u) - Untagged
```

```
VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________________
101      VNI-101      L2       NW   ----      ----                         ---- ----
4.4.4.4          2.2.2.2
101      VNI-101      L2       NW   ----      ----                         ---- ----
4.4.4.4          1.1.1.1
101      VNI-101      L2       NW   ----      ----                         ---- ----
4.4.4.4          5.5.5.5
201      VNI-201      L2       NW   ----      ----                         ---- ----
4.4.4.4          2.2.2.2
201      VNI-201      L2       NW   ----      ----                         ---- ----
4.4.4.4          1.1.1.1
201      VNI-201      L2       NW   ----      ----                         ---- ----
4.4.4.4          5.5.5.5
201      VNI-201      --       AC   xe3       --- Single Homed Port ---     20   ----
----             ----

Total number of entries are 7
VTEP4#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.101     0000.5555.1020 Static Local   ----
101      11.11.11.51      0000.2222.1010 Static Remote  ----
101      11.11.11.10      0010.9400.0002 Dynamic Remote ----
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.201     0000.4444.1010 Static Remote  ----
Total number of entries are 7
VTEP4#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________________
_____________
201      2121::1                          0000.0000.1111 Static Local     ----
101      1111::10                         0010.9400.0002 Dynamic Remote   ----
101      1111::1                          0000.0000.1111 Static Local     ----
Total number of entries are 3
VTEP4#show ipv4 route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
```

```
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:07:24
C            21.21.21.0/24 is directly connected, irb2001, 00:07:14
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:10:16

Gateway of last resort is not set
VTEP4#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C       ::1/128 via ::, lo.L3VRF1, 00:10:21
C       1111::/64 via ::, irb1001, 00:07:29
C       2121::/64 via ::, irb2001, 00:07:19
C       fe80::/64 via ::, irb2001, 00:07:19
VTEP4#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001
 1000         201          irb2001
```

VTEP5

```
VTEP5#show running-config qos
qos enable
!
qos profile queue-color-to-cos QUEUE_COS
 queue 2 color all cos 6
!
qos profile dscp-to-queue DSCP_QUEUE
 dscp 34 queue 2
!
!
!
VTEP5#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
evpn irb-forwarding anycast-gateway-mac 0000.0000.1111
```

```
!
nvo vxlan vtep-ip-global 5.5.5.5
!
nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP_QUEUE
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
 vni-name VNI-201
!
nvo vxlan access-if port-vlan xe1 10
 map vnid 101
 mac 0000.4444.1010 ip 11.11.11.201
 map qos-profile queue-color-to-cos QUEUE_COS
!
!
VTEP5#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
5.5.5.5          2.2.2.2          Installed        01:11:17       01:11:17
5.5.5.5          4.4.4.4          Installed        01:14:50       01:14:50
5.5.5.5          1.1.1.1          Installed        01:14:50       01:14:50

Total number of entries are 3
VTEP5#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
           AC - Access Port
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
--------------------------------------------------------------------------------
--------------------------------------
101      VNI-101      L2       NW   ----      ----                         ---- ----
5.5.5.5          2.2.2.2
101      VNI-101      L2       NW   ----      ----                         ---- ----
5.5.5.5          4.4.4.4
101      VNI-101      L2       NW   ----      ----                         ---- ----
5.5.5.5          1.1.1.1
101      VNI-101      --       AC   xe1       --- Single Homed Port ---     10   ----
----             ----
201      VNI-201      L2       NW   ----      ----                         ---- ----
5.5.5.5          2.2.2.2
```

```
201      VNI-201      L2       NW   ----      ----                          ---- ----
5.5.5.5          4.4.4.4
201      VNI-201      L2       NW   ----      ----                          ---- ----
5.5.5.5          1.1.1.1

Total number of entries are 7
VTEP5#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       0000.0000.1111 Static Local   ----
201      21.21.21.101     0000.5555.1020 Static Remote  ----
101      11.11.11.51      0000.2222.1010 Static Remote  ----
101      11.11.11.10      0010.9400.0002 Dynamic Local  ----
101      11.11.11.1       0000.0000.1111 Static Local   ----
101      11.11.11.201     0000.4444.1010 Static Local   ----
Total number of entries are 7
VTEP5#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________
_____________
201      2121::1                          0000.0000.1111 Static Local     ----
101      1111::10                         0010.9400.0002 Dynamic Local    ----
101      1111::1                          0000.0000.1111 Static Local     ----
Total number of entries are 3
VTEP5#show ipv4 route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
C            11.11.11.0/24 is directly connected, irb1001, 00:05:08
C            21.21.21.0/24 is directly connected, irb2001, 00:04:57
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:07:30

Gateway of last resort is not set
VTEP5#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
```

```
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C       ::1/128 via ::, lo.L3VRF1, 00:07:35
C       1111::/24 via ::, irb1001, 00:05:13
C       2121::/64 via ::, irb2001, 00:05:02
C       fe80::/64 via ::, irb2001, 00:05:02
VTEP5#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001
 1000         201          irb2001

VTEP5#show interface xe1 counters queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes      |  Dropped pkts   |
Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
q0                 (E) 12517376 0                 0                   0                 0
q1                 (E) 12517376 0                 0                   0                 0
q2                 (E) 12517376 1170696           1224549062          0                 0
q3                 (E) 12517376 0                 0                   0                 0
q4                 (E) 12517376 0                 0                   0                 0
q5                 (E) 12517376 0                 0                   0                 0
q6                 (E) 12517376 0                 0                   0                 0
q7                 (E) 12517376 0                 0                   0                 0

VTEP5#show qos-profile QUEUE_COS
profile name: QUEUE_COS
profile type: queue-color-to-cos
profile attached to 1 instances
configured mapping:
 queue 2 color all cos 6
Detailed mapping:
 ----------------+-------- | ----------------+-------- | ----------------+--------
      INPUT      | OUTPUT  |      INPUT      | OUTPUT  |      INPUT      | OUTPUT
 ----------------+-------- | ----------------+-------- | ----------------+--------
  Queue | Color  | COS     |  Queue | Color  | COS     |  Queue | Color  | COS
 -------+--------+-------- | -------+--------+-------- | -------+--------+--------
  0       green    0       |  0       yellow   0       |  0       red      0
  1       green    1       |  1       yellow   1       |  1       red      1
  2       green    6       |  2       yellow   6       |  2       red      6
  3       green    3       |  3       yellow   3       |  3       red      3
  4       green    4       |  4       yellow   4       |  4       red      4
  5       green    5       |  5       yellow   5       |  5       red      5
  6       green    6       |  6       yellow   6       |  6       red      6
```

```
  7       green    7       |  7       yellow   7       |  7       red      7

VTEP5#show qos-profile DSCP_QUEUE
profile name: DSCP_QUEUE
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
 dscp 34 queue 2
Detailed mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------- | -------+---------------------------
  INPUT |         OUTPUT             |  INPUT |         OUTPUT             |  INPUT |
OUTPUT              |  INPUT |          OUTPUT
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------- | -------+---------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0       0       green    0         |  16      2       green    16        |  32      4
green    32         |  48       6        green    48
  1       0       green    1         |  17      2       green    17        |  33      4
green    33         |  49       6        green    49
  2       0       green    2         |  18      2       green    18        |  34      2
green    34         |  50       6        green    50
  3       0       green    3         |  19      2       green    19        |  35      4
green    35         |  51       6        green    51
  4       0       green    4         |  20      2       yellow   20        |  36      4
yellow   36         |  52       6        green    52
  5       0       green    5         |  21      2       green    21        |  37      4
green    37         |  53       6        green    53
  6       0       green    6         |  22      2       yellow   22        |  38      4
yellow   38         |  54       6        green    54
  7       0       green    7         |  23      2       green    23        |  39      4
green    39         |  55       6        green    55
  8       1       green    8         |  24      3       green    24        |  40      5
green    40         |  56       7        green    56
  9       1       green    9         |  25      3       green    25        |  41      5
green    41         |  57       7        green    57
  10      1       green    10        |  26      3       green    26        |  42      5
green    42         |  58       7        green    58
  11      1       green    11        |  27      3       green    27        |  43      5
green    43         |  59       7        green    59
  12      1       yellow   12        |  28      3       yellow   28        |  44      5
green    44         |  60       7        green    60
  13      1       green    13        |  29      3       green    29        |  45      5
green    45         |  61       7        green    61
  14      1       yellow   14        |  30      3       yellow   30        |  46      5
green    46         |  62       7        green    62
  15      1       green    15        |  31      3       green    31        |  47      5
green    47         |  63       7        green    63
```

Distributed Gateway

In distributed gateway approach, VTEP will act as default gateways for one or more VNIDs,

Each VTEP having its own default gateway IP and MAC configuration for a given VNID.

IRB QoS Configuration for Distributed

Configure from Base Configuration-L2 VXLAN section, then configure below commands for centralized distributed approach.

Note: For L3 traffic, when l3vni is sent in the traffic, then dscp-to-queue qos profile mapped at tunnel ingress takes effect.

VTEP4

Unconfigure vnid 101 from nvo vxlan.

`(config)#nvo vxlan irb`	Enable VXLAN irb
`(config)#ip vrf L3VRF1`	Create mac routing/forwarding instance with L3VRF1 name and enter into vrf mode
`(config-vrf)#rd 41000:11`	Assign RD value
`(config-vrf)# route-target both 100:100`	Assign route-target value for same for import and export.
`(config-vrf)# l3vni 1000`	Configure L3VNI as 1000 for L3VRF1
`(config-vrf)#exit`	Exit from vrf mode
`(config)# interface irb2001`	Configure IRB interface 2001
`(config-if)ip vrf forwarding L3VRF1`	Configure L3VRF1
`(config-if)ip address 21.21.21.1/24`	Configure ip address
`(config-if)ipv6 address 2121::1/64`	Configure ipv6 address
`(config-if)exit`	Exit from interface config mode
`(config)router bgp 5000`	Enter into bgp router mode
`(config-router)#address-family ipv4 vrf L3VRF1`	Enter into address-family mode for L3VRF1
`(config-router-af)#redistribute connected`	Redistribute connected
`(config-router-af)#exit-address-family`	Exit form address-family
`(config-router)#address-family ipv6 vrf L3VRF1`	Enter into address-family mode for L3VRF1
`(config-router-af)#redistribute connected`	Redistribute connected
`(config-router-af)#exit-address-family`	Exit form address-family
`(config)#nvo vxlan id 201 ingress-replication inner-vid-disabled`	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF2`	Assign vrf for evpn-bgp to carry EVPN route
`(config-nvo)# evpn irb2001`	Configure irb2001 under vxlan id 201
`(config-nvo)#exit`	Exit from VXLAN tenant mode and enter into configuration mode.

(config)#qos remark dscp	Enable qos dscp remark for dscp-to-dscp mapping if required
(config)#qos profile dscp-to-queue DSCP_QUEUE	Configure qos profile to remark at vxlan ingress tunnel. Here dscp to <queue, dscp> is not applicable. I.e ingress remarking of dscp at the ingress tunnel is not applicable.
(config-ingress-dscp-map)#dscp 56 queue 6	Configure particular outer dscp value to queue value. Here classification at the ingress vtep is based on overlay dscp value
(config-ingress-dscp-map)#exit	Exit from qos profile config mode
(config-egress-dscp-map)#dscp 20 dscp 32	Configure particular dscp value to dscp value
(config-egress-dscp-map)#exit	Exit from qos profile config mode
(config)#qos profile queue-color-to-cos QUEUE_COS	Configure qos profile for remark at vxlan access-if
(config-egress-cos-map)#queue 6 cos 2	Configure particular queue value to COS value
(config-egress-cos-map)#exit	Exit from qos profile config mode
(config)#nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP_QUEUE	Map the qos profile to vxlan tunnel ingress
(config)#nvo vxlan access-if port-vlan xe3 20	Enter vxlan access-if mode
(config-nvo-acc-if)#map qos-profile queue-color-to-cos QUEUE_COS	Map qos profile
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration
(config-nvo-acc-if)#end	Exit from global configuration mode

VTEP5

Unconfigure vnid 201 from nvo vxlan.

(config)#nvo vxlan irb	Enable VXLAN IRB
(config)#ip vrf L3VRF1	Create mac routing/forwarding instance with L3VRF1 name and enter into vrf mode
(config-vrf)#rd 51000:11	Assign RD value
(config-vrf)# route-target both 100:100	Assign route-target value for same for import and export.
(config-vrf)# l3vni 1000	Configure L3VNI as 1000 for L3VRF1
(config-vrf)#exit	Exit from vrf mode
(config)# interface irb1001	Configure IRB interface 1001
(config-if)ip vrf forwarding L3VRF1	Configure L3VRF1
(config-if)ip address 11.11.11.1/24	Configure ip address
(config-if)ipv6 address 1111::1/64	Configure ipv6 address
(config-if)exit	Exit from interface config mode
(config)router bgp 5000	Enter into bgp router mode
(config-router)#address-family ipv4 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit form address-family

(config-router)#address-family ipv6 vrf L3VRF1	Enter into address-family mode for L3VRF1
(config-router-af)#redistribute connected	Redistribute connected
(config-router-af)#exit-address-family	Exit form address-family
(config)# nvo vxlan id 101 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp L2VRF1	Assign vrf for evpn-bgp to carry EVPN route
(config-nvo)# evpn irb1001	Configure irb1001 under vxlan id 101
(config-nvo)#exit	Exit from VXLAN tenant mode and enter into configuration mode.
(config)#qos profile dscp-to-queue DSCP_QUEUE	Configure qos profile for dscp to queue for ingress traffic
(config-ingress-dscp-map)# dscp 20 queue 4	Configure particular dscp value to queue value
(config)#interface irb1001	Enter IRB L3 interface
(config-irb-if)#qos map-profile dscp-to-queue DSCP_QUEUE	Map qos profile
(config-irb-if)#commit	Commit the candidate configuration to the running configuration
(config-irb-if)#end	Exit from global conf mode

Validations

Send traffic from TS1-11 to VTEP5 access-if with dscp value 20 COS value 1(vlan10) and verify traffic received at TS2-21 with dscp value 32 and COS value 2(vlan20) at the VTEP4 access-if.

VTEP5

```
VTEP5#show running-config qos
qos enable
!
qos profile dscp-to-queue DSCP_QUEUE
 dscp 20 queue 4
!
!
!
!
interface irb1001
 qos map-profile dscp-to-queue DSCP_QUEUE
!
VTEP5#show run nvo vxlan
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
```

```
nvo vxlan vtep-ip-global 5.5.5.5
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 evpn irb1001
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 vni-name VNI-201
!
nvo vxlan access-if port-vlan xe1 10
 map vnid 101
 mac 0000.4444.1010 ip 11.11.11.201
!
!
VTEP5#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
5.5.5.5          2.2.2.2          Installed        00:29:54       00:29:54
5.5.5.5          4.4.4.4          Installed        00:29:54       00:29:54
5.5.5.5          1.1.1.1          Installed        00:29:54       00:29:54

Total number of entries are 3
VTEP5#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
           AC - Access Port
          (u) - Untagged

VNID    VNI-Name    VNI-Type Type Interface ESI                           VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
101     VNI-101     L2       NW   ----      ----                          ---- ----
5.5.5.5          2.2.2.2
101     VNI-101     L2       NW   ----      ----                          ---- ----
5.5.5.5          4.4.4.4
101     VNI-101     L2       NW   ----      ----                          ---- ----
5.5.5.5          1.1.1.1
101     VNI-101     --       AC   xe1       --- Single Homed Port ---     10   ----
----             ----
201     VNI-201     L2       NW   ----      ----                          ---- ----
5.5.5.5          2.2.2.2
201     VNI-201     L2       NW   ----      ----                          ---- ----
5.5.5.5          4.4.4.4
201     VNI-201     L2       NW   ----      ----                          ---- ----
5.5.5.5          1.1.1.1
1000    ----        L3       NW   ----      ----                          ---- ----
5.5.5.5          4.4.4.4
```

```
Total number of entries are 8
VTEP5#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       3c2c.991a.da7a Static Remote  ----
201      21.21.21.101     0000.5555.1020 Static Remote  ----
101      11.11.11.51      0000.2222.1010 Static Remote  ----
101      11.11.11.1       04f8.f82f.8eee Static Local   ----
101      11.11.11.201     0000.4444.1010 Static Local   ----
Total number of entries are 6
VTEP5#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
===========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________
_____________
201      2121::1                        3c2c.991a.da7a Static Remote    ----
101      1111::1                        04f8.f82f.8eee Static Local     ----
Total number of entries are 2
VTEP5#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1000         101          irb1001

VTEP5#show ipv4 route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
B            4.4.4.4/32 [0/0] is directly connected, tunvxlan1000, 00:04:09
C            11.11.11.0/24 is directly connected, irb1001, 00:02:38
B            21.21.21.0/24 [200/0] via 4.4.4.4 (recursive is directly connected,
tunvxlan1000), 00:04:09
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:04:46

Gateway of last resort is not set
VTEP5#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
```

```
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C       ::1/128 via ::, lo.L3VRF1, 00:04:50
B       ::ffff:404:404/128 [0/0] via ::, tunvxlan1000, 00:04:13
C       1111::/64 via ::, irb1001, 00:02:42
B       2121::/64 [200/0] via ::ffff:404:404 (recursive via ::, unknown), 00:04:13
C       fe80::/64 via ::, irb1001, 00:02:42
VTEP5#show interface ce53 counter queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts     |     Tx bytes      |  Dropped pkts   |
Dropped bytes   |
+--------------------+--------+-----------------+-------------------+-----------------
+-------------------+
q0                 (E) 50069504 0                 0                   0                 0
q1                 (E) 50069504 0                 0                   0                 0
q2                 (E) 50069504 0                 0                   0                 0
q3                 (E) 50069504 0                 0                   0                 0
q4                 (E) 50069504 498796            698314400           0                 0
q5                 (E) 50069504 0                 0                   0                 0
q6                 (E) 50069504 0                 0                   0                 0
q7                 (E) 50069504 0                 0                   0                 0

VTEP5#show qos-profile DSCP_QUEUE
profile name: DSCP_QUEUE
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
 dscp 20 queue 4
Detailed mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------- | -------+---------------------------
  INPUT |        OUTPUT            |  INPUT |        OUTPUT            |  INPUT |
OUTPUT             |  INPUT |         OUTPUT
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------- | -------+---------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0      0      green    0         |  16     2      green    16        |  32     4
green    32        |  48      6       green    48
  1      0      green    1         |  17     2      green    17        |  33     4
green    33        |  49      6       green    49
  2      0      green    2         |  18     2      green    18        |  34     4
green    34        |  50      6       green    50
  3      0      green    3         |  19     2      green    19        |  35     4
green    35        |  51      6       green    51
```

```
  4       0       green    4         |  20      4       yellow   20        |  36      4
yellow    36         |  52       6        green    52
  5       0       green    5         |  21      2       green    21        |  37      4
green     37         |  53       6        green    53
  6       0       green    6         |  22      2       yellow   22        |  38      4
yellow    38         |  54       6        green    54
  7       0       green    7         |  23      2       green    23        |  39      4
green     39         |  55       6        green    55
  8       1       green    8         |  24      3       green    24        |  40      5
green     40         |  56       7        green    56
  9       1       green    9         |  25      3       green    25        |  41      5
green     41         |  57       7        green    57
  10      1       green    10        |  26      3       green    26        |  42      5
green     42         |  58       7        green    58
  11      1       green    11        |  27      3       green    27        |  43      5
green     43         |  59       7        green    59
  12      1       yellow   12        |  28      3       yellow   28        |  44      5
green     44         |  60       7        green    60
  13      1       green    13        |  29      3       green    29        |  45      5
green     45         |  61       7        green    61
  14      1       yellow   14        |  30      3       yellow   30        |  46      5
green     46         |  62       7        green    62
  15      1       green    15        |  31      3       green    31        |  47      5
green     47         |  63       7        green    63

VTEP5#show qos-profile interface irb1001
profile name: DSCP_QUEUE
 profile type: dscp-to-queue (Ingress)
 mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------- | -------+---------------------------
  INPUT |        OUTPUT              |  INPUT |        OUTPUT              |  INPUT |
OUTPUT              |  INPUT |         OUTPUT
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------- | -------+---------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0       0       green    0         |  16      2       green    16        |  32      4
green     32         |  48       6        green    48
  1       0       green    1         |  17      2       green    17        |  33      4
green     33         |  49       6        green    49
  2       0       green    2         |  18      2       green    18        |  34      4
green     34         |  50       6        green    50
  3       0       green    3         |  19      2       green    19        |  35      4
green     35         |  51       6        green    51
  4       0       green    4         |  20      4       yellow   20        |  36      4
yellow    36         |  52       6        green    52
  5       0       green    5         |  21      2       green    21        |  37      4
green     37         |  53       6        green    53
  6       0       green    6         |  22      2       yellow   22        |  38      4
yellow    38         |  54       6        green    54
  7       0       green    7         |  23      2       green    23        |  39      4
green     39         |  55       6        green    55
```

```
  8       1       green    8         |  24      3       green    24        |  40      5
green     40         |  56       7        green     56
  9       1       green    9         |  25      3       green    25        |  41      5
green     41         |  57       7        green     57
  10      1       green    10        |  26      3       green    26        |  42      5
green     42         |  58       7        green     58
  11      1       green    11        |  27      3       green    27        |  43      5
green     43         |  59       7        green     59
  12      1       yellow   12        |  28      3       yellow   28        |  44      5
green     44         |  60       7        green     60
  13      1       green    13        |  29      3       green    29        |  45      5
green     45         |  61       7        green     61
  14      1       yellow   14        |  30      3       yellow   30        |  46      5
green     46         |  62       7        green     62
  15      1       green    15        |  31      3       green    31        |  47      5
green     47         |  63       7        green     63
```

VTEP4

```
VTEP4#show run qos
qos enable
qos remark dscp
!
!
qos profile queue-color-to-cos QUEUE_COS
 queue 6 color all cos 2
!
qos profile dscp-to-queue DSCP_QUEUE
 dscp 56 queue 6
!
!
!
!
VTEP4#show run nvo vxlan
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
nvo vxlan vtep-ip-global 4.4.4.4
!
nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP_QUEUE
!
nvo vxlan id 101 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF1
 vni-name VNI-101
!
nvo vxlan id 201 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp L2VRF2
 evpn irb2001
```

```
 vni-name VNI-201
!
nvo vxlan access-if port-vlan xe3 20
 map vnid 201
 mac 0000.5555.1020 ip 21.21.21.101
 map qos-profile queue-color-to-cos QUEUE_COS
!
!
VTEP4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
4.4.4.4          2.2.2.2          Installed        00:28:32       00:28:32
4.4.4.4          1.1.1.1          Installed        00:28:38       00:28:38
4.4.4.4          5.5.5.5          Installed        00:28:31       00:28:31

Total number of entries are 3
VTEP4#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
           AC - Access Port
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
101      VNI-101      L2       NW   ----      ----                         ---- ----
4.4.4.4          2.2.2.2
101      VNI-101      L2       NW   ----      ----                         ---- ----
4.4.4.4          1.1.1.1
101      VNI-101      L2       NW   ----      ----                         ---- ----
4.4.4.4          5.5.5.5
201      VNI-201      L2       NW   ----      ----                         ---- ----
4.4.4.4          2.2.2.2
201      VNI-201      L2       NW   ----      ----                         ---- ----
4.4.4.4          1.1.1.1
201      VNI-201      L2       NW   ----      ----                         ---- ----
4.4.4.4          5.5.5.5
201      VNI-201      --       AC   xe3       --- Single Homed Port ---     20   ----
----             ----
1000     ----         L3       NW   ----      ----                         ---- ----
4.4.4.4          5.5.5.5

Total number of entries are 8
VTEP4#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      21.21.21.51      0000.2222.1020 Static Remote  ----
201      21.21.21.1       3c2c.991a.da7a Static Local   ----
```

```
201      21.21.21.101     0000.5555.1020 Static Local   ----
101      11.11.11.51      0000.2222.1010 Static Remote  ----
101      11.11.11.1       04f8.f82f.8eee Static Remote  ----
101      11.11.11.201     0000.4444.1010 Static Remote  ----
Total number of entries are 6
VTEP4#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
==========================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________________
_____________
201      2121::1                        3c2c.991a.da7a Static Local     ----
101      1111::1                        04f8.f82f.8eee Static Remote    ----
Total number of entries are 2
VTEP4#show ipv4 route vrf L3VRF1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "L3VRF1"
B            5.5.5.5/32 [0/0] is directly connected, tunvxlan1000, 00:01:03
B            11.11.11.0/24 [200/0] via 5.5.5.5 (recursive is directly connected,
tunvxlan1000), 00:01:03
C            21.21.21.0/24 is directly connected, irb2001, 00:03:31
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:15:13

Gateway of last resort is not set
VTEP4#show ipv6 route vrf L3VRF1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "L3VRF1"
C       ::1/128 via ::, lo.L3VRF1, 00:15:18
B       ::ffff:505:505/128 [0/0] via ::, tunvxlan1000, 00:01:08
B       1111::/64 [200/0] via ::ffff:505:505 (recursive via ::, unknown), 00:01:08
C       2121::/64 via ::, irb2001, 00:03:36
C       fe80::/64 via ::, irb2001, 00:03:36
VTEP4#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
```

```
 1000         201          irb2001

VTEP4#show interface xe3 counter queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+--------------------+--------+----------------+------------------+----------------
+-------------------+
|   Queue/Class-map  | Q-Size |     Tx pkts    |     Tx bytes     |  Dropped pkts   |
Dropped bytes   |
+--------------------+--------+----------------+------------------+----------------
+-------------------+
q0                (E) 12517376 0                0                  0                 0
q1                (E) 12517376 0                0                  0                 0
q2                (E) 12517376 0                0                  0                 0
q3                (E) 12517376 0                0                  0                 0
q4                (E) 12517376 0                0                  0                 0
q5                (E) 12517376 0                0                  0                 0
q6                (E) 12517376 2219303          3209162748         0                 0
q7                (E) 12517376 0                0                  0                 0

VTEP4#show qos-profile QUEUE_COS
profile name: QUEUE_COS
profile type: queue-color-to-cos
profile attached to 1 instances
configured mapping:
 queue 6 color all cos 2
Detailed mapping:
 ----------------+-------- | ----------------+-------- | ----------------+--------
      INPUT      | OUTPUT  |      INPUT      | OUTPUT  |      INPUT      | OUTPUT
 ----------------+-------- | ----------------+-------- | ----------------+--------
  Queue | Color  | COS     |  Queue | Color  | COS     |  Queue | Color  | COS
 -------+--------+-------- | -------+--------+-------- | -------+--------+--------
  0       green    0       |  0       yellow   0       |  0       red      0
  1       green    1       |  1       yellow   1       |  1       red      1
  2       green    2       |  2       yellow   2       |  2       red      2
  3       green    3       |  3       yellow   3       |  3       red      3
  4       green    4       |  4       yellow   4       |  4       red      4
  5       green    5       |  5       yellow   5       |  5       red      5
  6       green    2       |  6       yellow   2       |  6       red      2
  7       green    7       |  7       yellow   7       |  7       red      7

VTEP4#show qos-profile DSCP_QUEUE
profile name: DSCP_QUEUE
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
 dscp 56 queue 6
Detailed mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------ | -------+---------------------------
  INPUT |        OUTPUT            |  INPUT |        OUTPUT            |  INPUT |
OUTPUT             |  INPUT |         OUTPUT
```

```
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------ | -------+---------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0      0      green    0        | 16     2      green    16       | 32     4
green    32       |  48      6       green    48
  1      0      green    1        | 17     2      green    17       | 33     4
green    33       |  49      6       green    49
  2      0      green    2        | 18     2      green    18       | 34     4
green    34       |  50      6       green    50
  3      0      green    3        | 19     2      green    19       | 35     4
green    35       |  51      6       green    51
  4      0      green    4        | 20     2      yellow   20       | 36     4
yellow   36       |  52      6       green    52
  5      0      green    5        | 21     2      green    21       | 37     4
green    37       |  53      6       green    53
  6      0      green    6        | 22     2      yellow   22       | 38     4
yellow   38       |  54      6       green    54
  7      0      green    7        | 23     2      green    23       | 39     4
green    39       |  55      6       green    55
  8      1      green    8        | 24     3      green    24       | 40     5
green    40       |  56      6       green    56
  9      1      green    9        | 25     3      green    25       | 41     5
green    41       |  57      7       green    57
  10     1      green    10       | 26     3      green    26       | 42     5
green    42       |  58      7       green    58
  11     1      green    11       | 27     3      green    27       | 43     5
green    43       |  59      7       green    59
  12     1      yellow   12       | 28     3      yellow   28       | 44     5
green    44       |  60      7       green    60
  13     1      green    13       | 29     3      green    29       | 45     5
green    45       |  61      7       green    61
  14     1      yellow   14       | 30     3      yellow   30       | 46     5
green    46       |  62      7       green    62
  15     1      green    15       | 31     3      green    31       | 47     5
green    47       |  63      7       green    63

VTEP4#show qos-profile interface irb2001
profile name: default
 profile type: dscp-to-queue (Ingress)
 mapping:
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------ | -------+---------------------------
  INPUT |        OUTPUT             |  INPUT |        OUTPUT             |  INPUT |
OUTPUT              |  INPUT |         OUTPUT
 -------+--------------------------- | -------+--------------------------- | -------+--
------------------------ | -------+---------------------------
  DSCP  | Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP  |  DSCP  |
Queue | Color  | Out DSCP  |  DSCP  | Queue | Color  | Out DSCP
 -------+-------+--------+---------- | -------+-------+--------+---------- | -------+--
-----+--------+---------- | -------+-------+--------+----------
  0      0      green    0        | 16     2      green    16       | 32     4
green    32       |  48      6       green    48
```

```
  1      0      green    1        |  17     2      green    17       |  33     4
green    33        |  49      6       green    49
  2      0      green    2        |  18     2      green    18       |  34     4
green    34        |  50      6       green    50
  3      0      green    3        |  19     2      green    19       |  35     4
green    35        |  51      6       green    51
  4      0      green    4        |  20     2      yellow   20       |  36     4
yellow   36        |  52      6       green    52
  5      0      green    5        |  21     2      green    21       |  37     4
green    37        |  53      6       green    53
  6      0      green    6        |  22     2      yellow   22       |  38     4
yellow   38        |  54      6       green    54
  7      0      green    7        |  23     2      green    23       |  39     4
green    39        |  55      6       green    55
  8      1      green    8        |  24     3      green    24       |  40     5
green    40        |  56      7       green    56
  9      1      green    9        |  25     3      green    25       |  41     5
green    41        |  57      7       green    57
  10     1      green    10       |  26     3      green    26       |  42     5
green    42        |  58      7       green    58
  11     1      green    11       |  27     3      green    27       |  43     5
green    43        |  59      7       green    59
  12     1      yellow   12       |  28     3      yellow   28       |  44     5
green    44        |  60      7       green    60
  13     1      green    13       |  29     3      green    29       |  45     5
green    45        |  61      7       green    61
  14     1      yellow   14       |  30     3      yellow   30       |  46     5
green    46        |  62      7       green    62
  15     1      green    15       |  31     3      green    31       |  47     5
green    47        |  63      7       green    63
```

CHAPTER 7 Single Home VxLAN IRB with OSPF or ISIS

Overview

Single Home Virtual Extensible LAN (VxLAN) with Integrated Routing (IRB) using Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (ISIS) protocols provides the solution for connecting and managing virtual networks within a data center or network infrastructure.

This feature offers a solution for networks where the interconnection of VLANs is required. These protocols can be configured on IRB interfaces within layer 3 switches or routers. This configuration enables dynamic routing, facilitating the exchange of routing information with other devices in the network. By assigning IP addresses to the IRB interfaces, they serve as the default gateways for devices within the respective VLANs.

Both OSPF and ISIS routing updates are dynamically exchanged over IRB interfaces, ensuring up-to-date routing tables and optimized traffic routing across different VLANs and networks.

This feature offers flexibility in configuring network topologies, and ensures compatibility and interoperability within diverse network environments.

Feature Characteristics

The OSPF and ISIS support over the IRB Interface feature has the following characteristics:

- Enables the control of Receive (RX)/ Transmit (TX) of OSPF and ISIS packets on IRB interfaces, providing effective management of IRB interfaces interactions with OSPF and ISIS for optimized network communication and routing.
- IRB interfaces process configured MTU size packets.
- Maintains consistency in CLI commands with SVI interfaces for OSPF and ISIS configurations, simplifying network management tasks.

Benefits

The OSPF and ISIS support over the IRB Interface has the following benefits:

- Enables seamless inter-subnet communication across different VNIDs and subnets within the same customer network.
- Promotes seamless connectivity between devices, irrespective of whether they are connected through IRB or SVI interfaces, and simplifies network management.
- The network gains greater adaptability to various scenarios and evolving requirements, offering greater versatility in its operations.

Prerequisites

- Router must be up and running.
- Maintain synchronization with VRF changes by performing IRB shut/no shut actions when specific events occur within the IPVRF. These events may involve adding or removing Route Targets (RTs), updating Route Distinguishers (RDs), or modifying Layer 3 Virtual Network Identifiers (L3VNIs).

Topology for OSPF

The network topology includes various network elements such as routers, customer edge (CE) devices, Service Aggregator (SA) devices, and Provider Edge (PE) routers. The feature enables OSPF on the IRB interfaces, allowing for efficient routing and communication between network devices within the topology.

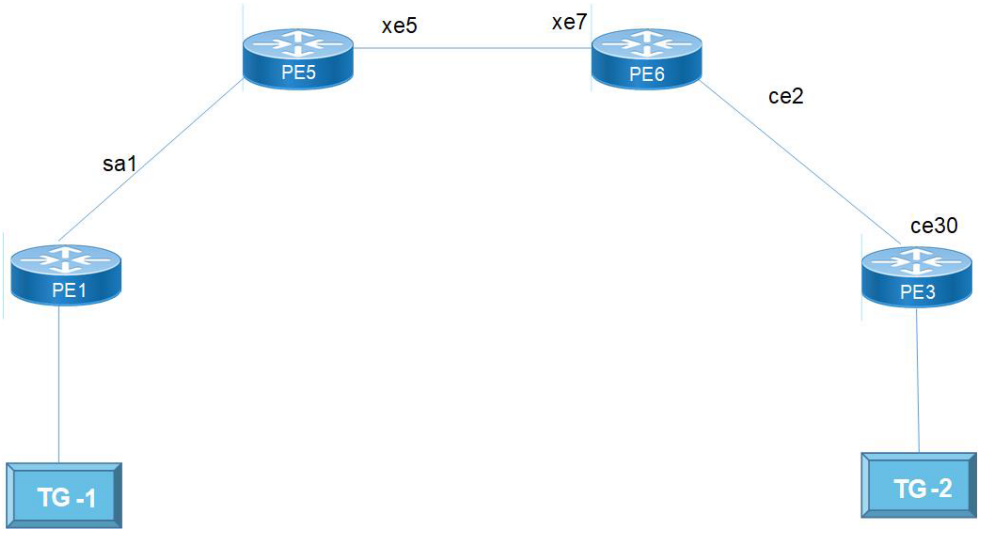

Single Home VxLAN IRB with OSPF

Configuration

Perform the following configurations to set up different interfaces, routing protocols, and BGP parameters to enable VXLAN, IRB, and EVPN functionality in the network.

Configure OSPF

PE1

`PEI(Config)# terminal`	Enters the configuration mode.
`PE1(config)#interface sa1`	Configure the sa1 interface as a network interface.
`PE1(config-if)# ip address 10.1.1.1/24`	Assigns an IP address to the sa1 interface with a subnet mask of /24.
`PE1(config-if)# ip ospf cost 10`	Configures the OSPF cost for the sa1 interface, setting it to 10.
`PE1(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa1 interface.
`PE1(config)#interface xe1`	Enters the interface xe1 mode.
`PE1(config-if)# static-channel-group 1`	Assigns the static channel group 1 to the xe1 interface.
`PE1(config-irb-if)#interface lo`	Configures the loopback (lo) interface.
`PE1(config-if)# ip address 1.1.1.1/32 secondary`	Assigns the primary IP address 1.1.1.1/32 to the loopback interface and specifies it as secondary.

`PE1(config)#router ospf 1`	Enters the OSPF configuration mode for OSPF process 1.
`PE1(config-router)# ospf router-id 1.1.1.1`	Sets the OSPF router ID to 1.1.1.1 for OSPF process 1.
`PE1(config-router)# network 1.1.1.1/32 area 0.0.0.0`	Advertises the network 1.1.1.1/32 into OSPF area 0.0.0.0.
`PE1(config-router)# network 10.1.1.0/24 area 0.0.0.0`	Advertises the network 10.1.1.0/24 into OSPF area 0.0.0.0.
`PE1(config)#nvo vxlan enable`	Enables the VXLAN feature on the device.
`PE1(config)#nvo vxlan irb`	Enables VXLAN IRB functionality.
`PE1(config-vrf)#mac vrf L2VRF1`	Configures a MAC VRF named L2VRF1.
`PE1(config-vrf)# rd 1.1.1.1:11`	Sets the Route Distinguisher (RD) to 1.1.1.1:11 for the VRF.
`PE1(config-vrf)# route-target both 9.9.9.9:100`	Configures both import and export route targets for the VRF.
`PE1(config-vrf)#ip vrf L3VRF1`	Configures an IP VRF named L3VRF1.
`PE1(config-vrf)# rd 51000:11`	Sets the RD value to 51000:11 for the L3VRF1.
`PE1(config-vrf)# route-target both 100:100`	Configures both import and export route targets for L3VRF1.
`PE1(config-vrf)# l3vni 1000`	Configures the L3 Virtual Network Identifier (L3VNI) with the value 1000.
`PE1(config)#interface irb1001`	Configures the IRB interface for L3VRF1.
`PE1(config-irb-if)# ip vrf forwarding L3VRF1`	Assigns the L3VRF1 to the IRB interface.
`PE1(config-irb-if)# ip address 11.11.11.1/24`	Assigns an IP address 11.11.11.1/24 to the IRB interface.
`PE1(config-irb)#interface irb2001`	Configures the IRB interface for IPv6 in L3VRF1.
`PE1(config-irb-if)# ip vrf forwarding L3VRF1`	Assigns the L3VRF1 to the IPv6 IRB interface.
`PE1(config-irb-if)# ipv6 address 2001::1/64`	Assigns an IP address 11.11.11.1/24 to the IRB interface.
`PE1(config-irb-if)#mtu 9000`	Sets the Maximum Transmission Unit (MTU) for this IRB interface to 9000 bytes.
`PE1(config-router)#router ospf 2 L3VRF1`	Configures OSPF on the L3VRF1.
`PE1(config-router)# network 11.11.11.0/24 area 0.0.0.0`	Advertises the network 11.11.11.0/24 into OSPF area 0.0.0.0.
`PE1(config-router)#router ipv6 vrf ospf L3VRF1`	Configures OSPFv3 on the L3VRF1.
`PE1(config-router)# router-id 1.1.1.1`	Configures the router ID as 1.1.1.1.
`PE1(config-irb)#interface irb2001`	Configures the IPv6 IRB interface.
`PE1(config-irb-if)# ipv6 router ospf area 0.0.0.0 tag L3VRF1 instance-id 0`	Attaches the OSPFv3 instance ID to the IPv6 IRB interface.
`PE1(config)#nvo vxlan vtep-ip-global 1.1.1.1`	Configures the global VTEP IP address as 1.1.1.1.
`PE1(config)#nvo vxlan id 101 ingress-replication`	Configures the VXLAN ID as 101 for ingress replication.
`PE1(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Maps the EVPN-BGP host reachability protocol to L2VRF1.
`PE1(config-nvo)# evpn irb1001`	Maps the IRB interface 1001 to EVPN.
`PE1(config-nvo)# vni-name VNI-101`	Configures the VNI name as VNI-101.
`PE1(config)#nvo vxlan id 2001 ingress-replication`	Configures the VXLAN ID as 2001 for ingress replication.
`PE1(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Maps the EVPN-BGP host reachability protocol to L2VRF1.

`PE1(config-nvo)# evpn irb2001`	Maps the IPv6 IRB interface to EVPN.
`PE1(config)#interface xe2`	Configures the xe2 interface.
`PE1(config-if)# switchport`	Configures the port as a Layer 2 (L2) switchport.
`PE1(config-if)# load-interval 30`	Configures the load-interval of 30 minutes for monitoring traffic on the xe2 interface.
`PE1(config)#nvo vxlan access-if port-vlan xe2 100`	Configures a VxLAN network virtualization overlay (NVO) on the interface xe2 with VLAN ID 100
`PE1(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN Network Identifier (VNID) 101.
`PE1(config-nvo-acc-if)#nvo vxlan access-if port-vlan xe2 2001`	Configures another VxLAN NVO on the same interface xe2, but this time with VLAN ID 2001
`PE1(config-nvo-acc-if)# map vnid 2001`	Maps VLAN 2001 to a different VxLAN VNID.
`PE1(config-router)#router bgp 100`	Configures the BGP process with AS number 100.
`PE1(config-router)# bgp router-id 1.1.1.1`	Assigns the router ID as 1.1.1.1 for the BGP instance.
`PE1(config-router)# neighbor 4.4.4.4 remote-as 100`	Configures neighbor 4.4.4.4 with a remote AS number of 100.
`PE1(config-router)# neighbor 4.4.4.4 update-source lo`	Configures the update source for neighbor 4.4.4.4 to be the loopback interface.
`PE1(config-router)# neighbor 4.4.4.4 advertisement-interval 0`	Configures the advertisement interval for neighbor 4.4.4.4 as 0.
`PE1(config-router)# address-family l2vpn evpn`	Configures the address-family for L2VPN EVPN.
`PE1(config-router-af)# neighbor 4.4.4.4 activate`	Activates the neighbor for the L2VPN EVPN address-family.
`PE1(config-router-af)# exit-address-family`	Exits from the address family configuration.
`PE1(config-router)# address-family ipv4 vrf L3VRF1`	Configures the IPv4 address-family for VRF L3VRF1.
`PE1(config-router-af)# redistribute connected`	Configures the redistribution of connected routes within the IPv4 address-family.
`PE1(config-router-af)# exit-address-family`	Exits the IPv4 address-family configuration.
`PE1(config-router)# address-family ipv6 vrf L3VRF1`	Configures the IPv6 address-family for VRF L3VRF1.
`PE1(config-router-af)# redistribute connected`	Configures the redistribution of connected routes within the IPv6 address-family.
`PE1(config-router-af)# exit-address-family`	Exits the IPv6 address-family configuration.

PE5

`PE5#configure terminal`	Enters the configuration mode
`PE5(config)#interface sa1`	Configure the sa1 interface as a network interface.
`PE5(config-if)# ip address 10.1.1.1/24`	Assigns an IP address to the sa1 interface with a subnet mask of /24.
`PE5(config-if)# ip ospf cost 10`	Configures the OSPF cost for the sa1 interface, setting it to 10.
`PE5(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa1 interface.
`PE5(config)#interface xe1`	Configure network interface towards PE6.
`PE5(config-if)# static-channel-group 1`	Assigns the static channel group 1 to the xe1 interface.

PE5(config)#interface xe5	configures the xe5 interface.
PE5(config-if)#ip address 30.1.1.1/24	Assigns the primary IP address 1.1.1.1/32 to the loopback interface and specifies it as secondary.
PE5(config)#ip ospf cost 10	Configures the OSPF cost for the xe5 interface, setting it to 10.
PE5(config-router)# ospf router-id 1.1.1.1	Assigns an IP address (30.1.1.1) to the xe5 interface with a subnet mask of /24.
PE5(config)#load-interval 30	Configures the load-interval for monitoring traffic on the xe5 interface.
PE5(config)#router ospf 1	Enters the OSPF configuration mode for OSPF process 1.
PE5(config-router)# network 30.1.1.0/24 area 0.0.0.0	Advertises the network 30.1.1.0/24 into OSPF area 0.0.0.0.
PE5(config-router)# network 10.1.1.0/24 area 0.0.0.0	Advertises the network 10.1.1.0/24 into OSPF area 0.0.0.0.

PE3

PE3#configure terminal	Enters the configuration mode
PE3(config)#interface ce30	Configure the ce30 interface as a network interface.
PE3(config-if)# ip address 40.1.1.2/24	Assigns an IP address to the ce30 interface with a subnet mask of /24.
PE3(config-if)# ip ospf cost 10	Configures the OSPF cost for the sa1 interface, setting it to 10.
PE3(config-if)# load-interval 30	Configures the load-interval for monitoring traffic on the sa1 interface.
PE3(config)#interface lo	Configure the loopback interface.
PE3(config-if)#ip address 4.4.4.4/32 secondary	Assign an secondary IP to an loopback interface.
PE3(config)#ip ospf cost 10	Configures the OSPF cost for the xe7interface, setting it to 10.
PE3(config)#load-interval 30	Configures the load-interval for monitoring traffic on the xe5 interface.
PE3(config)#router ospf 1	Enters the OSPF configuration mode for OSPF process 1.
PE3(config-router)# ospf router-id 4.4.4.4	Configures the router id to an ospf instance.
PE3(config-router)# network 4.4.4.4/32 area 0.0.0.0	Advertises the loopback address.
PE3(config-router)# network 40.1.1.0/24 area 0.0.0.0	Advertises the network interface IP address.
PE3(config)#nvo vxlan enable	Enables VXLAN on the device, allowing it to participate in VXLAN networks.
PE3(config)#nvo vxlan irb	Enables VXLAN IRB functionality, that allows routing between VXLAN and non-VXLAN networks.
PE3(config-vrf)#mac vrf L2VRF1	Configures a L2 MAC VRF instance named L2VRF1, which is a logical network segment for L2 traffic isolation.
PE3(config-vrf)# rd 4.4.4.4:11	Configures a RD for the L2VRF1, with the value 4.4.4.4:11.
PE3(config-vrf)# route-target both 9.9.9.9:100	Configures a route target for the VRF.

`PE3(config-vrf)#ip vrf L3VRF1`	Configures a L3 VRF named L3VRF1.
`PE3(config-vrf)# rd 56000:11`	Configures a RD for the L3VRF1, with the value 56000:11.
`PE3(config-vrf)# route-target both 100:100`	Configures a route target for the VRF.
`PE3(config-vrf)# l3vni 1000`	Configures a L3VNI with the ID 1000 for the VRF.
`PE3(config)#interface irb1001`	Configures the IRB interface with the ID 1001.
`PE3(config-irb-if)# ip vrf forwarding L3VRF1`	Associates the IRB interface with the L3VRF1, ensuring that traffic from this interface is isolated within that VRF.
`PE3(config-irb-if)# ip address 12.12.12.1/24`	Assigns an IP address 12.12.12.1 with a subnet mask of /24 to the IRB interface, enabling it for L3 routing.
`PE3(config-irb-if)# mtu 1500`	Configures the MTU for the interface irb1001 to 1500 bytes.
`PE3(config)#interface irb2001`	Configures another IRB interface with the ID 2001.
`PE3(config-irb-if)# ip vrf forwarding L3VRF1`	Associates the IRB interface with the L3VRF1.
`PE3(config-irb-if)# ipv6 address 2002::1/64`	Assigns an IPv6 address 2002::1 with a subnet mask of /64 to the IRB interface, enabling it for IPv6 routing.
`PE3(config-irb-if)# mtu 1500`	Configures the MTU for the interface irb2001 to 1500 bytes.
`PE3(config-router)#router ospf 2 L3VRF1`	Configures the OSPF routing process on OSPF instance 2 for the L3VRF1.
`PE3(config-router)# network 12.12.12.0/24 area 0.0.0.0`	Advertises the network 12.12.12.0/24 to OSPF area 0.0.0.0.
`PE3(config-router)#router ipv6 vrf ospf L3VRF1`	Configures the OSPFv3 routing process on OSPFv3 instance for the L3VRF1.
`PE3(config-router)# router-id 4.4.4.4`	Sets the router ID for the OSPF/OSPFv3 instances to 4.4.4.4.
`PE3(config)#nvo vxlan vtep-ip-global 4.4.4.4`	Configures the global VTEP IP address as 4.4.4.4 for VXLAN.
`PE3(config)#nvo vxlan id 102 ingress-replication`	Configures the VXLAN with VNI ID 102 for ingress replication.
`PE3(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Maps the VXLAN configuration with the EVPN-BGP protocol and associates it with the L2VRF1.
`PE3(config-nvo)# evpn irb1001`	Maps the IRB interface irb1001 to the VXLAN.
`PE3(config-nvo)# vni-name VNI-101`	Configures the VNI name as VNI-101.
`PE3(config)#nvo vxlan id 2002 ingress-replication`	Configures another VXLAN with VNI ID 2002 for ingress replication.
`PE3(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Maps the VXLAN configuration with the EVPN-BGP protocol and associates it with the L2VRF1.
`PE3(config-nvo)# evpn irb2001`	Maps the IPv6 IRB interface irb2001 to the VXLAN.
`PE3(config)#interface sa4`	Configures interface sa4.
`PE3(config-if)# switchport`	Configures the interface as a switchport.
`PE3(config-if)# load-interval 30`	Sets the load interval for the interface to 30 seconds.
`PE3(config-if)# mtu 1500`	Configures the MTU for the interface to 1500 bytes.
`PE3(config)#interface xe1`	Configures interface xe1.
`PE3(config-if)# static-channel-group 4`	Assigns a static channel group to interface xe1.
`PE3(config)#nvo vxlan access-if port-vlan sa4 100`	Configures a VxLAN nNVO on the interface xe2 with VLAN ID 100

PE3(config-nvo-acc-if)# map vnid 101	Maps VLAN 100 to the VxLAN VNID 101.
PE3(config-nvo-acc-if)#nvo vxlan access-if port-vlan sa4 2001	Configures another VxLAN NVO on the same interface xe2.
PE3(config-nvo-acc-if)# map vnid 2001	Maps VLAN 2001 to a different VxLAN VNID, in this case, VNID 2001.
PE3(config-router)#router bgp 100	Configures the BGP with AS number 100.
PE3(config-router)# bgp router-id 4.4.4.4	Sets the BGP router ID to 4.4.4.4.
PE3(config-router)# neighbor 1.1.1.1 remote-as 100	Configures a BGP neighbor with the remote AS number 100 and the IP address 1.1.1.1.
PE3(config-router)# neighbor 1.1.1.1 update-source lo	Specifies the BGP neighbor to use the loopback interface as the source for updates.
PE3(config-router)# neighbor 1.1.1.1 advertisement-interval 0	Configures the advertisement interval for BGP neighbor updates.
PE3(config-router)# address-family l2vpn evpn	Configures the BGP address family for Layer 2 VPN EVPN.
PE3(config-router-af)# neighbor 1.1.1.1 activate	Activates the BGP neighbor for the specified address family.
PE3(config-router-af)# exit-address-family	Exits the BGP address family configuration.
PE3(config-router)# address-family ipv4 vrf L3VRF1	Configures the BGP address family for IPv4 within VRF L3VRF1.
PE3(config-router-af)# redistribute connected	Configures BGP to redistribute connected routes into the BGP process.
PE3(config-router-af)# exit-address-family	Exits the BGP address family configuration for IPv4.
PE3(config-router)# address-family ipv6 vrf L3VRF1	Configures the BGP address family for IPv6 within VRF L3VRF1.
PE3(config-router-af)# redistribute connected	Configures BGP to redistribute connected routes into the BGP process.
PE3(config-router-af)# exit-address-family	Exits the BGP address family configuration for IPv6.

PE6

PE6#configure terminal	Enters the configuration mode.
PE6(config)#interface ce2	Configure the ce2 interface as a network interface.
PE6(config-if)# ip address 10.1.1.1/24	Assigns an IP address to the sa1 interface with a subnet mask of /24.
PE6(config-if)# ip ospf cost 10	Configures the OSPF cost for the sa1 interface, setting it to 10.
PE6(config-if)# load-interval 30	Configures the load-interval for monitoring traffic on the sa1 interface.
PE6(config)#interface xe7	Configure network interface towards PE5.
PE6(config-if)# static-channel-group 1	Assigns the static channel group 1 to the xe1 interface.
PE6(config-if)#ip address 30.1.1.1/24	Assign IP address to network interface.
PE6(config)#ip ospf cost 10	Configures the OSPF cost for the xe7interface, setting it to 10.

`PE6(config)#load-interval 30`	Configures the load-interval for monitoring traffic on the xe5 interface.
`PE6(config)#router ospf 1`	Enters the OSPF configuration mode for OSPF process 1.
`PE6(config-router)# network 30.1.1.0/24 area 0.0.0.0`	Advertises the network 30.1.1.0/24 into OSPF area 0.0.0.0.
`PE6(config-router)# network 40.1.1.0/24 area 0.0.0.0`	Advertises the network 40.1.1.0/24 into OSPF area 0.0.0.0.

Topology for ISIS

The network topology includes various network elements such as routers, customer edge (CE) devices, Service Aggregator (SA) devices, and Provider Edge (PE) routers. The feature enables OSPF and ISIS support on the IRB interfaces, allowing for efficient routing and communication between network devices within the topology.

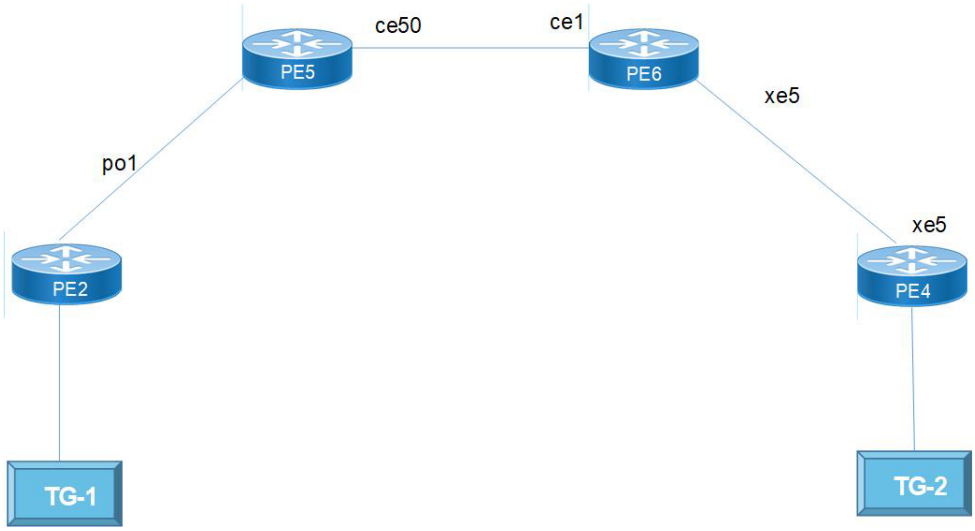

Single Home VxLAN IRB with ISIS

Configure ISIS

PE2

`PE2(config-if)# interface po1`	Enters configuration mode for po 1.
`PE2(config-if)# ip address 20.1.1.1/24`	Assigns the IP address 20.1.1.1 with a subnet mask of 255.255.255.0 to the interface.
`PE2(config-if)#ip router isis 1`	Enables ISIS routing protocol on the interface with process ID 1.
`PE2(config-if)#load-interval 30`	Sets the interval for which interface statistics are collected to 30 seconds.
`PE2(config)#nvo vxlan enable`	Enables the VXLAN feature on the device.
`PE2(config)#nvo vxlan irb`	Enables VXLAN IRB functionality.
`PE2(config-vrf)#mac vrf L2VRF2`	Enters the configuration mode for a MAC VRF named L2VRF2.

`PE2(config-vrf)# rd 2.2.2.2:11`	Sets the route distinguisher (RD) for the VRF to 2.2.2.2:11.
`PE2(config-vrf)#route-target both 10.10.10.10:100`	Specifies import and export route targets for the VRF.
`PE2(config-vrf)#ip vrf L3VRF2`	Enters the configuration mode for an IP VRF named L3VRF2.
`PE2(config-vrf)#rd 61000:11`	Sets the RD for the IP VRF to 61000:11
`PE2(config-vrf)# route-target both 101:101`	Specifies import and export route targets for the IP VRF.
`PE2(config-vrf)# l3vni 2000`	Configures the Layer 3 VNI (Virtual Network Identifier) for the IP VRF.
`PE2(config)#interface irb2001`	Enters the configuration mode for interface IRB2001.
`PE2(config-irb-if)# ip vrf forwarding L3VRF2`	Associates the interface with the IP VRF L3VRF2.
`PE2(config-irb-if)# ip address 13.13.13.1/24`	Configures an IP address with a subnet mask of /24 on IRB2001.
`PE2(config-irb-if)#mtu 9000`	Sets the Maximum Transmission Unit (MTU) for the interface to 9000 bytes.
`PE2(config-irb-if)#ip router isis 2`	Associates the interface with ISIS routing process 2.
`PE2(config-irb)#interface irb3001`	Enters the configuration mode for interface IRB3001.
`PE2(config-irb-if)# ip vrf forwarding L3VRF2`	Associates the interface with the IP VRF L3VRF2.
`PE2(config-irb-if)# ipv6 address 3001::1/64`	Configures an IPv6 address on IRB3001 with the specified prefix length.
`PE2(config-irb-if)#mtu 9000`	Sets the MTU for the interface to 9000 bytes.
`PE2(config-irb)#ipv6 router isis 3`	Associates the interface with IPv6 ISIS routing process 3.
`PE2(config)#router isis 2 L3VRF2`	Enters the configuration mode for ISIS routing process 2 within VRF L3VRF2.
`PE2(config-router)#is-type level-1-2`	Specifies the ISIS level type as level-1-2.
`PE2(config-router)#metric-style wide`	Configures a wide metric style for ISIS.
`PE2(config-router)# dynamic-hostname`	Enables dynamic hostname assignment for the ISIS router.
`PE2(config-router)# bfd all-interfaces`	Enables Bidirectional Forwarding Detection (BFD) on all interfaces within ISIS.
`PE2(config-router)#net 49.0000.0000.0221.00`	Configures the network entity title (NET) for the ISIS process.
`PE2(config)#router isis 3 L3VRF2`	Enters the configuration mode for ISIS routing process 3 within VRF L3VRF2.
`PE2(config-router)#is-type level-1-2`	Specifies the ISIS level type as level-1-2.
`PE2(config-router)# metric-style wide`	Configures a wide metric style for ISIS.
`PE2(config-router)# dynamic-hostname`	Enables dynamic hostname assignment for the ISIS router.
`PE2(config-router)#bfd all-interfaces`	Enables BFD on all interfaces within ISIS.
`PE2(config-router)# net 49.0000.0000.0222.00`	Configures the network entity title (NET) for ISIS routing with the specified value.
`PE2(config)#nvo vxlan vtep-ip-global 2.2.2.2`	Configures the global VxLAN VTEP IP address to 2.2.2.2.
`PE2(config)#nvo vxlan id 201 ingress-replication`	Configures a VxLAN with VNI 201 and specifies ingress-replication for multicast traffic handling.
`PE2(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Specifies the EVPN-BGP host-reachability-protocol for the VxLAN with the VRF L2VRF2
`PE2(config-nvo)# evpn irb2001`	Enables EVPN IRB (Integrated Routing and Bridging) for VxLAN interface IRB2001.
`PE2(config-nvo)# vni-name VNI-201`	Assigns a name VNI-201 to the VxLAN VNI 201.

`PE2(config)#nvo vxlan id 3001 ingress-replication`	Configures another VxLAN with VNI 3001 and specifies ingress-replicationr for multicast traffic handling.
`PE2(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Specifies the EVPN-BGP host-reachability-protocol for the VxLAN with the VRF L2VRF2.
`PE2(config-nvo)# evpn irb3001`	Enables EVPN IRB for VxLAN interface IRB3001.
`PE2(config-if)#interface xe11`	Enters the configuration mode for the interface 11.
`PE2(config-if)#switchport`	Configures the interface as a Layer 2 switchport.
`PE2(config-if)#load-interval 30`	Sets the interval for which interface statistics are collected to 30 seconds.
`PE2(config)#nvo vxlan access-if port-vlan xe11 100`	Configures a VxLAN network virtualization overlay (NVO) on the interface xe2 with VLAN ID 100
`PE2(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN Network Identifier (VNID) 101.
`PE2(config-nvo-acc-if)#nvo vxlan access-if port-vlan xe11 2001`	Configures another VxLAN NVO on the same interface xe2, but this time with VLAN ID 2001
`PE2(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN Network Identifier (VNID) 101.
`PE2(config-nvo-acc-if)#nvo vxlan access-if port-vlan xe11 2001`	Configures another VxLAN NVO on the same interface xe2, but this time with VLAN ID 2001
`PE2(config-nvo-acc-if)# map vnid 2001`	Maps VLAN 2001 to a different VxLAN VNID.
`PE2(config-if)#router isis 1`	Starts the ISIS routing process with process ID 1.
`PE2(config-if)#is-type level-1-2`	Specifies that the router participates in both Level 1 and Level 2 routing.
`PE2(config-if)#metric-style wide`	Configures the metric style to be wide, enabling more flexibility in metric calculations.
`PE2(config-if)#dynamic-hostname`	Enables the dynamic hostname feature for ISIS.
`PE2(config-if)#bfd all-interfaces`	Configures Bidirectional Forwarding Detection on all interfaces.
`PE2(config-if)#net 49.0000.0000.0001.00`	Specifies the network entity title (NET) for ISIS.

BGP Configuration

`PE2(config)#router bgp 100`	Starts the BGP routing process with an autonomous system number (AS) of 100.
`PE2(config-router)#bgp router-id 2.2.2.2`	Sets the BGP router ID to 2.2.2.2.
`PE2(config-router)#neighbor 3.3.3.3 remote-as 100`	Configures a BGP neighbor with the IP address 3.3.3.3 and specifies the remote AS number as 100.
`PE2(config-router)#neighbor 3.3.3.3 update-source lo`	Specifies that loopback interface (lo) is the source for BGP updates to the neighbor.
`PE2(config-router)#neighbor 3.3.3.3 advertisement-interval 0`	Sets the advertisement interval to 0, which means updates will be sent immediately.
`PE2(config-router)#address-family ipv4 unicast`	Enters the configuration mode for the IPv4 unicast address family within the router configuration.
`PE2(config-router-af)#network 2.2.2.2/32`	Specifies that network 2.2.2.2 with a /32 subnet mask is part of the IPv4 unicast address family.
`PE2(config-router-af)#neighbor 3.3.3.3 activate`	Activates the neighbor with the IP address 3.3.3.3 for the IPv4 unicast address family.
`PE2(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv4 unicast address family.

`PE2(config-router)#address-family l2vpn evpn`	Enters the configuration mode for the L2VPN EVPN address family within the router configuration.
`PE2(config-router-af)#neighbor 3.3.3.3 activate`	Activates the neighbor with the IP address 3.3.3.3 for the L2VPN EVPN address family.
`PE2(config-router-af)#exit-address-family`	Exits the configuration mode for the L2VPN EVPN address family.
`PE2(config-router)#address-family ipv4 vrf L3VRF2`	Enters the configuration mode for the IPv4 address family within the VRF named L3VRF2.
`PE2(config-router-af)#redistribute connected`	Configures the redistribution of directly connected routes into the IPv4 address family for the specified VRF.
`PE2(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv4 address family within the VRF L3VRF2.
`PE2(config-router-af)#address-family ipv6 vrf L3VRF2`	Enters the configuration mode for the IPv6 address family within the VRF named L3VRF2.
`PE2(config-router-af)#redistribute connected`	Configures the redistribution of directly connected routes into the IPv6 address family for the specified VRF.
`PE2(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv6 address family within the VRF L3VRF2.

PE5

`PE5(config-if)#interface po1`	Enters the configuration mode for po1.
`PE5(config-if)#ip address 20.1.1.2/24`	Assigns the IP address 20.1.1.2 with a subnet mask of /24 to this interface.
`PE5(config-if)#ip router isis 1`	Specifies that ISIS routing process 1 is enabled on this interface.
`PE5(config-if)#load-interval 30`	Sets the load interval to 30 seconds for monitoring the interface.
`PE5(config-if)#interface po2`	Enters the configuration mode for po2.
`PE5(config-if)#ip address 70.1.1.2/24`	Assigns the IP address 70.1.1.2 with a subnet mask of /24 to this interface.
`PE5(config-if)#load-interval 30`	Sets the load interval to 30 seconds for monitoring the interface.
`PE5(config-if)#interface sa1`	Assigns the IP address 10.1.1.2 with a subnet mask of /24 to this interface.
`PE5(config-if)#ip ospf cost 10`	Sets the OSPF cost for this interface to 10.
`PE5(config-if)#load-interval 30`	Sets the load interval to 30 seconds for monitoring the interface.
`PE5(config-if)#interface ce50`	Enters the configuration mode for ce50.
`PE5(config-if)#ip address 50.1.1.1/24`	Assigns the IP address 50.1.1.1 with a subnet mask of /24 to this interface.
`PE5(config-if)#ip router isis 1`	Specifies that ISIS routing process 1 is enabled on this interface.
`PE5(config-if)#load-interval 30`	Sets the load interval to 30 seconds for monitoring the interface.
`PE5(config-if)#router ospf 1`	Enters ISIS configuration mode with process ID 1.
`PE5(config-if)#network 10.1.1.0/24 area 0.0.0.0`	Specifies that the network 10.1.1.0 with subnet mask 255.255.255.0 belongs to OSPF area 0.0.0.0.

PE5(config-if)#network 30.1.1.0/24 area 0.0.0.0	Specifies another network, 30.1.1.0 with subnet mask 255.255.255.0, also belonging to OSPF area 0.0.0.0.
PE5(config-if)#network 70.1.1.0/24 area 0.0.0.0	Specifies a third network, 70.1.1.0 with subnet mask 255.255.255.0, in OSPF area 0.0.0.0.
PE5(config-if)#router isis 1	Enters ISIS configuration mode with process ID 1.
PE5(config-if)#is-type level-1-2	Configures this ISIS router to support both Level 1 and Level 2 routing.
PE5(config-if)#metric-style wide	Configures ISIS to use the wide metric style, which allows for greater flexibility in metric values.
PE5(config-if)#dynamic-hostname	Allows the hostname to be dynamically generated.
PE5(config-if)#bfd all-interfaces	Enables Bidirectional Forwarding Detection on all interfaces.
PE5(config-if)#net 49.0000.0005.0001.00	Sets the NET for this router.
PE5(config-if)#exit	Exits from the router mode.

PE 6

PE6#configure terminal	Enters the configuration maode.
PE6(config-if)#interface sa2	Enters configuration mode for interface sa2.
PE6(config-if)#ip address 80.1.1.2/24	Assigns the IP address 80.1.1.2 with a subnet mask of 255.255.255.0 to interface sa2.
PE6(config-if)#ip router isis 1	Associates ISIS routing protocol with this interface using process ID 1.
PE6(config-if)#load-interval 30	Sets the load-interval to 30 seconds.
PE6(config-if)#interface ce1	Enters configuration mode for interface ce1.
PE6(config-if)#ip address 50.1.1.2/24	Assigns the IP address 50.1.1.2 with a subnet mask of 255.255.255.0 to interface ce1.
PE6(config-if)#ip router isis 1	Associates ISIS routing protocol with this interface using process ID 1.
PE6(config-if)#load-interval 30	Sets the load-interval to 30 seconds.
PE6(config-if)#interface ce2	Enters configuration mode for interface ce2.
PE6(config-if)#speed 40g	Sets the interface speed to 40 gigabits per second.
PE6(config-if)#ip address 40.1.1.1/24	Assigns the IP address 40.1.1.1 with a subnet mask of 255.255.255.0 to interface ce2.
PE6(config-if)#ip ospf cost 10	Sets the OSPF cost for this interface to 10.
PE6(config-if)#load-interval 30	Sets the load-interval to 30 seconds.
PE6(config-if)#router ospf 1	Enters ISIS configuration mode with process ID 1.
PE6(config-if)#network 30.1.1.0/24 area 0.0.0.0	Specifies another network, 30.1.1.0 with subnet mask 255.255.255.0, also belonging to OSPF area 0.0.0.0.
PE6(config-if)#network 40.1.1.0/24 area 0.0.0.0	Specifies a third network, 40.1.1.0/24 with subnet mask 255.255.255.0, in OSPF area 0.0.0.0.
PE6(config-if)#router isis 1	Enters ISIS configuration mode with process ID 1.
PE6(config-if)#is-type level-1-2	Configures this ISIS router to support both Level 1 and Level 2 routing.
PE6(config-if)#metric-style wide	Configures ISIS to use the wide metric style, which allows for greater flexibility in metric values.

PE6(config-if)#dynamic-hostname	Allows the hostname to be dynamically generated.
PE6(config-if)#bfd all-interfaces	Enable BFD on all network interfaces.

PE4

PE4#configure terminal	Enters the configuration mode.
PE4(config-if)# interface xe5	Enters configuration mode for xe5.
PE4(config-if)# ip address 60.1.1.2/24	Assigns the IP address 60.1.1.2 with a subnet mask of 255.255.255.0 to the interface.
PE4(config-if)#ip router isis 1	Enables ISIS routing protocol on the interface with process ID 1.
PE4(config-if)#load-interval 30	Sets the interval for which interface statistics are collected to 30 seconds.
PE4(config)#nvo vxlan enable	Enables the VXLAN feature on the device.
PE4(config)#nvo vxlan irb	Enables VXLAN IRB functionality.
PE4(config-vrf)#mac vrf L2VRF2	Configures a VRF instance named L2VRF2 and associates it with a specific RD
PE4(config-vrf)# rd 3.3.3.3:11	Sets the RD for the L2VRF2 VRF to 3.3.3.3:11.
PE4(config-vrf)#route-target both 10.10.10.10:100	Associates a route target with the L2VRF2 VRF for VPN route distribution.
PE4(config-vrf)#ip vrf L3VRF2	Configures another VRF named L3VRF2.
PE4(config-vrf)#rd 63000:11	Sets the RD for the L3VRF2 VRF to 63000:11.
PE4(config-vrf)# route-target both 101:101	Associates a route target with the L3VRF2 VRF for VPN route distribution.
PE4(config-vrf)# l3vni 2000	Configures the L3VNI for the L3VRF2 VRF.
PE4(config)#interface irb2001	Configuring an IRB interface with the number 2001.
PE4(config-irb-if)# ip vrf forwarding L3VRF2	Associates the IRB interface with the L3VRF2 VRF.
PE4(config-irb-if)# ip address 14.14.14.1/24	Assigns an IP address to the IRB interface.
PE4(config-irb-if)#mtu 9000	Sets the MTU for the IRB interface.
PE4(config-irb-if)#ip router isis 2	Associates the IRB interface with ISIS routing.
PE4(config-irb)#interface irb3002	Configures another IRB interface with the number 3002.
PE4(config-irb-if)# ip vrf forwarding L3VRF2	Associates the second IRB interface with the "L3VRF2" VRF.
PE4(config-irb-if)# ipv6 address 3002::1/64	Assigns an IPv6 address to the second IRB interface.
PE4(config-irb-if)#mtu 9000	Sets the MTU for the second IRB interface.
PE4(config-irb)#ipv6 router isis 3	Associates the IRB interfaces with IPv6 and ISIS routing.
PE4(config)#router isis 2 L3VRF2	Configures ISIS routing with the VRF L3VRF2.
PE4(config-router)#is-type level-1-2	Sets the ISIS level type to level-1-2.
PE4(config-router)# metric-style wide	Configures a wide metric style for ISIS.
PE4(config-router)# dynamic-hostname	Enables dynamic hostname assignment for the ISIS router.
PE4(config-router)#bfd all-interfaces	Enables BFD on all interfaces within ISIS.
PE4(config-router)# net 49.0000.0000.0441.00	Configures the network entity title (NET) for ISIS routing with the specified value.
PE4(config)#router isis 3 L3VRF2	Configures ISIS routing with the VRF L3VRF2.
PE4(config-router)#is-type level-1-2	Sets the ISIS level type to level-1-2.

`PE4(config-router)# metric-style wide`	Configures a wide metric style for ISIS.
`PE4(config-router)# dynamic-hostname`	Enables dynamic hostname assignment for the ISIS router.
`PE4(config-router)#bfd all-interfaces`	Enables BFD on all interfaces within ISIS.
`PE4(config-router)# net 49.0000.0000.0442.00`	Configures the network entity title (NET) for ISIS routing with the specified value.
`PE4(config)#nvo vxlan vtep-ip-global 3.3.3.3`	Configures the global VxLAN VTEP IP address to 3.3.3.3.
`PE4(config)#nvo vxlan id 201 ingress-replication`	Configures a VxLAN with VNI 201 and specifies ingress-replication for multicast traffic handling.
`PE4(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Specifies the EVPN-BGP host-reachability-protocol for the VxLAN with the VRF L2VRF2
`PE4(config-nvo)# evpn irb2001`	Enables EVPN IRB (Integrated Routing and Bridging) for VxLAN interface IRB2001.
`PE4(config-nvo)# vni-name VNI-201`	Assigns a name VNI-201 to the VxLAN VNI 201.
`PE4(config)#nvo vxlan id 3002 ingress-replication`	Configures another VxLAN with VNI 3002 and specifies ingress-replicationr for multicast traffic handling.
`PE4(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Specifies the EVPN-BGP host-reachability-protocol for the VxLAN with the VRF L2VRF2.
`PE4(config-nvo)# evpn irb3002`	Enables EVPN IRB for VxLAN interface IRB3002
`PE4(config-if)#interface xe5`	Enters the configuration mode for the interface 5.
`PE4(config-if)#switchport`	Configures the interface as a L2 switchport.
`PE4(config-if)#load-interval 30`	Sets the interval for which interface statistics are collected to 30 seconds.
`PE4(config)#nvo vxlan access-if port-vlan xe5 100`	Configures a VxLAN network virtualization overlay (NVO) on the interface xe2 with VLAN ID 100
`PE4(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN Network Identifier (VNID) 101.
`PE4(config-nvo-acc-if)#nvo vxlan access-if port-vlan xe5 2001`	Configures another VxLAN NVO on the same interface xe2, but this time with VLAN ID 2001
`PE4(config-nvo-acc-if)# map vnid 2001`	Maps VLAN 2001 to a different VxLAN VNID.
`PE4(config-if)#router isis 1`	Starts the ISIS routing process with process ID 1.
`PE4(config-if)#is-type level-1-2`	Specifies that the router participates in both Level 1 and Level 2 routing.
`PE4(config-if)#metric-style wide`	Configures the metric style to be wide, enabling more flexibility in metric calculations.
`PE4(config-if)#dynamic-hostname`	Enables the dynamic hostname feature for ISIS.
`PE4(config-if)#bfd all-interfaces`	Configures Bidirectional Forwarding Detection on all interfaces.
`PE4(config-if)#net 49.0000.0003.0001.00`	Specifies the network entity title (NET) for ISIS.

BGP Configuration

`PE4(config)#router bgp 100`	Starts the BGP routing process with an autonomous system number (AS) of 100.
`PE4(config-router)#bgp router-id 3.3.3.3`	Sets the BGP router ID to 3.3.3.3
`PE4(config-router)#neighbor 2.2.2.2 remote-as 100`	Configures a BGP neighbor with the IP address 2.2.2.2 and specifies the remote AS number as 100.
`PE4(config-router)#neighbor 2.2.2.2 update-source lo`	Specifies that loopback interface (lo) is the source for BGP updates to the neighbor.

`PE4(config-router)#neighbor 2.2.2.2 advertisement-interval 0`	Sets the advertisement interval to 0, which means updates will be sent immediately.
`PE4(config-router)#address-family ipv4 unicast`	Enters the configuration mode for the IPv4 unicast address family within the router configuration.
`PE4(config-router-af)#network 3.3.3.3/32`	Specifies that network 3.3.3.3 with a /32 subnet mask is part of the IPv4 unicast address family.
`PE4(config-router-af)#neighbor 2.2.2.2 activate`	Activates the neighbor with the IP address 2.2.2.2 for the IPv4 unicast address family.
`PE4(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv4 unicast address family.
`PE4(config-router)#address-family l2vpn evpn`	Enters the configuration mode for the L2VPN EVPN address family within the router configuration.
`PE4(config-router-af)#neighbor 2.2.2.2 activate`	Activates the neighbor with the IP address 2.2.2.2 for the L2VPN EVPN address family.
`PE4(config-router-af)#exit-address-family`	Exits the configuration mode for the L2VPN EVPN address family.
`PE4(config-router)#address-family ipv4 vrf L3VRF2`	Enters the configuration mode for the IPv4 address family within the VRF named L3VRF2.
`PE4(config-router-af)#redistribute connected`	Configures the redistribution of directly connected routes into the IPv4 address family for the specified VRF.
`PE4(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv4 address family within the VRF L3VRF2.
`PE4(config-router-af)#address-family ipv6 vrf L3VRF2`	Enters the configuration mode for the IPv6 address family within the VRF named L3VRF2.
`PE4(config-router-af)#redistribute connected`	Configures the redistribution of directly connected routes into the IPv6 address family for the specified VRF.
`PE4(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv6 address family within the VRF L3VRF2.

Implementation Examples

Scenario: Configure OSPF and ISIS protocols on an IRB interface with an assigned IP address.

New CLI Commands

No CLI commands are introduced.

Validation

OSPF Validation

```
PE1#show ip ospf neighbor
Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
       Instance ID
50.1.1.1          1   Full/DR         00:00:38    10.1.1.2        sa1
```

```
           0

Total number of full neighbors: 1
OSPF process 2 VRF(L3VRF1):
Neighbor ID     Pri   State            Dead Time   Address         Interface
       Instance ID
192.0.0.1         0   Full/DROther     00:00:34    11.11.11.2      irb1001
           0
PE1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================================
1.1.1.1          4.4.4.4          Installed     00:15:59      00:15:59

Total number of entries are 2
PE1# show evpn irb-status
IRB is ACTIVE in Hardware
PE1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
101      11.11.11.1       9819.2ccd.9301 Static Local   ----
101      11.11.11.2       0010.9400.0001 Dynamic Local  ----
Total number of entries are 2
PE1#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, 00:53:03
O            4.4.4.4/32 [110/31] via 10.1.1.2, sa1, 00:16:29
O            7.7.7.7/32 [110/12] via 10.1.1.2, sa1, 00:44:26
C            10.1.1.0/24 is directly connected, sa1, 00:50:10
O            30.1.1.0/24 [110/20] via 10.1.1.2, sa1, 00:44:22
O            40.1.1.0/24 [110/30] via 10.1.1.2, sa1, 00:17:14
O            70.1.1.0/24 [110/11] via 10.1.1.2, sa1, 00:45:18
C            127.0.0.0/8 is directly connected, lo, 00:53:03
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 00:53:03
C            127.0.0.0/8 is directly connected, lo.management, 00:53:03
IP Route Table for VRF "L2VRF1"
IP Route Table for VRF "L3VRF1"
B            4.4.4.4/32 [0/0] is directly connected, tunvxlan2, 00:16:25
```

```
B            7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:44:21
C            11.11.11.0/24 is directly connected, irb1001, 00:53:03
B            12.12.12.0/24 [200/0] via 4.4.4.4 (recursive is directly connected,
 tunvxlan2), 00:16:26
B            16.16.16.0/24 [200/0] via 7.7.7.7 (recursive is directly connected,
 tunvxlan2), 00:44:21
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:53:03

Gateway of last resort is not set


PE1#show bgp l2vpn evpn
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i
- internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop          Metric    LocPrf    Weight     Path Peer       Encap

RD[7100:11]
*>i  [5]:[0]:[0]:[24]:[16.16.16.0]:[0.0.0.0]:[1000]
                      7.7.7.7              0        100       0    i  7.7.7.7       VXLAN
*>i  [5]:[0]:[0]:[64]:[7002::]:[::]:[1000]
                      7.7.7.7              0        100       0    i  7.7.7.7       VXLAN

RD[56000:11]
*>i  [5]:[0]:[0]:[24]:[12.12.12.0]:[0.0.0.0]:[1000]
                      4.4.4.4              0        100       0    ?  4.4.4.4       VXLAN
*>i  [5]:[0]:[0]:[64]:[2002::]:[::]:[1000]
                      4.4.4.4              0        100       0    ?  4.4.4.4       VXLAN

RD[1.1.1.1:11] VRF[L2VRF1]:
*>   [2]:[0]:[101]:[48,0010:9400:0001]:[0]:[101]
                     1.1.1.1              0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[101]:[48,0010:9400:0001]:[32,11.11.11.2]:[101]
                     1.1.1.1              0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[101]:[48,9819:2ccd:9301]:[32,11.11.11.1]:[101]
                     1.1.1.1              0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[102]:[48,0010:9400:0002]:[0]:[102]
                      4.4.4.4              0        100       0    i  4.4.4.4       VXLAN
* i  [2]:[0]:[102]:[48,0010:9400:0002]:[32,12.12.12.2]:[102]
                      4.4.4.4              0        100       0    i  4.4.4.4       VXLAN
```

```
* i  [2]:[0]:[102]:[48,5c07:5813:425e]:[32,12.12.12.1]:[102]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>   [2]:[0]:[2001]:[48,0010:9400:0009]:[0]:[2001]
                    1.1.1.1             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[2001]:[48,0010:9400:0009]:[128,2001::2][2001]
                    1.1.1.1             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[2001]:[48,9819:2ccd:9301]:[128,2001::1][2001]
                    1.1.1.1             0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[2002]:[48,0010:9400:000a]:[0]:[2002]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
* i  [2]:[0]:[2002]:[48,0010:9400:000a]:[128,2002::2][2002]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
* i  [2]:[0]:[2002]:[48,5c07:5813:425e]:[128,2002::1][2002]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>   [3]:[101]:[32,1.1.1.1]
                    1.1.1.1             0       100      32768  i  ---------      VXLAN
* i  [3]:[102]:[32,4.4.4.4]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>   [3]:[2001]:[32,1.1.1.1]
                    1.1.1.1             0       100      32768  i  ---------      VXLAN
* i  [3]:[2002]:[32,4.4.4.4]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN

RD[4.4.4.4:11]
*>i  [2]:[0]:[102]:[48,0010:9400:0002]:[0]:[102]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[102]:[48,0010:9400:0002]:[32,12.12.12.2]:[102]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[102]:[48,5c07:5813:425e]:[32,12.12.12.1]:[102]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[2002]:[48,0010:9400:000a]:[0]:[2002]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[2002]:[48,0010:9400:000a]:[128,2002::2][2002]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[2002]:[48,5c07:5813:425e]:[128,2002::1][2002]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>i  [3]:[102]:[32,4.4.4.4]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN
*>i  [3]:[2002]:[32,4.4.4.4]
                     4.4.4.4             0       100       0    i  4.4.4.4       VXLAN

Total number of prefixes 28

PE3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================
4.4.4.4          1.1.1.1          Installed     00:18:19      00:18:19

Total number of entries are 1
```

```
PE3#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
       Instance ID
40.1.1.2          1   Full/DR         00:00:36    40.1.1.1        ce30
           0

Total number of full neighbors: 1
OSPF process 2 VRF(L3VRF1):
Neighbor ID     Pri   State           Dead Time   Address         Interface
       Instance ID
192.0.0.2         0   Full/DROther    00:00:36    12.12.12.2      irb1001
           0
PE3#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            1.1.1.1/32 [110/31] via 40.1.1.1, ce30, 00:18:35
C            4.4.4.4/32 is directly connected, lo, 00:19:22
O            7.7.7.7/32 [110/22] via 40.1.1.1, ce30, 00:18:35
O            10.1.1.0/24 [110/30] via 40.1.1.1, ce30, 00:18:35
O            30.1.1.0/24 [110/20] via 40.1.1.1, ce30, 00:18:35
C            40.1.1.0/24 is directly connected, ce30, 00:19:21
O            70.1.1.0/24 [110/21] via 40.1.1.1, ce30, 00:18:35
C            127.0.0.0/8 is directly connected, lo, 00:20:05
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 00:19:19
C            127.0.0.0/8 is directly connected, lo.management, 00:20:05
IP Route Table for VRF "L3VRF1"
B            1.1.1.1/32 [0/0] is directly connected, tunvxlan2, 00:18:31
B            11.11.11.0/24 [200/0] via 1.1.1.1 (recursive is directly connected,
 tunvxlan2), 00:18:32
C            12.12.12.0/24 is directly connected, irb1001, 00:19:28
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:19:29
IP Route Table for VRF "L2VRF1"

Gateway of last resort is not set
PE3# show bgp l2vpn evpn
BGP table version is 4, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i
- internal,
```

```
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop       Metric    LocPrf     Weight  Path Peer      Encap

RD[51000:11]
*>i  [5]:[0]:[0]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                      1.1.1.1             0        100       0    ?  1.1.1.1       VXLAN
*>i  [5]:[0]:[0]:[64]:[2001::]:[::]:[1000]
                      1.1.1.1             0        100       0    ?  1.1.1.1       VXLAN

RD[1.1.1.1:11]
*>i  [2]:[0]:[101]:[48,0010:9400:0001]:[0]:[101]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[101]:[48,0010:9400:0001]:[32,11.11.11.2]:[101]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[101]:[48,9819:2ccd:9301]:[32,11.11.11.1]:[101]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[2001]:[48,0010:9400:0009]:[0]:[2001]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[2001]:[48,0010:9400:0009]:[128,2001::2][2001]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[2001]:[48,9819:2ccd:9301]:[128,2001::1][2001]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [3]:[2001]:[32,1.1.1.1]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN

RD[4.4.4.4:11] VRF[L2VRF1]:
* i  [2]:[0]:[101]:[48,0010:9400:0001]:[0]:[101]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
* i  [2]:[0]:[101]:[48,0010:9400:0001]:[32,11.11.11.2]:[101]
                      1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
* i  [2]:[0]:[101]:[48,9819:2ccd:9301]:[32,11.11.11.1]:[101]
                       1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>   [2]:[0]:[102]:[48,0010:9400:0002]:[0]:[102]
                     4.4.4.4             0        100       32768  i  ---------      VXLAN
*>   [2]:[0]:[102]:[48,0010:9400:0002]:[32,12.12.12.2]:[102]
                     4.4.4.4             0        100       32768  i  ---------      VXLAN
*>   [2]:[0]:[102]:[48,5c07:5813:425e]:[32,12.12.12.1]:[102]
                      4.4.4.4             0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[2001]:[48,0010:9400:0009]:[0]:[2001]
```

```
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
* i  [2]:[0]:[2001]:[48,0010:9400:0009]:[128,2001::2][2001]
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
* i  [2]:[0]:[2001]:[48,9819:2ccd:9301]:[128,2001::1][2001]
                       1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
*>   [2]:[0]:[2002]:[48,0010:9400:000a]:[0]:[2002]
                     4.4.4.4              0        100       32768  i  ---------      VXLAN
*>   [2]:[0]:[2002]:[48,0010:9400:000a]:[128,2002::2][2002]
                     4.4.4.4              0        100       32768  i  ---------      VXLAN
*>   [2]:[0]:[2002]:[48,5c07:5813:425e]:[128,2002::1][2002]
                     4.4.4.4              0        100       32768  i  ---------      VXLAN
* i  [3]:[101]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
*>   [3]:[102]:[32,4.4.4.4]
                     4.4.4.4              0        100       32768  i  ---------      VXLAN
* i  [3]:[2001]:[32,1.1.1.1]
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
*>   [3]:[2002]:[32,4.4.4.4]
                     4.4.4.4              0        100       32768  i  ---------      VXLAN

Total number of prefixes 26
```

ISIS Validation

```
PE2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
2.2.2.2          3.3.3.3          Installed     00:00:10      00:00:10

Total number of entries are 1
PE2#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA                State  Holdtime  Type Protocol
PE5            po1         b86a.9725.a7f2      Up     28        L1   IS-IS
                                                 Up     28        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA                State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0003      Up     28        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
```

```
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA                State  Holdtime  Type Protocol
Spirent-1      irb3001     0010.9400.000c      Up     28        L2   IS-IS
PE2#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            2.2.2.2/32 is directly connected, lo, 02:13:57
i L2         3.3.3.3/32 [115/30] via 20.1.1.2, po1, 00:00:32
i L1         7.7.7.7/32 [115/40] via 20.1.1.2, po1, 01:05:49
C            20.1.1.0/24 is directly connected, po1, 02:13:21
i L1         50.1.1.0/24 [115/20] via 20.1.1.2, po1, 01:06:05
i L1         60.1.1.0/24 [115/30] via 20.1.1.2, po1, 00:00:47
i L1         80.1.1.0/24 [115/30] via 20.1.1.2, po1, 01:05:49
C            127.0.0.0/8 is directly connected, lo, 02:13:57
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 02:13:57
C            127.0.0.0/8 is directly connected, lo.management, 02:13:57
IP Route Table for VRF "L3VRF2"
B            3.3.3.3/32 [0/0] is directly connected, tunvxlan2, 00:00:28
C            13.13.13.0/24 is directly connected, irb2001, 02:13:57
B            14.14.14.0/24 [200/0] via 3.3.3.3 (recursive is directly connected,
 tunvxlan2), 00:00:28
C            127.0.0.0/8 is directly connected, lo.L3VRF2, 02:13:57
IP Route Table for VRF "L2VRF2"

Gateway of last resort is not set
PE2# show bgp l2vpn evpn
BGP table version is 2, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i
- internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop          Metric    LocPrf    Weight   Path Peer        Encap
```

```
RD[63000:11]
*>i  [5]:[0]:[0]:[24]:[14.14.14.0]:[0.0.0.0]:[2000]
                     3.3.3.3             0        100       0    ?  3.3.3.3       VXLAN
*>i  [5]:[0]:[0]:[64]:[3002::]:[::]:[2000]
                     3.3.3.3             0        100       0    ?  3.3.3.3       VXLAN

RD[2.2.2.2:11] VRF[L2VRF2]:
*>   [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0        100      32768  i  ---------     VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0        100      32768  i  ---------     VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0        100      32768  i  ---------     VXLAN
* i  [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0        100      32768  i  ---------     VXLAN
*>   [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0        100      32768  i  ---------     VXLAN
*>   [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0        100      32768  i  ---------     VXLAN
* i  [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100      32768  i  ---------     VXLAN
* i  [3]:[201]:[32,3.3.3.3]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [3]:[3001]:[32,2.2.2.2]
                     2.2.2.2             0        100      32768  i  ---------     VXLAN
* i  [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN

RD[3.3.3.3:11]
*>i  [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                     3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
```

```
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [3]:[201]:[32,3.3.3.3]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [3]:[3002]:[32,3.3.3.3]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN

Total number of prefixes 26
PE2# show nvo vxlan arp-
arp-cache  arp-nd
PE2# show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________
201      13.13.13.1       e8c5.7a76.581d Static Local   ----
201      13.13.13.2       0010.9400.0003 Dynamic Local  ----
201      14.14.14.1       e8c5.7aa8.7cb3 Static Remote  ----
201      14.14.14.2       0010.9400.0005 Dynamic Remote ----
Total number of entries are 4
PE2# show evpn irb-status
IRB is ACTIVE in Hardware
PE2#

PE4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=============================================================================
3.3.3.3          7.7.7.7          Installed     00:01:28      00:01:28
3.3.3.3          2.2.2.2          Installed     00:01:28      00:01:28

Total number of entries are 2
PE4#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA               State  Holdtime  Type Protocol
PE6            xe5         00e0.4b71.f12c     Up     25        L1   IS-IS
                                              Up     25        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA               State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0005     Up     28        L2   IS-IS
```

```
Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb3002     0010.9400.000b      Up     28        L2   IS-IS
PE4#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L2         2.2.2.2/32 [115/30] via 60.1.1.1, xe5, 00:01:46
C            3.3.3.3/32 is directly connected, lo, 02:09:52
i L1         7.7.7.7/32 [115/30] via 60.1.1.1, xe5, 00:01:46
i L1         20.1.1.0/24 [115/30] via 60.1.1.1, xe5, 00:01:46
i L1         50.1.1.0/24 [115/20] via 60.1.1.1, xe5, 00:01:46
C            60.1.1.0/24 is directly connected, xe5, 00:02:02
i L1         80.1.1.0/24 [115/20] via 60.1.1.1, xe5, 00:01:46
C            127.0.0.0/8 is directly connected, lo, 02:09:52
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 02:09:52
C            127.0.0.0/8 is directly connected, lo.management, 02:09:52
IP Route Table for VRF "L3VRF2"
B            2.2.2.2/32 [0/0] is directly connected, tunvxlan2, 00:01:42
B            7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:01:42
B            13.13.13.0/24 [200/0] via 2.2.2.2 (recursive is directly connected,
tunvxlan2), 00:01:42
C            14.14.14.0/24 is directly connected, irb2001, 02:09:52
B            17.17.17.0/24 [200/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:01:42
C            127.0.0.0/8 is directly connected, lo.L3VRF2, 02:09:52
IP Route Table for VRF "L2VRF2"

Gateway of last resort is not set
PE4# show bgp l2vpn evpn
BGP table version is 3, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
```

```
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer  Encap

RD[7400:11]
*>i  [5]:[0]:[0]:[24]:[17.17.17.0]:[0.0.0.0]:[2000]
                     7.7.7.7             0        100       0    i  7.7.7.7         VXLAN
*>i  [5]:[0]:[0]:[64]:[8002::]:[::]:[2000]
                     7.7.7.7             0        100       0    i  7.7.7.7         VXLAN

RD[61000:11]
*>i  [5]:[0]:[0]:[24]:[13.13.13.0]:[0.0.0.0]:[2000]
                     2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN
*>i  [5]:[0]:[0]:[64]:[3001::]:[::]:[2000]
                     2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[3001]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN

RD[3.3.3.3:11] VRF[L2VRF2]:
* i  [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                    3.3.3.3              0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                      3.3.3.3             0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                      3.3.3.3             0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
```

```
                     2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
* i  [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
* i  [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
*>   [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                     3.3.3.3             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                     3.3.3.3             0       100      32768  i  ----------     VXLAN
*>   [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                     3.3.3.3             0       100      32768  i  ----------     VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2        VXLAN
*>   [3]:[201]:[32,3.3.3.3]
                     3.3.3.3             0       100      32768  i  ---------      VXLAN
* i  [3]:[3001]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2        VXLAN
*>   [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3             0       100      32768  i  ----------     VXLAN

Total number of prefixes 28
```

ISIS Validation

```
PE2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
2.2.2.2          3.3.3.3          Installed     00:00:10      00:00:10

Total number of entries are 1
PE2#show clns neighbors
Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA                State  Holdtime  Type Protocol
PE5            po1         b86a.9725.a7f2      Up     28        L1   IS-IS
                                               Up     28        L2   IS-IS
Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA                State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0003      Up     28        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA                State  Holdtime  Type Protocol
```

```
Spirent-1      irb3001     0010.9400.000c      Up     28        L2   IS-IS
PE2#
PE2#
PE2#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            2.2.2.2/32 is directly connected, lo, 02:13:57
i L2         3.3.3.3/32 [115/30] via 20.1.1.2, po1, 00:00:32
i L1         7.7.7.7/32 [115/40] via 20.1.1.2, po1, 01:05:49
C            20.1.1.0/24 is directly connected, po1, 02:13:21
i L1         50.1.1.0/24 [115/20] via 20.1.1.2, po1, 01:06:05
i L1         60.1.1.0/24 [115/30] via 20.1.1.2, po1, 00:00:47
i L1         80.1.1.0/24 [115/30] via 20.1.1.2, po1, 01:05:49
C            127.0.0.0/8 is directly connected, lo, 02:13:57
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 02:13:57
C            127.0.0.0/8 is directly connected, lo.management, 02:13:57
IP Route Table for VRF "L3VRF2"
B            3.3.3.3/32 [0/0] is directly connected, tunvxlan2, 00:00:28
C            13.13.13.0/24 is directly connected, irb2001, 02:13:57
B            14.14.14.0/24 [200/0] via 3.3.3.3 (recursive is directly connected,
 tunvxlan2), 00:00:28
C            127.0.0.0/8 is directly connected, lo.L3VRF2, 02:13:57
IP Route Table for VRF "L2VRF2"

Gateway of last resort is not set
PE2# show bgp l2vpn evpn
BGP table version is 2, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i
- internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
Network          Next Hop            Metric    LocPrf       Weight     Path
 Peer          Encap
```

```
RD[63000:11]
*>i  [5]:[0]:[0]:[24]:[14.14.14.0]:[0.0.0.0]:[2000]
                      3.3.3.3             0        100       0    ?  3.3.3.3       VXLAN
*>i  [5]:[0]:[0]:[64]:[3002::]:[::]:[2000]
                      3.3.3.3             0        100       0    ?  3.3.3.3       VXLAN

RD[2.2.2.2:11] VRF[L2VRF2]:
*>   [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2              0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
* i  [3]:[201]:[32,3.3.3.3]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [3]:[3001]:[32,2.2.2.2]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
* i  [3]:[3002]:[32,3.3.3.3]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN

RD[3.3.3.3:11]
*>i  [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
```

```
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [3]:[201]:[32,3.3.3.3]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>i  [3]:[3002]:[32,3.3.3.3]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN

Total number of prefixes 26
PE2#        show nvo vxlan arp-
arp-cache  arp-nd
PE2#        show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________
201      13.13.13.1       e8c5.7a76.581d Static Local   ----
201      13.13.13.2       0010.9400.0003 Dynamic Local  ----
201      14.14.14.1       e8c5.7aa8.7cb3 Static Remote  ----
201      14.14.14.2       0010.9400.0005 Dynamic Remote ----
Total number of entries are 4
PE2# show evpn irb-status
IRB is ACTIVE in Hardware
PE2#

PE4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=============================================================================
3.3.3.3          7.7.7.7          Installed     00:01:28      00:01:28
3.3.3.3          2.2.2.2          Installed     00:01:28      00:01:28

Total number of entries are 2
PE4#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA               State  Holdtime  Type Protocol
PE6            xe5         00e0.4b71.f12c     Up     25        L1   IS-IS
                                              Up     25        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA               State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0005     Up     28        L2   IS-IS
```

```
Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb3002     0010.9400.000b      Up     28        L2   IS-IS
PE4#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L2         2.2.2.2/32 [115/30] via 60.1.1.1, xe5, 00:01:46
C            3.3.3.3/32 is directly connected, lo, 02:09:52
i L1         7.7.7.7/32 [115/30] via 60.1.1.1, xe5, 00:01:46
i L1         20.1.1.0/24 [115/30] via 60.1.1.1, xe5, 00:01:46
i L1         50.1.1.0/24 [115/20] via 60.1.1.1, xe5, 00:01:46
C            60.1.1.0/24 is directly connected, xe5, 00:02:02
i L1         80.1.1.0/24 [115/20] via 60.1.1.1, xe5, 00:01:46
C            127.0.0.0/8 is directly connected, lo, 02:09:52
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 02:09:52
C            127.0.0.0/8 is directly connected, lo.management, 02:09:52
IP Route Table for VRF "L3VRF2"
B            2.2.2.2/32 [0/0] is directly connected, tunvxlan2, 00:01:42
B            7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:01:42
B            13.13.13.0/24 [200/0] via 2.2.2.2 (recursive is directly connected,
tunvxlan2), 00:01:42
C            14.14.14.0/24 is directly connected, irb2001, 02:09:52
B            17.17.17.0/24 [200/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:01:42
C            127.0.0.0/8 is directly connected, lo.L3VRF2, 02:09:52
IP Route Table for VRF "L2VRF2"

Gateway of last resort is not set
PE4# show bgp l2vpn evpn
BGP table version is 3, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
```

```
4 - Ethernet Segment Route
5 - Prefix Route

   Network          Next Hop            Metric    LocPrf      Weight     Path  Peer
Encap

RD[7400:11]
*>i  [5]:[0]:[0]:[24]:[17.17.17.0]:[0.0.0.0]:[2000]
                     7.7.7.7             0        100       0    i  7.7.7.7         VXLAN
*>i  [5]:[0]:[0]:[64]:[8002::]:[::]:[2000]
                     7.7.7.7             0        100       0    i  7.7.7.7         VXLAN

RD[61000:11]
*>i  [5]:[0]:[0]:[24]:[13.13.13.0]:[0.0.0.0]:[2000]
                     2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN
*>i  [5]:[0]:[0]:[64]:[3001::]:[::]:[2000]
                     2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[3001]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN

RD[3.3.3.3:11] VRF[L2VRF2]:
* i  [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                    3.3.3.3              0       100      32768  i  ----------     VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                    3.3.3.3              0       100      32768  i  ----------     VXLAN
* i  [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                    3.3.3.3              0       100      32768  i  ----------     VXLAN
* i  [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
```

```
* i  [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                     3.3.3.3             0        100      32768  i  ----------     VXLAN
*>   [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                     3.3.3.3             0        100      32768  i  ----------     VXLAN
*>   [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                     3.3.3.3             0        100      32768  i  ----------     VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [3]:[201]:[32,3.3.3.3]
                     3.3.3.3             0        100      32768  i  ----------     VXLAN
* i  [3]:[3001]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3             0        100      32768  i  ----------     VXLAN

Total number of prefixes 28
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	**Description**
ECMP	Equal-Cost Multipath
EVPN	Ethernet Virtual Private Network
VxLAN	Virtual Extensible LAN
SR	Segment Routing
IRB	Integrated Routing
OSPF	Open Shortest Path First
ISIS	Intermediate System to Intermediate System

Glossary

The following provides definitions for key terms used throughout this document.

Single Home VxLAN	This refers to a Virtual Extensible LAN (VxLAN) deployment where a single data center or network site is connected to a single external network (usually the internet) for connectivity.
IRB	A networking feature that enables the integration of Layer 3 IP routing and Layer 2 MAC address bridging within the same interface, simplifying network management and resource utilization.

OSPF	A dynamic and efficient link-state routing protocol used to determine the best path for data packets in an IP network. It is characterized by rapid convergence and adaptability, making it suitable for large and dynamic networks.
ISIS	A routing protocol designed for scalability and stability in computer networks, commonly used in large Service Provider networks. It provides a robust framework for routing information exchange.
Layer 3 Routing	Network routing operations at the Network Layer (Layer 3) of the OSI model, focusing on routing IP packets between different subnets or networks.
Layer 2 Bridging	Network bridging operations at the Data Link Layer (Layer 2) of the OSI model, handling the forwarding of data frames based on MAC addresses within the same network segment.
EVPN	Ethernet VPN, a technology that provides advanced and efficient methods for Layer 2 and Layer 3 services in Ethernet networks, often used in data centers and service provider environments.

CHAPTER 8 Multi Home EVPN-VxLAN IRB with OSPF or ISIS

Overview

The support for Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (ISIS) protocols on Virtual Extensible LAN (VxLAN) Integrated Routing (IRB) interface is enhanced with multihoming switches that provides the solution for connecting and managing virtual networks within a data center or network infrastructure.

This feature offers flexibility in configuring network topologies, and ensures compatibility and interoperability within diverse network environments.

Note: Configure mutually exclusive secondary IP subnets between each anycast-IRB and CE within the same L2VNI in multi-homing scenarios or when the same IRB anycast interface is configured on multiple nodes. This ensures unique identification of the routing protocol peer.

Feature Characteristics

The OSPF and ISIS support over the IRB Interface with multihoming feature has the following characteristics:

- Connect the host node to two VTEPs with all-active redundancy mode. It helps forward all traffic from VTEP to the host when one VTEP goes down.

Benefits

The OSPF and ISIS support over the IRB Interface has the following benefits:

- Uninterrupted service between host and VTEP.

Prerequisites

- Router must be up and running.
- Maintain synchronization with VRF changes by performing IRB `shut/no shut` actions when specific events occur within the IPVRF. These events may involve adding or removing Route Targets (RTs), updating Route Distinguishers (RDs), or modifying Layer 3 Virtual Network Identifiers (L3VNIs).

Topology for OSPF

The network topology includes various network elements such as routers, customer edge (CE) devices, Service Aggregator (SA) devices, and Provider Edge (PE) routers. The feature enables OSPF on the IRB interfaces, allowing for efficient routing and communication between network devices within the topology.

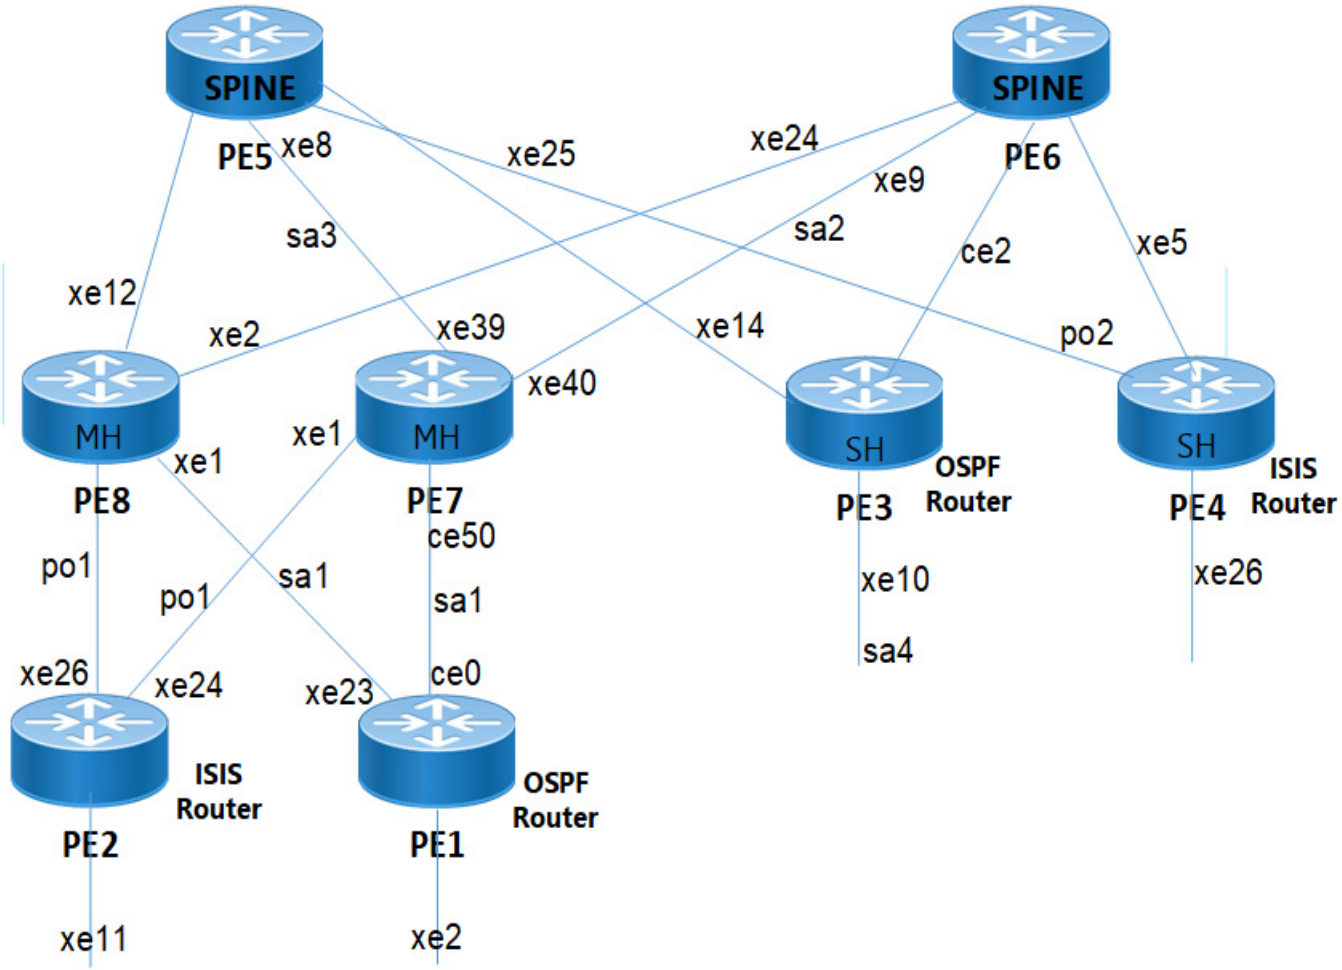

Figure 8-5: Multi Home EVPN-VxLAN IRB with OSPF or ISIS

Configuration

Perform the following configurations to set up different interfaces, routing protocols, and BGP parameters to enable VXLAN, IRB, and EVPN functionality with multihoming in the network.

Configure OSPF Router

Perform the following configurations to create multiple VxLAN interfaces and set up OSPF routing process on PE1 node.

PE1

`PE1(config)#vlan database`	Enter the VLAN configuration mode.
`PE1(config-vlan)# vlan 200-201 bridge 1 state enable`	Enable VLAN (200-201) on bridge 1. Specifying the enable state allows forwarding of frames on this VLAN-aware bridge.
`PE1(config-vlan)#interface sa1`	Enter sa1 interface mode.
`PE1(config-if)# switchport`	Configure port as L2.
`PE1(config-if)# bridge-group 1`	Associate the interface with bridge group 1.
`PE1(config-if)# switchport mode trunk`	Set the switching characteristics of this interface to trunk mode.
`PE1(config-if)# switchport trunk allowed vlan add 200-201`	Enable VLAN ID 200-201 on this port.
`PE1(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa1 interface.
`PE1(config-if)# exit`	Exit from sa1 interface configuration mode.
`PE1(config)#interface ce0`	Enter ce0 interface mode.
`PE1(config-if)# static-channel-group 1`	Make it member port of sa1
`PE1(config-if)#interface lo`	Configures the loopback (lo) interface.
`PE1(config-if)# ip address 1.1.1.1/32 secondary`	Assigns a secondary the IP address 1.1.1.1/32 to the loopback interface.
`PE1(config-if)# ipv6 address ::1/128`	Assigns the IPv6 address `::1/128` to the loopback interface.
`PE1(config-if)#interface vlan1.200`	Enter interface VLAN1.200 configuration mode.
`PE1(config-if)# ip address 101.11.11.1/24`	Assigns the primary IP address.
`PE1(config-if)# ip address 16.16.16.1/24 secondary`	Assigns the secondary IP address.
`PE1(config-if)# ip address 17.17.17.1/24 secondary`	Assigns the secondary IP address.
`PE1(config-if)# exit`	Exit from VLAN interface configuration mode.
`PE1(config)#interface xe2`	Enter xe2 interface configuration mode.
`PE1(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the xe2 interface.
`PE1(config-if)#interface xe2.103`	Enter VLAN ID 103 interface config mode.
`PE1(config-if)# encapsulation dot1q 103`	Setting Encapsulation to dot1q with VLAN ID 103.
`PE1(config-if)# ip address 20.1.1.1/24`	Assigns the primary IP address.
`PE1(config-if)# exit`	Exit from xe2 VLAN interface configuration mode.
`PE1(config)#interface xe23`	Enter xe23 interface configuration mode.
`PE1(config-if)# static-channel-group 1`	Configure static channel 1 member port
`PE1(config)#router ospf 1`	Enters the OSPF configuration mode for OSPF process 1.
`PE1(config-router)# network 16.16.16.0/24 area 0.0.0.0`	Advertises the network 16.16.16.0/24 into OSPF area 0.0.0.0.
`PE1(config-router)# network 17.17.17.0/24 area 0.0.0.0`	Advertises the network 17.17.17.0/24 into OSPF area 0.0.0.0.
`PE1(config-router)# network 20.1.1.0/24 area 0.0.0.0`	Advertises the network 20.1.1.0/24 into OSPF area 0.0.0.0.

Configure ISIS Router

Perform the following configurations to create multiple VxLAN interfaces and set up ISIS routing process on PE2 node.

PE2

`PE2(config)#vlan database`	Enter the VLAN configuration mode.
`PE2(config-vlan)# vlan 100-101 bridge 1 state enable`	Enable VLAN (100-101) on bridge 1. Specifying the enable state allows forwarding of frames on this VLAN-aware bridge
`PE2(config-vlan)#interface po1`	Enter po1 interface mode.
`PE2(config-if)# switchport`	Configure port as L2.
`PE2(config-if)# bridge-group 1`	Associate the interface with bridge group 1.
`PE2(config-if)# switchport mode trunk`	Set the switching characteristics of this interface to trunk mode.
`PE2(config-if)# switchport trunk allowed vlan add 100-101`	Enable VLAN ID 100-101 on this port.
`PE2(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the po1 interface.
`PE2(config-if)#interface lo`	Configure loopback interface
`PE2(config-if)# ip address 2.2.2.2/32 secondary`	Assign the secondary ip address to lo
`PE2(config-if)#interface vlan1.100`	Enter interface VLAN1.100 configuration mode.
`PE2(config-if)# ip address 11.11.11.1/24`	Assign primary IP address
`PE2(config-if)# ip address 12.1.1.1/24 secondary`	Assign secondary IP address towards MH node.
`PE2(config-if)# ip address 13.1.1.1/24 secondary`	Assign secondary address towards MH node
`PE2(config-if)# ip router isis 1`	Configure the isis
`PE2(config-if)# exit`	Exit from VLAN interface configuration mode.
`PE2(config)#interface xe11`	Enter xe11 interface configuration mode.
`PE2(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the xe11 interface.
`PE2(config-if)#interface xe11.105`	Configure subinterface xe11.105.
`PE2(config-if)# encapsulation dot1q 105`	Setting Encapsulation to dot1q with VLAN ID 105.
`PE2(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the xe11 interface.
`PE2(config-if)# ip address 10.1.1.1/24`	Assign an IP address
`PE2(config-if)# ip router isis 1`	Configure interface as ISIS router
`PE2(config-if)# exit`	Exit from xe11 interface configuration mode.
`PE2(config)#interface xe24`	Enter xe24 interface configuration mode.
`PE2(config-if)# channel-group 1 mode active`	Configure the member port for po1 interface
`PE2(config-if)# exit`	Exit from xe24 interface configuration mode.
`PE2(config-if)#interface xe26`	Enter xe26 interface configuration mode.

`PE2(config-if)# channel-group 1 mode active`	Configure the member port for po1 interface
`PE2(config)#router isis 1`	Configure the ISIS router.
`PE2(config-router)# is-type level-1-2`	Configure level1-2 ISIS mode.
`PE2(config-router)# metric-style wide`	Configure metric style as wide.
`PE2(config-router)# dynamic-hostname`	Configure the hostname to be advertised for an ISIS 1 instance.
`PE2(config-router)# bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
`PE2(config-router)# net 49.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`PE2(config-router)# redistribute connected`	Redistribute the connected routes.

Configure VxLAN IRB with Multihoming for OSPF and ISIS Router

Perform the following configurations to create VxLAN IRB interfaces with multihoming for OSPF and ISIS routing process on PE7 and PE8 node.

PE7

`PE7(config)#nvo vxlan enable`	Enable VxLAN.
`PE7(config)#nvo vxlan irb`	Enable VxLAN IRB
`PE7(config)#evpn esi hold-time 60`	Configure hold time for the tunnels to come up during VxLAN initialization before making the esi up.
`PE7(config)#evpn vxlan multihoming enable`	Enable VxLAN multihome
`PE7(config)#ip vrf management`	Enter into VRF configuration mode.
`PE7(config-vrf)#ip vrf L3VRF3`	Create a VRF routing information base called L3VRF3 for OSPF router.
`PE7(config-vrf)# rd 7100:11`	Specify a route distinguisher for the VRF.
`PE7(config-vrf)# route-target both 100:100`	Add import and export route-target extended communities to the VRF.
`PE7(config-vrf)# l3vni 1000`	Configure the L3 Virtual Network Identifier for an IP VRF.
`PE7(config-vrf)#ip vrf L3VRF4`	Create a VRF routing information base called L3VRF4 for ISIS router.
`PE7(config-vrf)# rd 7400:11`	Specify a route distinguisher for the VRF.
`PE7(config-vrf)# route-target both 101:101`	Add import and export route-target extended communities to the VRF.
`PE7(config-vrf)# l3vni 2000`	Configure the L3 Virtual Network Identifier for an IP VRF.
`PE7(config-vrf)#mac vrf L2VRF1`	Create a L2 MAC VRF to use in EVPN routes.
`PE7(config-vrf)# rd 7.7.7.7:11`	Specify a route distinguisher for the MAC VRF.
`PE7(config-vrf)# route-target both 7.7.7.7:100`	Add import and export route-target extended communities to the VRF.
`PE7(config-vrf)#mac vrf L2VRF2`	Create a L2 MAC VRF to use in EVPN routes.
`PE7(config-vrf)# rd 7.7.7.7:12`	Specify a route distinguisher for the VRF.

`PE7(config-vrf)# route-target both 7.7.7.7:101`	Add import and export route-target extended communities to the VRF
`PE7(config-vrf)#exit`	Exit the VRF configuration mode.
`PE7(config)#evpn irb-forwarding anycast-gateway-mac 0000.0000.1111`	Configure common Anycast MAC address for all the IRB interfaces.
`PE7(config)#interface irb701`	Enter into IRB interface configuration mode to configure IRB interfaces for OSPF router.
`PE7(config-irb-if)# ip vrf forwarding L3VRF3`	Associate IRB interface with L3VRF3.
`PE7(config-irb-if)# evpn irb-if-forwarding anycast-gateway-mac`	Map the global anycast IRB MAC address with L3VRF3.
`PE7(config-irb-if)# ip address 17.12.13.1/24 anycast`	Configure an Anycast IP address.
`PE7(config-irb-if)# ip address 16.16.16.2/24 secondary`	Configure a secondary IP address.
`PE7(config-irb-if)#exit`	Exit from the IRB interface.
`PE7(config-irb-if)#interface irb801`	Enter into IRB interface configuration mode to configure IRB interfaces for ISIS router.
`PE7(config-irb-if)# ip vrf forwarding L3VRF4`	Associate IRB interface with L3VRF4.
`PE7(config-irb-if)# evpn irb-if-forwarding anycast-gateway-mac`	Map the global anycast IRB MAC address with L3VRF4.
`PE7(config-irb-if)# ip address 12.13.14.2/24 anycast`	Configure an Anycast IP address.
`PE7(config-irb-if)# ip address 13.1.1.2/24 secondary`	Configure a secondary IP address.
`PE7(config-irb-if)# ip router isis 2`	Configure ISIS router on IRB interface.
`PE7(config-irb-if)#exit`	Exit from the IRB interface.
`PE7(config)#nvo vxlan vtep-ip-global 7.7.7.7`	Configure the source VTEP IP address of the VxLAN tunnel.
`PE7(config)#nvo vxlan id 701 ingress-replication inner-vid-disabled`	Add a tenant L2 VNID to the VxLAN. Specify • `ingress-replication` to use head end replication for forwarding BUM traffic • `inner-vid-disabled` to not carry VID out of network port.
`PE7(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Associate the L2VRF1reachable protocol to Ethernet-VPN over BGP
`PE7(config-nvo)# evpn irb701`	Configure default gateway behavior for IRB interface irb701.
`PE7(config-nvo)# vni-name VNI-701`	Configure VNI name.
`PE7(config-nvo)#exit`	Exit from NVO mode.
`PE7(config-nvo)#nvo vxlan id 801 ingress-replication inner-vid-disabled`	Add a tenant L2 VNID to the VxLAN. Specify • ingress-replication to use head end replication for forwarding BUM traffic • inner-vid-disabled to not carry VID out of network port
`PE7(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Associate the L2VRF2reachable protocol to Ethernet-VPN over BGP

`PE7(config-nvo)# evpn irb801`	Configure default gateway behavior for IRB interface irb801.
`PE7(config-nvo)# vni-name VNI-101`	Configure VNI name.
`PE7(config-nvo)#qos enable`	Enable QoS.
`PE7(config-nvo)#exit`	Exit from NVO mode.
`PE7(config)#interface po1`	Configure po1 interface.
`PE7(config-if)# switchport`	Configure port as L2.
`PE7(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the po1 interface.
`PE7(config-if)# evpn multi-homed system-mac 0000.0000.7782`	Configure system MAC as ESI value for po1 interface. VTEP1 and VTEP2 should have same ESI value.
`PE7(config-if)#interface sa1`	Configure sa1 interface.
`PE7(config-if)# switchport`	Configure port as L2.
`PE7(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the po1 interface.
`PE7(config-if)# evpn multi-homed esi 00:01:02:03:04:05:06:07:08`	Configure 9-octet ESI value for sa1 interface.
`PE7(config-if-es)#interface sa2`	Configure sa2 interface.
`PE7(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa2 interface.
`PE7(config-if)# ip address 80.1.1.1/24`	Configure IP address.
`PE7(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE7(config-if)# ip router isis 1`	Configure ISIS router on sa2 interface.
`PE7(config-if)#exit`	Exit from the sa2 interface.
`PE7(config)#interface sa3`	Configure sa3 interface.
`PE7(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa3 interface.
`PE7(config-if)# ip address 22.1.1.1/24`	Configure IP address.
`PE7(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE7(config-if)#interface ce50`	Configure ce50 interface.
`PE7(config-if)# static-channel-group 1`	Configure member port for sa3
`PE7(config-if)#exit`	Exit from ce50 interface configuration mode.
`PE7(config)#interface lo`	Configure loopback interface.
`PE7(config-if)# ip address 7.7.7.7/32 secondary`	Configure secondary IP address
`PE7(config-if)#exit`	Exit from the lo interface.
`PE7(config)#interface xe1`	Configure xe1 interface.
`PE7(config-if)# channel-group 1 mode active`	Configure member port of po1 interface.
`PE7(config-if)#interface xe39`	Configure xe39 interface.
`PE7(config-if)# static-channel-group 3`	Configure member port of sa3 interface.
`PE7(config-if)#interface xe40`	Configure xe40 interface.
`PE7(config-if)# static-channel-group 2`	Configure member port of sa2 interface.
`PE7(config-if)#exit`	Exit from the xe40 interface.
`PE7(config)#router ospf 1`	Configure OSPF router.

`PE7(config-router)# ospf router-id 7.7.7.7`	Configure router id for OSPF.
`PE7(config-router)# network 7.7.7.7/32 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 7.7.7.7/32
`PE7(config-router)# network 22.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on interface with IP address that match the network address 22.1.1.0/24 .
`PE7(config-router)#router ospf 2 L3VRF3`	Configure OSPF on IRB L3VRF3.
`PE7(config-router)# redistribute bgp`	Redistribute BGP routes.
`PE7(config-router)# network 16.16.16.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 16.16.16.0/24.
`PE7(config-router)#router isis 1`	Configure ISIS router.
`PE7(config-router)# is-type level-1-2`	Configure IS type as level 1 and level 2.
`PE7(config-router)# metric-style wide`	Configure metric-style as wide.
`PE7(config-router)# dynamic-hostname`	Configure the hostname to advertise for the ISIS router.
`PE7(config-router)# bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
`PE7(config-router)# net 49.0000.0007.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`PE7(config-router)# redistribute connected`	Redistribute connected routes.
`PE7(config-router)#router isis 2 L3VRF4`	Configure ISIS on IRB L3VRF4.
`PE7(config-router)# is-type level-1-2`	Configure IS type as level 1 and level 2.
`PE7(config-router)# metric-style wide`	Configure metric-style as wide.
`PE7(config-router)# dynamic-hostname`	Configure the hostname to advertise for the ISIS router.
`PE7(config-router)# bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
`PE7(config-router)# net 49.0000.0000.0475.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`PE7(config-router)# redistribute bgp`	Redistribute connected routes.
`PE7(config-router)#router bgp 700`	Start the eBGP router.
`PE7(config-router)# bgp router-id 7.7.7.7`	Configure BGP router id.
`PE7(config-router)# no bgp inbound-route-filter`	Disable the VPN/BGP inbound route-target filter.
`PE7(config-router)# neighbor 5.5.5.5 remote-as 500`	Configure BGP peering relationship with a customer edge router.
`PE7(config-router)# neighbor 6.6.6.6 remote-as 600`	Configure BGP peering relationship with a customer edge router.
`PE7(config-router)# neighbor 6.6.6.6 advertisement-interval 0`	Configure a minimum advertisement interval between the sending of BGP routing updates.
`PE7(config-router)# address-family l2vpn evpn`	Enter to the L2 VPN address family mode to configure the address-family specific parameters.
`PE7(config-router-af)# neighbor 5.5.5.5 activate`	Enable the exchange of specific address family routes with a neighboring router 5.5.5.5.

PE7(config-router-af)# neighbor 6.6.6.6 activate	Enable the exchange of specific address family routes with a neighboring router 6.6.6.6.
PE7(config-router-af)# exit-address-family	Exit from address family mode.
PE7(config-router)# address-family ipv4 vrf L3VRF3	Enter to the IRB IPv4 VRF address family mode to configure the address-family specific parameters.
PE7(config-router-af)# network 16.16.16.0/24	Enable OSPF routing on IRB interface with IP address that match the network address 16.16.16.0/24.
PE7(config-router-af)# redistribute ospf	Redistribute OSPF routes.
PE7(config-router-af)# exit-address-family	Exit from address family mode.
PE7(config-router)# address-family ipv4 vrf L3VRF4	Enter to the IRB IPv4 VRF address family mode to configure the address-family specific parameters.
PE7(config-router-af)# network 12.1.1.0/24	Enable OSPF routing on IRB interface with IP address that match the network address 12.1.1.0/24.
PE7(config-router-af)# redistribute isis	Redistribute ISIS routes.
PE7(config-router-af)# exit-address-family	Exit address family mode.
PE7(config)#nvo vxlan access-if port-vlan po1 100	Map the VLAN port for ISIS switch on po1 interface to identify the VxLAN traffic and to enter NVO access interface mode.
PE7(config-nvo-acc-if)# map vnid 801	Map the l2vnid to an access-port.
PE7(config-nvo-acc-if)#nvo vxlan access-if port-vlan sa1 200	Configure access-if port for OSPF switch.
PE7(config-nvo-acc-if)# map vnid 701	Map the l2vnid to an access-port.

PE8

PE8(config)#nvo vxlan enable	Enable VxLAN
PE8(config)#nvo vxlan irb	Enable VxLAN IRB
PE8(config)#evpn esi hold-time 60	Configure hold time for the tunnels to come up during VxLAN initialization before making the esi up.
PE8(config)#evpn vxlan multihoming enable	Enable VxLAN multihome
PE8(config)#ip vrf managemen	Enter into VRF configuration mode.
PE8(config-vrf)#ip vrf L3VRF3	Create a VRF routing information base called L3VRF3 for OSPF router.
PE8(config-vrf)# rd 8100:11	Specify a route distinguisher for the VRF.
PE8(config-vrf)# route-target both 100:100	Add import and export route-target extended communities to the VRF.
PE8(config-vrf)# l3vni 1000	Configure the L3 Virtual Network Identifier for an IP VRF.
PE8(config-vrf)#ip vrf L3VRF4	Create a VRF routing information base called L3VRF4 for ISIS router.
PE8(config-vrf)# rd 8400:11	Specify a route distinguisher for the VRF.
PE8(config-vrf)# route-target both 101:101	Add import and export route-target extended communities to the VRF.
PE8(config-vrf)# l3vni 2000	Configure the L3 Virtual Network Identifier for an IP VRF.
PE8(config-vrf)#mac vrf L2VRF1	Create a L2 MAC VRF to use in EVPN routes.
PE8(config-vrf)# rd 8.8.8.8:11	Specify a route distinguisher for the MAC VRF.

`PE8(config-vrf)# route-target both 7.7.7.7:100`	Add import and export route-target extended communities to the VRF.
`PE8(config-vrf)#mac vrf L2VRF2`	Create a L2 MAC VRF to use in EVPN routes.
`PE8(config-vrf)# rd 8.8.8.8:12`	Specify a route distinguisher for the VRF.
`PE8(config-vrf)# route-target both 7.7.7.7:101`	Add import and export route-target extended communities to the VRF.
`PE7(config-vrf)#exit`	Exit the VRF configuration mode.
`PE8(config)#evpn irb-forwarding anycast-gateway-mac 0000.0000.1111`	Configure common Anycast MAC address for all the IRB interfaces.
`PE8(config-if)#interface irb701`	Enter into IRB interface configuration mode to configure IRB interfaces for OSPF router.
`PE8(config-irb-if)# ip vrf forwarding L3VRF3`	Associate IRB interface with L3VRF3.
`PE8(config-irb-if)# evpn irb-if-forwarding anycast-gateway-mac`	Map the global anycast IRB MAC address with L3VRF3.
`PE8(config-irb-if)# ip address 17.12.13.1/24 anycast`	Configure an Anycast IP address.
`PE8(config-irb-if)# ip address 17.17.17.2/24 secondary`	Configure a secondary IP address. PE7(config-irb-if)#exitExit from the IRB interface.
`PE8(config-irb-if)#interface irb801`	Enter into IRB interface configuration mode to configure IRB interfaces for ISIS router.
`PE8(config-irb-if)# ip vrf forwarding L3VRF4`	Associate IRB interface with L3VRF4.
`PE8(config-irb-if)# evpn irb-if-forwarding anycast-gateway-mac`	Map the global anycast IRB MAC address with L3VRF4.
`PE8(config-irb-if)# ip address 12.13.14.2/24 anycast`	Configure an Anycast IP address.
`PE8(config-irb-if)# ip address 12.1.1.2/24 secondary`	Configure an Anycast IP address.
`PE8(config-irb-if)# ip router isis 2`	Configure ISIS router on IRB interface.
`PE7(config-irb-if)#exit`	Exit from the IRB interface.
`PE8(config)#nvo vxlan vtep-ip-global 8.8.8.8`	Configure the source VTEP IP address of the VxLAN tunnel.
`PE8(config)#nvo vxlan id 701 ingress-replication inner-vid-disabled`	Add a tenant L2 VNID to the VxLAN. Specify ingress-replication to use head end replication for forwarding BUM traffic inner-vid-disabled to not carry VID out of network port
`PE8(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Associate the L2VRF1reachable protocol to Ethernet-VPN over BGP
`PE8(config-nvo)# evpn irb701`	Configure default gateway behavior for IRB interface irb701.
`PE8(config-nvo)# vni-name VNI-701`	Configure VNI name.
`PE7(config-nvo)#exit`	Exit from NVO mode.
`PE8(config-nvo)#nvo vxlan id 801 ingress-replication inner-vid-disabled`	Add a tenant L2 VNID to the VxLAN. Specify ingress-replication to use head end replication for forwarding BUM traffic inner-vid-disabled to not carry VID out of network port
`PE8(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Associate the L2VRF2reachable protocol to Ethernet-VPN over BGP

`PE8(config-nvo)# evpn irb801`	Configure default gateway behavior for IRB interface irb801.
`PE8(config-nvo)# vni-name VNI-101`	Configure VNI name.
`PE8(config-nvo)#qos enable`	Enable QoS.
`PE8(config-nvo)#exit`	Exit from NVO mode.
`PE8(config-vlan)#interface po1`	Configure po1 interface.
`PE8(config-if)# switchport`	Configure port as L2.
`PE8(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the po1 interface.
`PE8(config-if)# evpn multi-homed system-mac 0000.0000.7782`	Configure system MAC as ESI value for po1 interface. VTEP1 and VTEP2 should have same ESI value.
`PE8(config-if-es)#interface sa1`	Configure sa1 interface.
`PE8(config-if)# switchport`	Configure port as L2.
`PE8(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the po1 interface.
`PE8(config-if)# evpn multi-homed esi 00:01:02:03:04:05:06:07:08`	Configure 9-octet ESI value for sa1 interface.
`PE8(config-irb-if)#interface lo`	Configure loopback interface
`PE8(config-if)# ip address 8.8.8.8/32 secondary`	Configure secondary IP address
`PE8(config-if)#exit`	Exit from the lo interface.
`PE8(config)#interface xe1`	Configure xe1 interface.
`PE8(config-if)# static-channel-group 1`	Configure member port of sa3 interface.
`PE8(config-if)#interface xe2`	Configure xe2 interface.
`PE8(config-if)# ip address 90.1.1.1/24`	Configure IP address
`PE8(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE8(config-if)# ip router isis 1`	Configure ISIS router on IRB interface.
`PE8(config-if)#exit`	Exit from the lo interface.
`PE8(config)#interface xe12`	Configure xe1 interface.
`PE8(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the po1 interface.
`PE8(config-if)# ip address 21.1.1.1/24`	Configure IP address
`PE8(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE8(config-if)#exit`	Exit from the lo interface.
`PE8(config-if)#interface xe26`	Configure xe26 interface.
`PE8(config-if)# channel-group 1 mode active`	Configure member port of xe26 interface.
`PE8(config)#router ospf 1`	Configure OSPF router.
`PE8(config-router)# ospf router-id 8.8.8.8`	Configure router id for OSPF.
`PE8(config-router)# network 8.8.8.8/32 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 8.8.8.8/32.
`PE8(config-router)# network 21.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 21.1.1.0/24.
`PE8(config-router)#router ospf 2 L3VRF3`	Configure OSPF on IRB L3VRF3.

`PE8(config-router)# redistribute bgp`	Redistribute BGP routes.
`PE8(config-router)# network 17.17.17.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 17.17.17.0/24
`PE8(config-router)#router isis 1`	Configure ISIS router.
`PE8(config-router)# is-type level-1-2`	Configure IS type as level 1 and level 2.
`PE8(config-router)# metric-style wide`	Configure metric-style as wide.
`PE8(config-router)# dynamic-hostname`	Configure the hostname to advertise for the ISIS router.
`PE8(config-router)# bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
`PE8(config-router)# net 49.0000.0008.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`PE8(config-router)# redistribute connected`	Redistribute connected routes.
`PE8(config-router)#router isis 2 L3VRF4`	Configure ISIS on IRB L3VRF4.
`PE8(config-router)# is-type level-1-2`	Configure IS type as level 1 and level 2.
`PE8(config-router)# metric-style wide`	Configure metric-style as wide.
`PE8(config-router)# dynamic-hostname`	Configure the hostname to advertise for the ISIS router.
`PE8(config-router)# bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
`PE8(config-router)# net 49.0000.0000.0485.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`PE8(config-router)# redistribute bgp`	Redistribute connected routes.
`PE8(config-router)#router bgp 800`	Start the eBGP router.
`PE8(config-router)# bgp router-id 8.8.8.8`	Configure BGP router id.
`PE8(config-router)# no bgp inbound-route-filter`	Disable the VPN/BGP inbound route-target filter.
`PE8(config-router)# neighbor 5.5.5.5 remote-as 500`	Configure BGP peering relationship with a customer edge router.
`PE8(config-router)# neighbor 6.6.6.6 remote-as 600`	Configure BGP peering relationship with a customer edge router.
`PE8(config-router)# neighbor 6.6.6.6 update-source lo`	Configure a minimum advertisement interval between the sending of BGP routing updates.
`PE8(config-router)# address-family l2vpn evpn`	Enter to the L2 VPN address family mode to configure the address-family specific parameters.
`PE8(config-router-af)# neighbor 5.5.5.5 activate`	Enable the exchange of specific address family routes with a neighboring router 5.5.5.5.
`PE8(config-router-af)# neighbor 6.6.6.6 activate`	Enable the exchange of specific address family routes with a neighboring router 6.6.6.6.
`PE8(config-router-af)# exit-address-family`	Exit from address family mode.
`PE8(config-router)# address-family ipv4 vrf L3VRF3`	Enter to the IRB IPv4 VRF address family mode to configure the address-family specific parameters.
`PE8(config-router-af)# network 16.16.16.0/24`	Enable OSPF routing on IRB interface with IP address that match the network address 16.16.16.0/24.
`PE8(config-router-af)# redistribute ospf`	Redistribute OSPF routes.
`PE8(config-router-af)# exit-address-family`	Exit from address family mode.
`PE8(config-router)# address-family ipv4 vrf L3VRF4`	Enter to the IRB IPv4 VRF address family mode to configure the address-family specific parameters.

`PE8(config-router-af)# network 12.1.1.0/24`	Enable OSPF routing on IRB interface with IP address that match the network address 12.1.1.0/24.
`PE8(config-router-af)# redistribute isis`	Redistribute ISIS routes.
`PE8(config-router-af)# exit-address-family`	Exit address family mode.
`PE8(config-router)# exit`	Exit from router mode configuration.
`PE8(config)#nvo vxlan access-if port-vlan po1 100`	Map the VLAN port for ISIS switch on po1 interface to identify the VxLAN traffic and to enter NVO access interface mode.
`PE8(config-nvo-acc-if)# map vnid 801`	Map the l2vnid to an access-port.
`PE8(config-nvo-acc-if)#nvo vxlan access-if port-vlan sa1 200`	Configure access-if port for OSPF switch.
`PE8(config-nvo-acc-if)# map vnid 701`	Map the l2vnid to an access-port.
`PE8(config-nvo-acc-if)#end`	End global configuration.

Configure Spine Node with OSPF Router

Perform the following configurations to make the node as spine with routing protocol as OSPF.

PE5

`PE5(config-if)#interface sa3`	Configure sa3 interface.
`PE5(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa3 interface.
`PE5(config-if)# ip address 22.1.1.2/24`	Configure IP address
`PE5(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE5(config-if)#interface lo`	Configure loopback interface
`PE5(config-if)# ip address 5.5.5.5/32 secondary`	Configure secondary IP address.
`PE7(config-if)#exit`	Exit from the lo interface.
`PE5(config-if)#interface xe8`	Configure xe8 interface.
`PE5(config-if)# static-channel-group 3`	Configure member port of sa3 interface.
`PE5(config-if)#interface xe12`	Configure xe12 interface.
`PE5(config-if)# ip address 21.1.1.2/24`	Configure IP address
`PE5(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE5(config-if)#interface xe14`	Configure xe14 interface.
`PE5(config-if)# ip address 23.1.1.2/24`	Configure IP address
`PE5(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE5(config-if)#interface xe25`	Configure xe25 interface.
`PE5(config-if)# ip address 24.1.1.2/24`	Configure IP address
`PE5(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE5(config)#router ospf 1`	Configure OSPF router.

`PE5(config-router)# network 5.5.5.5/32 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 5.5.5.5/32.
`PE5(config-router)# network 21.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 21.1.1.0/24.
`PE5(config-router)# network 22.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 22.1.1.0/24.
`PE5(config-router)# network 23.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 23.1.1.0/24.
`PE5(config-router)# network 24.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 24.1.1.0/24.
`PE5(config-router)#router bgp 500`	Start the eBGP router 500.
`PE5(config-router)# bgp router-id 5.5.5.5`	Configure BGP router id 5.5.5.5.
`PE5(config-router)# no bgp inbound-route-filter`	Disable the VPN/BGP inbound route-target filter.
`PE5(config-router)# neighbor 3.3.3.3 remote-as 300`	Configure BGP 3.3.3.3 peering relationship with a customer edge router.
`PE5(config-router)# neighbor 4.4.4.4 remote-as 400`	Configure BGP 4.4.4.4 peering relationship with a customer edge router.
`PE5(config-router)# neighbor 7.7.7.7 remote-as 700`	Configure BGP 7.7.7.7 peering relationship with a customer edge router.
`PE5(config-router)# neighbor 8.8.8.8 remote-as 800`	Configure BGP 8.8.8.8 peering relationship with a customer edge router.
`PE5(config-router)# neighbor 3.3.3.3 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 3.3.3.3 to establish the TCP connections.
`PE5(config-router)# neighbor 4.4.4.4 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 4.4.4.4 to establish the TCP connections.
`PE5(config-router)# neighbor 7.7.7.7 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 7.7.7.7 to establish the TCP connections.
`PE5(config-router)# neighbor 8.8.8.8 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 8.8.8.8 to establish the TCP connections.
`PE5(config-router)# address-family l2vpn evpn`	Enter to the L2 VPN address family mode to configure the address-family specific parameters.
`PE5(config-router-af)# neighbor 3.3.3.3 activate`	Enable the exchange of specific address family routes with a neighboring router 3.3.3.3.
`PE5(config-router-af)# neighbor 4.4.4.4 activate`	Enable the exchange of specific address family routes with a neighboring router 4.4.4.4.
`PE5(config-router-af)# neighbor 7.7.7.7 activate`	Enable the exchange of specific address family routes with a neighboring router 7.7.7.7.
`PE5(config-router-af)# neighbor 8.8.8.8 activate`	Enable the exchange of specific address family routes with a neighboring router 8.8.8.8.
`PE5(config-router-af)# exit-address-family`	Exit address family configuration mode.
`PE5(config-router)# exit`	Exit router configuration mode.
`PE5(config)#end`	Exit global configuration mode.

Configure Spine Node with ISIS Router

Perform the following configurations to make the node as spine with routing protocol as ISIS.

PE6

PE6(config-if)#interface sa2	Configure sa2 interface.
PE6(config-if)# load-interval 30	Configures the load-interval for monitoring traffic on the sa2 interface.
PE6(config-if)# ip address 80.1.1.2/24	Configure IP address.
PE6(config-if)# mtu 9000	Configure the Maximum Transmission Unit (MTU).
PE6(config-if)# ip router isis 1	Configure ISIS router on sa2 interface.
PE6(config-if)#interface ce2	Configure ce2 interface.
PE6(config-if)# ip address 101.1.1.1/24	Configure IP address
PE6(config-if)# mtu 9000	Configure the Maximum Transmission Unit (MTU).
PE6(config-if)# ip router isis 1	Configure ISIS router on sa2 interface.
PE6(config-if)#interface lo	Configure loopback interface
PE6(config-if)# ip address 6.6.6.6/32 secondary	Configure secondary IP address.
PE6(config-if)# ip router isis 1	Configure ISIS router on sa2 interface.
PE6(config-if)#interface xe5	Configure xe5 interface.
PE6(config-if)# ip address 102.1.1.1/24	Configure IP address
PE6(config-if)# mtu 9000	Configure the Maximum Transmission Unit (MTU).
PE6(config-if)# ip router isis 1	Configure ISIS router on xe5 interface.
PE6(config-if)#interface xe9	Configure xe9 interface.
PE6(config-if)# static-channel-group 2	Configure member port of sa2 interface.
PE6(config-if)#interface xe24	Configure xe24 interface.
PE6(config-if)# ip address 90.1.1.2/24	Configure IP address
PE6(config-if)# mtu 9000	Configure the Maximum Transmission Unit (MTU).
PE6(config-if)# ip router isis 1	Configure ISIS router on xe24 interface.
PE6(config)#router isis 1	Configure ISIS router on xe24 interface.
PE6(config-router)# is-type level-1-2	Configure IS type as level 1 and level 2.
PE6(config-router)# metric-style wide	Configure metric-style as wide.
PE6(config-router)# dynamic-hostname	Configure the hostname to advertise for the ISIS router.
PE6(config-router)# bfd all-interfaces	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
PE6(config-router)# net 49.0000.0006.0001.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
PE6(config-router)#router bgp 600	Start the eBGP router.
PE6(config-router)# bgp router-id 6.6.6.6	Configure BGP router id.
PE6(config-router)# no bgp inbound-route-filter	Disable the VPN/BGP inbound route-target filter.
PE6(config-router)# neighbor 3.3.3.3 remote-as 300	Configure BGP peering relationship with a customer edge router.
PE6(config-router)# neighbor 4.4.4.4 remote-as 400	Configure BGP peering relationship with a customer edge router.

`PE6(config-router)# neighbor 7.7.7.7 remote-as 700`	Configure BGP peering relationship with a customer edge router.
`PE6(config-router)# neighbor 8.8.8.8 remote-as 800`	Configure BGP peering relationship with a customer edge router.
`PE6(config-router)# neighbor 3.3.3.3 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 3.3.3.3 to establish the TCP connections
`PE6(config-router)# neighbor 4.4.4.4 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 4.4.4.4 to establish the TCP connections
`PE6(config-router)# neighbor 7.7.7.7 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 7.7.7.7 to establish the TCP connections
`PE6(config-router)# neighbor 8.8.8.8 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 8.8.8.8 to establish the TCP connections
`PE6(config-router)# address-family l2vpn evpn`	Enter to the L2 VPN address family mode to configure the address-family specific parameters.
`PE6(config-router-af)# neighbor 3.3.3.3 activate`	Enable the exchange of specific address family routes with a neighboring router 3.3.3.3.
`PE6(config-router-af)# neighbor 4.4.4.4 activate`	Enable the exchange of specific address family routes with a neighboring router 4.4.4.4.
`PE6(config-router-af)# neighbor 7.7.7.7 activate`	Enable the exchange of specific address family routes with a neighboring router 7.7.7.7.
`PE6(config-router-af)# neighbor 8.8.8.8 activate`	Enable the exchange of specific address family routes with a neighboring router 8.8.8.8.
`PE6(config-router-af)# exit-address-family`	Exit address family configuration mode.
`PE6(config-router)# exit`	Exit router configuration mode.
`PE6(config)#end`	Exit global configuration mode.

Configure VxLAN IRB with Single Homing OSPF and ISIS Router

Perform the following configurations to create VxLAN IRB interfaces with singlehoming for OSPF and ISIS routing process on PE3 and PE4 nodes.

PE3

`PE3(config)#nvo vxlan enable`	Enable VxLAN.
`PE3(config)#nvo vxlan irb`	Enable VxLAN IRB
`PE3(config)#ip vrf management`	Enter into VRF configuration mode.
`PE3(config-vrf)#ip vrf L3VRF1`	Create a VRF routing information base called L3VRF3 for OSPF router.
`PE3(config-vrf)# rd 56000:11`	Specify a route distinguisher for the VRF.
`PE3(config-vrf)# route-target both 100:100`	Add import and export route-target extended communities to the VRF.
`PE3(config-vrf)# l3vni 1000`	Configure the L3 Virtual Network Identifier for an IP VRF.
`PE3(config-vrf)#mac vrf L2VRF1`	Create a L2 MAC VRF to use in EVPN routes.
`PE3(config-vrf)# rd 4.4.4.4:11`	Specify a route distinguisher for the MAC VRF.

PE3(config-vrf)# route-target both 9.9.9.9:100	Add import and export route-target extended communities to the VRF.
PE3(config-if)#interface irb1001	Enter into IRB interface configuration mode to configure IRB interfaces for OSPF router.
PE3(config-irb-if)# ip vrf forwarding L3VRF1	Associate IRB interface with L3VRF3.
PE3(config-irb-if)# ip address 40.1.1.1/24	Configure an Anycast IP address.
PE3(config-irb-if)# ip ospf cost 1	Specify the cost of the link-state metric in a router-LSA.
PE3(config-vrf)#nvo vxlan vtep-ip-global 3.3.3.3	Configure the source VTEP IP address of the VxLAN tunnel.
PE3(config)#nvo vxlan id 102 ingress-replication	Add a tenant L2 VNID to the VxLAN. Specify ingress-replication to use head end replication for forwarding BUM traffic inner-vid-disabled to not carry VID out of network port
PE3(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1	Associate the L2VRF1 reachable protocol to Ethernet-VPN over BGP
PE3(config-nvo)# evpn irb1001	Configure default gateway behavior for IRB interface irb1001.
PE3(config-nvo)# vni-name VNI-102	Configure VNI name.
PE3(config-nvo)#nvo vxlan id 2002 ingress-replication	Add a tenant L2 VNID to the VxLAN. Specify ingress-replication to use head end replication for forwarding BUM traffic inner-vid-disabled to not carry VID out of network port
PE3(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1	Associate the L2VRF1 reachable protocol to Ethernet-VPN over BGP
PE3(config-nvo)# evpn irb2002	Configure default gateway behavior for IRB interface irb72002.
PE3(config-nvo)#qos enable	Enable QoS.
PE3(config)#interface sa4	Configure sa4 interface.
PE3(config-if)# switchport	Configure port as L2.
PE3(config-if)# load-interval 30	Configures the load-interval for monitoring traffic on the po1 interface.
PE3(config)#nvo vxlan access-if port-vlan sa4 100	Map the VLAN port for ISIS switch on sa4 interface to identify the VxLAN traffic and to enter NVO access interface mode.
PE3(config-acc-if-evpn)#map vpn-id 102	Map the l2vnid to an access-port.
PE3(config-if)#interface ce3	Configure ce3 interface.
PE3(config-if)# load-interval 30	Configures the load-interval for monitoring traffic on the ce3 interface.
PE3(config-if)# ip address 101.1.1.2/24	Configure IP address.
PE3(config-if)# mtu 9000	Configure the Maximum Transmission Unit (MTU).
PE3(config-if)# ip router isis 1	Configure ISIS router on ce3 interface.
PE3(config-irb-if)#interface lo	Configure loopback interface
PE3(config-if)# ip address 3.3.3.3/32 secondary	Configure secondary IP address
PE3(config-if)# ip router isis 1	Configure ISIS router on ce3 interface.
PE3(config-if)#interface xe10	Configure xe10 interface.
PE3(config-if)# static-channel-group 4	Configure member port of ce3 interface.

`PE3(config-if)#interface xe14`	Configure xe14 interface.
`PE3(config-if)# speed 10g`	Set the link speed of the xe14 interface.
`PE3(config-if)# ip address 23.1.1.1/24`	Configure IP address.
`PE3(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE3(config)#router ospf 1`	Configure OSPF router.
`PE3(config-router)# ospf router-id 3.3.3.3`	Configure router id for OSPF.
`PE3(config-router)# network 3.3.3.3/32 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 3.3.3.3/32.
`PE3(config-router)# network 23.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 23.1.1.0/24.
`PE3(config-router)# network 101.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 101.1.1.0/24.
`PE3(config-router)#router ospf 2 L3VRF1`	Configure OSPF on IRB L3VRF1.
`PE3(config-router)# redistribute bgp`	Redistribute BGP into OSPF.
`PE3(config-router)# network 40.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 40.1.1.0/24.
`PE3(config-router)#router isis 1`	Configure ISIS router on xe24 interface.
`PE3(config-router)# is-type level-1-2`	Configure IS type as level 1 and level 2.
`PE3(config-router)# metric-style wide`	Configure metric-style as wide.
`PE3(config-router)# dynamic-hostname`	Configure the hostname to advertise for the ISIS router.
`PE3(config-router)# bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
`PE3(config-router)# net 49.0000.0003.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`PE3(config-router)# redistribute connected`	Redistribute connected routes.
`PE3(config-router)#router bgp 300`	Start the eBGP router.
`PE3(config-router)# bgp router-id 3.3.3.3`	Configure BGP router id.
`PE3(config-router)# no bgp inbound-route-filter`	Disable the VPN/BGP inbound route-target filter.
`PE3(config-router)# neighbor 5.5.5.5 remote-as 500`	Configure BGP peering relationship with a customer edge router.
`PE3(config-router)# neighbor 6.6.6.6 remote-as 600`	Configure BGP peering relationship with a customer edge router.
`PE3(config-router)# neighbor 5.5.5.5 update-source lo`	Specifies that loopback interface (lo) is the source for the BGP 5.5.5.5 to establish the TCP connections
`PE3(config-router)# neighbor 6.6.6.6 advertisement-interval 0`	Configure a minimum advertisement interval between the sending of BGP routing updates.
`PE3(config-router)# address-family l2vpn evpn`	Enter to the L2 VPN address family mode to configure the address-family specific parameters.
`PE3(config-router-af)# neighbor 5.5.5.5 activate`	Enable the exchange of specific address family routes with a neighboring router 5.5.5.5.

`PE3(config-router-af)# neighbor 6.6.6.6 activate`	Enable the exchange of specific address family routes with a neighboring router 6.6.6.6.
`PE3(config-router-af)# exit-address-family`	Exit address family configuration mode.
`PE3(config-router)# address-family ipv4 vrf L3VRF1`	Enter to the IRB IPv4 VRF address family mode to configure the address-family specific parameters.
`PE3(config-router-af)# redistribute connected`	Redistribute connected routes.
`PE3(config-router-af)# redistribute ospf`	Redistribute OSPF routes.
`PE3(config-router-af)# exit-address-family`	Exit address family configuration mode.
`PE3(config-router)# exit`	Exit router configuration mode.
`PE3(config)#end`	Exit global configuration mode.

PE4

`PE4(config)#nvo vxlan enable`	Enable VxLAN.
`PE4(config)#nvo vxlan irb`	Enable VxLAN IRB
`PE4(config)#ip vrf management`	Enter into VRF configuration mode.
`PE4(config-vrf)#ip vrf L3VRF2`	reate a VRF routing information base called L3VRF2 for OSPF router.
`PE4(config-vrf)# rd 63000:11`	Specify a route distinguisher for the VRF.
`PE4(config-vrf)# route-target both 101:101`	dd import and export route-target extended communities to the VRF.
`PE4(config-vrf)# l3vni 2000`	Configure the L3 Virtual Network Identifier for an IP VRF.
`PE4(config-vrf)#mac vrf L2VRF2`	Create a L2 MAC VRF to use in EVPN routes.
`PE4(config-vrf)# rd 3.3.3.3:11`	Specify a route distinguisher for the MAC VRF.
`PE4(config-vrf)# route-target both 10.10.10.10:100`	Add import and export route-target extended communities to the VRF.
`PE4(config-if)#interface irb2001`	Enter into IRB interface configuration mode to configure IRB interfaces for OSPF router.
`PE4(config-irb-if)# ip vrf forwarding L3VRF2`	Associate IRB interface with L3VRF3.
`PE4(config-irb-if)# ip address 50.50.50.1/24`	Configure an Anycast IP address.
`PE4(config-irb-if)# ip router isis 2`	Configure ISIS router on IRB interface.
`PE4(config-vrf)#nvo vxlan vtep-ip-global 4.4.4.4`	Configure the source VTEP IP address of the VxLAN tunnel.
`PE4(config)#nvo vxlan id 201 ingress-replication`	Add a tenant L2 VNID to the VxLAN. Specify ingress-replication to use head end replication for forwarding BUM traffic
`PE4(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Associate the L2VRF1 reachable protocol to Ethernet-VPN over BGP
`PE4(config-nvo)# evpn irb2001`	Configure default gateway behavior for IRB interface irb2001.
`PE4(config-nvo)# vni-name VNI-201`	Configure VNI name.
`PE4(config-nvo)#nvo vxlan id 3002 ingress-replication`	Add a tenant L2 VNID to the VxLAN. Specify ingress-replication to use head end replication for forwarding BUM traffic.

`PE4(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Associate the L2VRF2 reachable protocol to Ethernet-VPN over BGP.
`PE4(config-nvo)# evpn irb3002`	Configure default gateway behavior for IRB interface irb3002.
`PE4(config-nvo)#nvo vxlan id 3003 ingress-replication`	Add a tenant L2 VNID to the VxLAN. Specify ingress-replication to use head end replication for forwarding BUM traffic.
`PE4(config-nvo)#qos enable`	Enable QoS.
`PE4(config-irb-if)#interface lo`	Configure loopback interface
`PE4(config-if)# ip address 4.4.4.4/32 secondary`	Configure secondary IP address
`PE4(config-if)# ip router isis 1`	Configure ISIS router on lo interface.
`PE4(config-if)#interface xe5`	Configure xe5 interface.
`PE4(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the xe5 interface.
`PE4(config-if)# ip address 102.1.1.2/24`	Configure IP address.
`PE4(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE4(config-if)# ip router isis 1`	Configure ISIS router on ce3 interface.
`PE4(config-if)#interface xe25`	Configure xe25 interface.
`PE4(config-if)# ip address 24.1.1.1/24`	Configure IP address.
`PE4(config-if)# mtu 9000`	Configure the Maximum Transmission Unit (MTU).
`PE4(config-if)#interface xe26`	Configure xe26 interface.
`PE4(config-if)# switchport`	Configure port as L2.
`PE4(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the xe26 interface.
`PE4(config)#nvo vxlan access-if port-vlan xe26 200`	Map the VLAN port for ISIS switch on xe26 interface to identify the VxLAN traffic and to enter NVO access interface mode.
`PE4(config-acc-if-evpn)#  map vpn-id 201`	Map the l2vnid to an access-port.
`PE4(config-if)# exit`	Exit the interface configuration mode.
`PE4(config)#router ospf 1`	Configure OSPF router.
`PE4(config-router)# network 4.4.4.4/32 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 4.4.4.4/32.
`PE4(config-router)# network 24.1.1.0/24 area 0.0.0.0`	Enable OSPF routing with area ID 0.0.0.0 on IRB interface with IP address that match the network address 24.1.1.0/24.
`PE4(config-router)#router isis 1`	Configure ISIS router on xe24 interface.
`PE4(config-router)# is-type level-1-2`	Configure IS type as level 1 and level 2.
`PE4(config-router)# metric-style wide`	Configure metric-style as wide.
`PE4(config-router)# dynamic-hostname`	Configure the hostname to advertise for the ISIS router.
`PE4(config-router)# bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
`PE4(config-router)# net 49.0000.0004.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`PE4(config-router)# redistribute connected`	Redistribute connected routes.

PE4(config-router)#!	
PE4(config-router)#router isis 2 L3VRF2	Configure ISIS on IRB L3VRF2.
PE4(config-router)# is-type level-1-2	Configure IS type as level 1 and level 2.
PE4(config-router)# metric-style wide	Configure metric-style as wide.
PE4(config-router)# dynamic-hostname	Configure the hostname to advertise for the ISIS router.
PE4(config-router)# bfd all-interfaces	Enable the Bidirectional Forwarding Detection (BFD) feature on all the interfaces enabled with this ISIS instance.
PE4(config-router)# net 49.0000.0000.0441.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
PE4(config-router)#router bgp 400	Start the eBGP router.
PE4(config-router)# bgp router-id 4.4.4.4	Configure BGP router id.
PE4(config-router)# no bgp inbound-route-filter	Disable the VPN/BGP inbound route-target filter.
PE4(config-router)# neighbor 5.5.5.5 remote-as 500	Configure BGP peering relationship with a customer edge router.
PE4(config-router)# neighbor 6.6.6.6 remote-as 600	Configure BGP peering relationship with a customer edge router.
PE4(config-router)# neighbor 5.5.5.5 update-source lo	Specifies that loopback interface (lo) is the source for the BGP 5.5.5.5 to establish the TCP connections
PE4(config-router)# neighbor 6.6.6.6 advertisement-interval 0	Configure a minimum advertisement interval between the sending of BGP routing updates.
PE4(config-router)# address-family l2vpn evpn	Enter to the L2 VPN address family mode to configure the address-family specific parameters.
PE4(config-router-af)# neighbor 5.5.5.5 activate	Enable the exchange of specific address family routes with a neighboring router 5.5.5.5.
PE4(config-router-af)# neighbor 6.6.6.6 activate	Enable the exchange of specific address family routes with a neighboring router 6.6.6.6.
PE4(config-router-af)# exit-address-family	Exit address family configuration mode.
PE4(config-router)# address-family ipv4 vrf L3VRF2	Enter to the IRB IPv4 VRF address family mode to configure the address-family specific parameters.
PE4(config-router-af)# redistribute connected	Redistribute connected routes.
PE4(config-router-af)# exit-address-family	Exit address family configuration mode.
PE4(config-router)# exit	Exit router configuration mode.
PE4(config)#end	Exit global configuration mode.

Validation

```
PE1#show ip ospf neighbor
```

```
Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface         Instance ID
17.12.13.1        1   Full/DR         00:00:35    16.16.16.2      vlan1.200         0
17.17.17.2        1   Full/DR         00:00:31    17.17.17.2      vlan1.200         0

PE1#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, 00:23:27
O E2         3.3.3.3/32 [110/1] via 16.16.16.2, vlan1.200, 00:18:02
                        [110/1] via 17.17.17.2, vlan1.200
O E2         7.7.7.7/32 [110/1] via 17.17.17.2, vlan1.200, 00:18:45
O E2         8.8.8.8/32 [110/1] via 16.16.16.2, vlan1.200, 00:19:14
C            16.16.16.0/24 is directly connected, vlan1.200, 00:21:24
C            17.17.1.0/24 is directly connected, vlan1.200, 00:21:24
C            17.17.17.0/24 is directly connected, vlan1.200, 00:21:24
C            20.1.1.0/24 is directly connected, xe2.103, 00:22:31
O E2         40.1.1.0/24 [110/1] via 16.16.16.2, vlan1.200, 00:18:02
                         [110/1] via 17.17.17.2, vlan1.200
C            101.11.11.0/24 is directly connected, vlan1.200, 00:21:24
C            127.0.0.0/8 is directly connected, lo, 00:23:27

IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 00:23:27
C            127.0.0.0/8 is directly connected, lo.management, 00:23:27

Gateway of last resort is not set
PE1#
PE1#ping 40.1.1.1
Press CTRL+C to exit
PING 40.1.1.1 (40.1.1.1) 56(84) bytes of data.
64 bytes from 40.1.1.1: icmp_seq=1 ttl=63 time=0.678 ms
64 bytes from 40.1.1.1: icmp_seq=2 ttl=63 time=0.568 ms
64 bytes from 40.1.1.1: icmp_seq=3 ttl=63 time=0.567 ms
64 bytes from 40.1.1.1: icmp_seq=4 ttl=63 time=0.657 ms

--- 40.1.1.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 59ms
rtt min/avg/max/mdev = 0.567/0.617/0.678/0.056 ms
PE1#
```

```
PE2#show clns neighbors

Total number of L1 adjacencies: 4
Total number of L2 adjacencies: 4
Total number of adjacencies: 8
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE7            vlan1.100   0000.0000.1111  Up     21        L1   IS-IS
                                           Up     21        L2   IS-IS
PE8            vlan1.100   0000.0000.1111  Up     20        L1   IS-IS
                                           Up     20        L2   IS-IS
PE7            vlan1.101   0000.0000.1111  Up     21        L1   IS-IS
                                           Up     21        L2   IS-IS
PE8            vlan1.101   0000.0000.1111  Up     20        L1   IS-IS
                                           Up     20        L2   IS-IS

PE2#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            2.2.2.2/32 is directly connected, lo, 00:25:56
i L2         4.4.4.4/32 [115/10] via 13.1.1.2, vlan1.100, 00:22:51
                        [115/10] via 12.1.1.2, vlan1.100
i L2         7.7.7.7/32 [115/10] via 12.1.1.2, vlan1.100, 00:24:20
i L2         8.8.8.8/32 [115/10] via 13.1.1.2, vlan1.100, 00:24:03
C            10.1.1.0/24 is directly connected, xe11.105, 00:25:23
C            11.11.11.0/24 is directly connected, vlan1.100, 00:25:22
C            12.1.1.0/24 is directly connected, vlan1.100, 00:25:22
i L1         12.13.14.0/24 [115/20] via 13.1.1.2, vlan1.100, 00:24:20
                           [115/20] via 12.1.1.2, vlan1.100
C            13.1.1.0/24 is directly connected, vlan1.100, 00:25:22
i L2         50.50.50.0/24 [115/10] via 13.1.1.2, vlan1.100, 00:22:51
                           [115/10] via 12.1.1.2, vlan1.100
C            127.0.0.0/8 is directly connected, lo, 00:25:56

IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 00:25:56
C            127.0.0.0/8 is directly connected, lo.management, 00:25:56

Gateway of last resort is not set
PE2#ping 50.50.50.1
Press CTRL+C to exit
PING 50.50.50.1 (50.50.50.1) 56(84) bytes of data.
```

```
64 bytes from 50.50.50.1: icmp_seq=1 ttl=63 time=0.491 ms
64 bytes from 50.50.50.1: icmp_seq=2 ttl=63 time=0.411 ms
64 bytes from 50.50.50.1: icmp_seq=3 ttl=63 time=0.628 ms
64 bytes from 50.50.50.1: icmp_seq=4 ttl=63 time=0.661 ms

--- 50.50.50.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 65ms
rtt min/avg/max/mdev = 0.411/0.547/0.661/0.105 ms
PE2#


PE7#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE6            sa2         e8c5.7a19.c3c8  Up     5         L1   IS-IS
                                           Up     5         L2   IS-IS

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 2
Total number of adjacencies: 4
Tag 2:  VRF : L3VRF4
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE2            irb801      e8c5.7a76.581d  Up     5         L1   IS-IS
                                           Up     5         L2   IS-IS
PE2            irb802      e8c5.7a76.581d  Up     5         L1   IS-IS
                                           Up     5         L2   IS-IS

PE7#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri  State           Dead Time   Address      Interface   Instance ID
5.5.5.5         1    Full/Backup     00:00:34    22.1.1.2     sa3         0

Total number of full neighbors: 1
OSPF process 2 VRF(L3VRF3):
Neighbor ID     Pri   State          Dead Time   Address       Interface    Instance ID
1.1.1.1         1   Full/Backup      00:00:34    16.16.16.1    irb701       0




PE7#show ip route 2023 Nov 27 13:04:32.790 : PE7 : HSL : NOTIF : [IF_PKT_ERRORS_4]:
Fragment packets received on xe43 (1 packets)

vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
```

```
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            3.3.3.3/32 [110/3] via 22.1.1.2, sa3, 00:23:42
O            4.4.4.4/32 [110/3] via 22.1.1.2, sa3, 00:23:27
O            5.5.5.5/32 [110/2] via 22.1.1.2, sa3, 00:26:13
i L1         6.6.6.6/32 [115/20] via 80.1.1.2, sa2, 00:25:44
C            7.7.7.7/32 is directly connected, lo, 00:29:00
O            8.8.8.8/32 [110/3] via 22.1.1.2, sa3, 00:26:09
O            21.1.1.0/24 [110/2] via 22.1.1.2, sa3, 00:26:13
C            22.1.1.0/24 is directly connected, sa3, 00:27:04
O            23.1.1.0/24 [110/2] via 22.1.1.2, sa3, 00:24:37
O            24.1.1.0/24 [110/2] via 22.1.1.2, sa3, 00:24:20
C            80.1.1.0/24 is directly connected, sa2, 00:25:59
i L1         90.1.1.0/24 [115/20] via 80.1.1.2, sa2, 00:25:44
O            101.1.1.0/24 [110/3] via 22.1.1.2, sa3, 00:23:42
i L1         102.1.1.0/24 [115/20] via 80.1.1.2, sa2, 00:23:31
C            127.0.0.0/8 is directly connected, lo, 00:29:00

IP Route Table for VRF "management"
C            10.12.93.0/24 is directly connected, eth0, 00:29:00
C            127.0.0.0/8 is directly connected, lo.management, 00:29:00

IP Route Table for VRF "L3VRF3"
B            3.3.3.3/32 [0/0] is directly connected, tunvxlan2, 00:23:42
B            8.8.8.8/32 [0/0] is directly connected, tunvxlan2, 00:25:09
C            16.16.16.0/24 is directly connected, irb701, 00:29:00
C            17.12.13.0/24 is directly connected, irb701, 00:29:00
O            17.17.17.0/24 [110/2] via 16.16.16.1, irb701, 00:26:33
O            20.1.1.0/24 [110/2] via 16.16.16.1, irb701, 00:26:33
B          40.1.1.0/24 [20/0] via 3.3.3.3 (recursive is directly connected, tunvxlan2),
00:23:56
C            127.0.0.0/8 is directly connected, lo.L3VRF3, 00:29:00

IP Route Table for VRF "L3VRF4"
i L2         2.2.2.2/32 [115/10] via 13.1.1.1, irb801, 00:25:22
B            4.4.4.4/32 [0/0] is directly connected, tunvxlan3, 00:23:27
B            8.8.8.8/32 [0/0] is directly connected, tunvxlan3, 00:25:09
i L1         10.1.1.0/24 [115/20] via 13.1.1.1, irb801, 00:25:24
i L1         11.11.11.0/24 [115/20] via 13.1.1.1, irb801, 00:25:24
i L1         12.1.1.0/24 [115/20] via 13.1.1.1, irb801, 00:25:24
C            12.13.14.0/24 is directly connected, irb801, 00:29:00
C            13.1.1.0/24 is directly connected, irb801, 00:29:00
B            50.50.50.0/24 [20/0] via 4.4.4.4 (recursive is directly connected,
tunvxlan3), 00:23:56
```

```
C            127.0.0.0/8 is directly connected, lo.L3VRF4, 00:29:00
IP Route Table for VRF "L2VRF1"
IP Route Table for VRF "L2VRF2"
IP Route Table for VRF "evpn-gvrf-1"

Gateway of last resort is not set

PE7# show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================
7.7.7.7          8.8.8.8          Installed     00:25:36      00:25:36
7.7.7.7          4.4.4.4          Installed     00:23:53      00:23:53
7.7.7.7          3.3.3.3          Installed     00:24:08      00:24:08

Total number of entries are 3
PE7#show nvo vxlan l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 2000         801          irb801
 2000         8001         irb802
 1000         701          irb701
 1000         7001         irb702

PE7#show 2023 Nov 27 13:05:12.791 : PE7 : HSL : NOTIF : [IF_PKT_ERRORS_4]: Fragment
packets received on xe43 (1 packets)

nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
701      VNI-701      L2       NW   ----      ----                            ----      ----      7.7.7.7          8.8.8.8
701      VNI-701      --       AC   sa1       00:00:01:02:03:04:05:06:07:08 200       DF        ----             ----
801      VNI-101      L2       NW   ----      ----                            ----      ----      7.7.7.7          8.8.8.8
801      VNI-101      --       AC   po1       00:00:00:00:00:77:82:00:00:00 100       DF        ----             ----
1000     ----         L3       NW   ----      ----                            ----      ----      7.7.7.7          8.8.8.8
1000     ----         L3       NW   ----      ----                            ----      ----      7.7.7.7          3.3.3.3
2000     ----         L3       NW   ----      ----                            ----      ----      7.7.7.7          8.8.8.8
2000     ----         L3       NW   ----      ----                            ----      ----      7.7.7.7          4.4.4.4
7001     ----         L2       NW   ----      ----                            ----      ----      7.7.7.7          8.8.8.8
7001     ----         --       AC   sa1       00:00:01:02:03:04:05:06:07:08 201       NON-DF    ----             ----
8001     ----         L2       NW   ----      ----                            ----      ----      7.7.7.7          8.8.8.8
8001     ----         --       AC   po1       00:00:00:00:00:77:82:00:00:00 101       NON-DF    ----             ----

Total number of entries are 12
PE7#

PE7#show bgp l2vpn evpn prefix-route
```

```
RD[8100:11]
ESI             Eth-Tag Prefix-Length   IP-Address      GW-IPAddress    L3VNID/LABEL    Nexthop         Encap     Router-Mac
0               0       24              17.17.17.0      0.0.0.0         1000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              17.17.17.0      0.0.0.0         1000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              30::            ::              1000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              30::            ::              1000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              7102::          ::              1000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              7102::          ::              1000            8.8.8.8         VXLAN     1444:8f53:3e9e

RD[8400:11]
ESI             Eth-Tag Prefix-Length   IP-Address      GW-IPAddress    L3VNID/LABEL    Nexthop         Encap     Router-Mac
0               0       24              10.1.1.0        0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              10.1.1.0        0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              11.11.11.0      0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              11.11.11.0      0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              12.1.1.0        0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              12.1.1.0        0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              12.13.14.0      0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              12.13.14.0      0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              13.1.1.0        0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       24              13.1.1.0        0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       32              2.2.2.2         0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       32              2.2.2.2         0.0.0.0         2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              2::             ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              2::             ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              40::            ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              40::            ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              50::            ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              50::            ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              8002::          ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              8002::          ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              8102::          ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e
0               0       64              8102::          ::              2000            8.8.8.8         VXLAN     1444:8f53:3e9e

RD[56000:11]
ESI             Eth-Tag Prefix-Length   IP-Address      GW-IPAddress    L3VNID/LABEL    Nexthop         Encap     Router-Mac
0               0       24              40.1.1.0        0.0.0.0         1000            3.3.3.3         VXLAN     e49d:73b1:c301
0               0       24              40.1.1.0        0.0.0.0         1000            3.3.3.3         VXLAN     e49d:73b1:c301
0               0       64              2002::          ::              1000            3.3.3.3         VXLAN     e49d:73b1:c301
0               0       64              2002::          ::              1000            3.3.3.3         VXLAN     e49d:73b1:c301

RD[63000:11]
ESI             Eth-Tag Prefix-Length   IP-Address      GW-IPAddress    L3VNID/LABEL    Nexthop         Encap     Router-Mac
0               0       24              50.50.50.0      0.0.0.0         2000            4.4.4.4         VXLAN     e8c5:7aa8:7cb3
0               0       24              50.50.50.0      0.0.0.0         2000            4.4.4.4         VXLAN     e8c5:7aa8:7cb3
0               0       64              3002::          ::              2000            4.4.4.4         VXLAN     e8c5:7aa8:7cb3
0               0       64              3002::          ::              2000            4.4.4.4         VXLAN     e8c5:7aa8:7cb3
PE7#PE8#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE6            xe2         e8c5.7a19.c3c1  Up     7         L1   IS-IS
                                           Up     7         L2   IS-IS

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 2
Total number of adjacencies: 4
Tag 2:  VRF : L3VRF4
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE2            irb801      e8c5.7a76.581d  Up     7         L1   IS-IS
                                           Up     7         L2   IS-IS
```

```
PE2            irb802      e8c5.7a76.581d      Up     7         L1   IS-IS
                                               Up     7         L2   IS-IS
PE8#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State            Dead Time   Address      Interface Instance ID
5.5.5.5           1   Full/Backup      00:00:31    21.1.1.2        xe12         0

Total number of full neighbors: 1
OSPF process 2 VRF(L3VRF3):
Neighbor ID     Pri   State           Dead Time   Address      Interface     Instance ID
1.1.1.1           1  Full/Backup     00:00:35    17.17.17.1    irb701       0
PE8#terminal width 511

PE8#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O             3.3.3.3/32 [110/3] via 21.1.1.2, xe12, 00:26:22
O             4.4.4.4/32 [110/3] via 21.1.1.2, xe12, 00:26:07
O             5.5.5.5/32 [110/2] via 21.1.1.2, xe12, 00:28:59
i L1          6.6.6.6/32 [115/20] via 90.1.1.2, xe2, 00:28:23
O             7.7.7.7/32 [110/3] via 21.1.1.2, xe12, 00:28:52
C             8.8.8.8/32 is directly connected, lo, 00:31:21
C             21.1.1.0/24 is directly connected, xe12, 00:29:44
O             22.1.1.0/24 [110/2] via 21.1.1.2, xe12, 00:28:59
O             23.1.1.0/24 [110/2] via 21.1.1.2, xe12, 00:27:17
O             24.1.1.0/24 [110/2] via 21.1.1.2, xe12, 00:27:00
i L1          80.1.1.0/24 [115/20] via 90.1.1.2, xe2, 00:28:23
C             90.1.1.0/24 is directly connected, xe2, 00:28:39
O             101.1.1.0/24 [110/3] via 21.1.1.2, xe12, 00:26:22
i L1          102.1.1.0/24 [115/20] via 90.1.1.2, xe2, 00:26:11
C             127.0.0.0/8 is directly connected, lo, 00:31:21

IP Route Table for VRF "management"
C             10.12.93.0/24 is directly connected, eth0, 00:31:21
C             127.0.0.0/8 is directly connected, lo.management, 00:31:21

IP Route Table for VRF "L3VRF3"
B             3.3.3.3/32 [0/0] is directly connected, tunvxlan2, 00:26:22
B             7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:28:17
B             16.16.16.0/24 [20/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:28:17
```

```
C            17.12.13.0/24 is directly connected, irb701, 00:31:21
C            17.17.17.0/24 is directly connected, irb701, 00:31:21
B          20.1.1.0/24 [20/0] via 7.7.7.7 (recursive is directly connected, tunvxlan2),
00:28:17
B          40.1.1.0/24 [20/0] via 3.3.3.3 (recursive is directly connected, tunvxlan2),
00:26:37
C            127.0.0.0/8 is directly connected, lo.L3VRF3, 00:31:21

IP Route Table for VRF "L3VRF4"
i L2         2.2.2.2/32 [115/10] via 12.1.1.1, irb801, 00:28:44
B            4.4.4.4/32 [0/0] is directly connected, tunvxlan3, 00:26:07
B            7.7.7.7/32 [0/0] is directly connected, tunvxlan3, 00:28:17
i L1         10.1.1.0/24 [115/20] via 12.1.1.1, irb801, 00:28:44
i L1         11.11.11.0/24 [115/20] via 12.1.1.1, irb801, 00:28:44
C            12.1.1.0/24 is directly connected, irb801, 00:31:21
C            12.13.14.0/24 is directly connected, irb801, 00:31:21
i L1         13.1.1.0/24 [115/20] via 12.1.1.1, irb801, 00:28:44
B            50.50.50.0/24 [20/0] via 4.4.4.4 (recursive is directly connected,
tunvxlan3), 00:26:37
C            127.0.0.0/8 is directly connected, lo.L3VRF4, 00:31:21
IP Route Table for VRF "L2VRF1"
IP Route Table for VRF "L2VRF2"
IP Route Table for VRF "evpn-gvrf-1"

Gateway of last resort is not set
PE8#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
8.8.8.8          7.7.7.7          Installed     00:28:24      00:28:24
8.8.8.8          3.3.3.3          Installed     00:26:28      00:26:28
8.8.8.8          4.4.4.4          Installed     00:26:13      00:26:13

Total number of entries are 3
PE8#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
___
701      VNI-701      L2       NW   ----      ----                            ----      ----      8.8.8.8          7.7.7.7
701      VNI-701      --       AC   sa1       00:00:01:02:03:04:05:06:07:08 200       NON-DF    ----             ----
801      VNI-101      L2       NW   ----      ----                            ----      ----      8.8.8.8          7.7.7.7
801      VNI-101      --       AC   po1       00:00:00:00:00:77:82:00:00:00 100       NON-DF    ----             ----
1000     ----         L3       NW   ----      ----                            ----      ----      8.8.8.8          7.7.7.7
1000     ----         L3       NW   ----      ----                            ----      ----      8.8.8.8          3.3.3.3
2000     ----         L3       NW   ----      ----                            ----      ----      8.8.8.8          7.7.7.7
2000     ----         L3       NW   ----      ----                            ----      ----      8.8.8.8          4.4.4.4
7001     ----         L2       NW   ----      ----                            ----      ----      8.8.8.8          7.7.7.7
7001     ----         --       AC   sa1       00:00:01:02:03:04:05:06:07:08 201       DF        ----             ----
8001     ----         L2       NW   ----      ----                            ----      ----      8.8.8.8          7.7.7.7
8001     ----         --       AC   po1       00:00:00:00:00:77:82:00:00:00 101       DF        ----             ----
```

```
Total number of entries are 12

PE8#show bgp l2vpn evpn prefix-route

RD[7100:11]
ESI            Eth-Tag Prefix-Length   IP-Address      GW-IPAddress   L3VNID/LABEL   Nexthop         Encap     Router-Mac
0              0       24              16.16.16.0      0.0.0.0        1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              16.16.16.0      0.0.0.0        1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              17.17.17.0      0.0.0.0        1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              17.17.17.0      0.0.0.0        1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              20.1.1.0        0.0.0.0        1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              20.1.1.0        0.0.0.0        1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              30::            ::             1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              30::            ::             1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              7002::          ::             1000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              7002::          ::             1000           7.7.7.7         VXLAN     e001:a666:056d

RD[7400:11]
ESI            Eth-Tag Prefix-Length   IP-Address     GW-IPAddress    L3VNID/LABEL   Nexthop         Encap     Router-Mac
0              0       24              10.1.1.0        0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              10.1.1.0        0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              11.11.11.0      0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              11.11.11.0      0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              12.1.1.0        0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              12.1.1.0        0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              12.13.14.0      0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              12.13.14.0      0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              13.1.1.0        0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       24              13.1.1.0        0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       32              2.2.2.2         0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       32              2.2.2.2         0.0.0.0        2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              2::             ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              2::             ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              40::            ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              40::            ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              50::            ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              50::            ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              8002::          ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              8002::          ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              8102::          ::             2000           7.7.7.7         VXLAN     e001:a666:056d
0              0       64              8102::          ::             2000           7.7.7.7         VXLAN     e001:a666:056d

RD[56000:11]
ESI            Eth-Tag Prefix-Length   IP-Address     GW-IPAddress    L3VNID/LABEL   Nexthop         Encap    Router-Mac
0              0       24              40.1.1.0        0.0.0.0        1000           3.3.3.3         VXLAN     e49d:73b1:c301
0              0       24              40.1.1.0        0.0.0.0        1000           3.3.3.3         VXLAN     e49d:73b1:c301
0              0       64              2002::          ::             1000           3.3.3.3         VXLAN     e49d:73b1:c301
0              0       64              2002::          ::             1000           3.3.3.3         VXLAN     e49d:73b1:c301

RD[63000:11]
ESI            Eth-Tag Prefix-Length   IP-Address     GW-IPAddress    L3VNID/LABEL   Nexthop        Encap     Router-Mac
0              0       24              50.50.50.0      0.0.0.0        2000           4.4.4.4         VXLAN     e8c5:7aa8:7cb3
0              0       24              50.50.50.0      0.0.0.0        2000           4.4.4.4         VXLAN     e8c5:7aa8:7cb3
0              0       64              3002::          ::             2000           4.4.4.4         VXLAN     e8c5:7aa8:7cb3
0              0       64              3002::          ::             2000           4.4.4.4         VXLAN     e8c5:7aa8:7cb3
PE8#


PE3#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address      Interface      Instance ID
5.5.5.5           1   Full/DR         00:00:29    23.1.1.2      xe14          0
```

```
Total number of full neighbors: 1
OSPF process 2 VRF(L3VRF1):
Neighbor ID     Pri   State           Dead Time   Address     Interface       Instance ID
192.0.0.2         0   Full/DROther     00:00:35    40.1.1.2   irb1001           0

PE3#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            3.3.3.3/32 is directly connected, lo, 00:37:01
O            4.4.4.4/32 [110/3] via 23.1.1.2, xe14, 00:35:25
O            5.5.5.5/32 [110/2] via 23.1.1.2, xe14, 00:35:44
i L1         6.6.6.6/32 [115/20] via 101.1.1.1, ce3, 00:36:20
O            7.7.7.7/32 [110/3] via 23.1.1.2, xe14, 00:35:44
O            8.8.8.8/32 [110/3] via 23.1.1.2, xe14, 00:35:44
O            21.1.1.0/24 [110/2] via 23.1.1.2, xe14, 00:35:44
O            22.1.1.0/24 [110/2] via 23.1.1.2, xe14, 00:35:44
C            23.1.1.0/24 is directly connected, xe14, 00:36:36
O            24.1.1.0/24 [110/2] via 23.1.1.2, xe14, 00:35:44
i L1         80.1.1.0/24 [115/20] via 101.1.1.1, ce3, 00:36:20
i L1         90.1.1.0/24 [115/20] via 101.1.1.1, ce3, 00:36:20
C            101.1.1.0/24 is directly connected, ce3, 00:36:35
i L1         102.1.1.0/24 [115/20] via 101.1.1.1, ce3, 00:35:30
C            127.0.0.0/8 is directly connected, lo, 00:37:01

IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 00:37:01
C            127.0.0.0/8 is directly connected, lo.management, 00:37:01

IP Route Table for VRF "L3VRF1"
B            7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:35:44
B            8.8.8.8/32 [0/0] is directly connected, tunvxlan2, 00:35:44
B            16.16.16.0/24 [20/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:36:15
B            17.17.17.0/24 [20/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:36:15
B            20.1.1.0/24 [20/0] via 7.7.7.7 (recursive is directly connected, tu
nvxlan2), 00:36:15
C            40.1.1.0/24 is directly connected, irb1001, 00:37:01
O            55.0.0.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
O            55.0.1.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
O            55.0.2.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
O            55.0.3.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
```

```
O            55.0.4.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
O            55.0.5.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
O            55.0.6.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
O            55.0.7.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
O            55.0.8.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
O            55.0.9.0/24 [110/2] via 40.1.1.2, irb1001, 00:01:38
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:37:01

IP Route Table for VRF "L2VRF1"

Gateway of last resort is not set
PE3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================
3.3.3.3          7.7.7.7          Installed     00:35:56      00:35:56
3.3.3.3          8.8.8.8          Installed     00:35:56      00:35:56

Total number of entries are 2
PE3#

PE4#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE6            xe5         e8c5.7a19.c3ae  Up     21        L1   IS-IS
                                           Up     21        L2   IS-IS
Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0007  Up     26        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0010.9400.0009 irb3002     0010.9400.0009  Up     25        L2   IS-IS
PE4#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
```

```
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O             3.3.3.3/32 [110/3] via 24.1.1.2, xe25, 00:36:12
C             4.4.4.4/32 is directly connected, lo, 00:37:52
O             5.5.5.5/32 [110/2] via 24.1.1.2, xe25, 00:36:12
i L1          6.6.6.6/32 [115/20] via 102.1.1.1, xe5, 00:36:48
O             7.7.7.7/32 [110/3] via 24.1.1.2, xe25, 00:36:12
O             8.8.8.8/32 [110/3] via 24.1.1.2, xe25, 00:36:12
O             21.1.1.0/24 [110/2] via 24.1.1.2, xe25, 00:36:12
O             22.1.1.0/24 [110/2] via 24.1.1.2, xe25, 00:36:12
O             23.1.1.0/24 [110/2] via 24.1.1.2, xe25, 00:36:12
C             24.1.1.0/24 is directly connected, xe25, 00:37:04
i L1          80.1.1.0/24 [115/20] via 102.1.1.1, xe5, 00:36:48
i L1          90.1.1.0/24 [115/20] via 102.1.1.1, xe5, 00:36:48
O             101.1.1.0/24 [110/3] via 24.1.1.2, xe25, 00:36:12
C             102.1.1.0/24 is directly connected, xe5, 00:37:04
C             127.0.0.0/8 is directly connected, lo, 00:37:52
IP Route Table for VRF "management"
C             10.12.98.0/24 is directly connected, eth0, 00:37:52
C             127.0.0.0/8 is directly connected, lo.management, 00:37:52
IP Route Table for VRF "L3VRF2"
Gateway of last resort is 7.7.7.7 to network 0.0.0.0

B*            0.0.0.0/0 [20/0] via 7.7.7.7 (recursive is directly connected, tunv
xlan2), 00:01:50
B             2.2.2.2/32 [20/0] via 7.7.7.7 (recursive is directly connected, tun
vxlan2), 00:36:41
B             7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:36:12
B             8.8.8.8/32 [0/0] is directly connected, tunvxlan2, 00:36:12
B             10.1.1.0/24 [20/0] via 7.7.7.7 (recursive is directly connected, tu
nvxlan2), 00:36:41
B             11.11.11.0/24 [20/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:36:41
B             12.1.1.0/24 [20/0] via 7.7.7.7 (recursive is directly connected, tu
nvxlan2), 00:36:41
B             12.13.14.0/24 [20/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:36:41
B             13.1.1.0/24 [20/0] via 7.7.7.7 (recursive is directly connected, tu
nvxlan2), 00:36:41
C             50.50.50.0/24 is directly connected, irb2001, 00:37:52
C             127.0.0.0/8 is directly connected, lo.L3VRF2, 00:37:52
IP Route Table for VRF "L2VRF2"

Gateway of last resort is not set
PE4#  show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================
```

```
4.4.4.4          7.7.7.7          Installed     00:36:18      00:36:18
4.4.4.4          8.8.8.8          Installed     00:36:18      00:36:18

Total number of entries are 2
PE4#
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
ECMP	Equal-Cost Multipath
EVPN	Ethernet Virtual Private Netwrok
VxLAN	Virtual Extensible LAN
SR	Segment Routing
IRB	Integrated Routing
OSPF	Open Shortest Path First
ISIS	Intermediate System to Intermediate System

Glossary

The following provides definitions for key terms used throughout this document.

Multi Home VxLAN	This refers to a Virtual Extensible LAN (VxLAN) deployment where a Multi data center or network site is connected to a Multi external network (usually the internet) for connectivity.
IRB	A networking feature that enables the integration of Layer 3 IP routing and Layer 2 MAC address bridging within the same interface, simplifying network management and resource utilization.
OSPF	A dynamic and efficient link-state routing protocol used to determine the best path for data packets in an IP network. It is characterized by rapid convergence and adaptability, making it suitable for large and dynamic networks.
ISIS	A routing protocol designed for scalability and stability in computer networks, commonly used in large Service Provider networks. It provides a robust framework for routing information exchange.
Layer 3 Routing	Network routing operations at the Network Layer (Layer 3) of the OSI model, focusing on routing IP packets between different subnets or networks.
Layer 2 Bridging	Network bridging operations at the Data Link Layer (Layer 2) of the OSI model, handling the forwarding of data frames based on MAC addresses within the same network segment.
EVPN	Ethernet VPN, a technology that provides advanced and efficient methods for Layer 2 and Layer 3 services in Ethernet networks, often used in data centers and service provider environments.

CHAPTER 9 Single Hop BFD over IRB

Overview

The Single-hop Bidirectional Forwarding Detection (BFD) with BGP over EVPN-VXLAN IRB feature provides a lightweight, software-based failure detection mechanism for routed interfaces in virtualized data center networks. This feature enables BFD sessions over Integrated Routing and Bridging (IRB) interfaces without requiring hardware BFD support, ensuring compatibility with legacy or older routers that do not support BFD in hardware.

With this enhancement, BFD operates in the control plane, where software handles the transmission (TX), reception (RX), and state machine logic. This approach facilitates faster failure detection in environments using EVPN-VXLAN overlays with BGP routing over IRB interfaces.

This feature is ideal for networks with critical interoperability and control-plane flexibility, especially when upgrading or migrating from older hardware platforms.

Feature Characteristics

The Single-hop BFD with BGP over EVPN-VXLAN IRB feature includes the following characteristics:

- **Software-based BFD operation**: Introduces the CLI command `bfd session software`, enabling BFD to run in the control plane.
- **Dedicated session state**: Each BFD session maintains its own local discriminator, session state, and control-plane state machine.
- **Interoperability**: Supports BFD over IRB interfaces for routers that do not have hardware BFD capability.
- **EVPN-VXLAN integration**: Compatible with BGP routing over EVPN-VXLAN IRB deployments, improving failure detection across virtualized networks.
- **No reliance on hardware**: The software control plane manages the entire BFD session lifecycle, including packet processing.

Benefits

Implementing this feature provides the following advantages:

- Enhanced compatibility: Enables BFD deployment in networks with legacy routers that lack hardware BFD support.
- Improved resiliency: Offers faster failure detection and improved convergence in BGP EVPN-VXLAN environments.
- Flexible deployment: Facilitates BFD sessions over IRB interfaces without requiring hardware upgrades or dependencies.
- Simplified management: Reduces configuration complexity by allowing consistent BFD behavior across both new and existing platforms.

Limitations

- CPU-dependent scaling: The number of concurrent BFD sessions is limited by the CPU performance of the device, as all processing occurs in the control plane.

- Recommended interval: It is recommended to configure a minimum BFD interval of 250 ms to avoid excessive load on the control plane and ensure stable operation..

Configuration

Perform configuration steps for enabling Single-hop BFD with BGP over EVPN-VXLAN IRB, based on the configuration you provided. This includes enabling software-based BFD sessions over the IRB interface and configuring BFD integration with BGP in a VXLAN-EVPN environment.

Topology

The network topology includes various network elements such as routers, Customer Edge (CE) devices, Service Aggregator (SA) devices, and Provider Edge (PE) routers. The feature enables OSPF on the IRB interfaces, allowing for efficient routing and communication between network devices within the topology.

Single Hop BFD over EVPN-VXLAN IRB

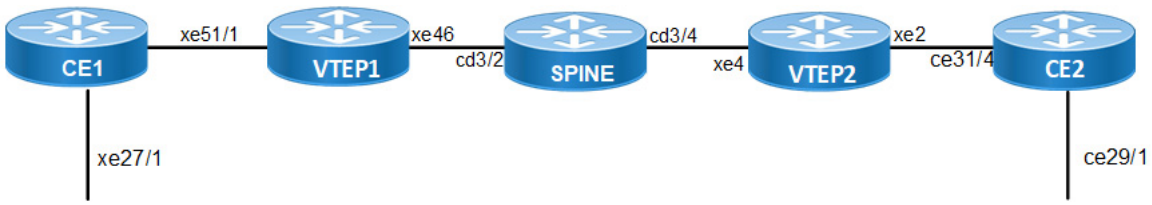

Perform the following configurations to enable single-hop BFD with BGP over EVPN-VXLAN IRB.

1. Enable the BFD filter profile and set default timers globally to support BFD across the device.

```
VTEP1(config)# hardware-profile filter bfd-group enable
VTEP1(config)# bfd interval 50 minrx 50 multiplier 50
```

2. Configure the IRB interface to participate in L3 forwarding within a VRF and enable software-based BFD for failure detection.

```
VTEP1(config)#interface irb1601
VTEP1(config-irb-if)# ip vrf forwarding L3_VRF1
VTEP1(config-irb-if)# evpn irb-if-forwarding anycast-gateway-mac
VTEP1(config-irb-if)# ip address 12.0.0.12/24 anycast
VTEP1(config-irb-if)# ip address 71.10.0.2/24 secondary
VTEP1(config-irb-if)# ipv6 address 1100::1/48
VTEP1(config-irb-if)# ipv6 address 2a02:2488:2:2906::2/64 anycast
VTEP1(config-irb-if)# bfd interval 50 minrx 50 multiplier 50
VTEP1(config-if)# bfd session software
VTEP1(config-if)# exit
```

3. Define the VRF and associate it with a Layer 3 VNI to enable inter-subnet routing in the overlay network..

```
VTEP1(config)#ip vrf L3_VRF1
VTEP1(config-vrf)# rd 1.1.1.1:10
VTEP1(config-vrf)# route-target both 10:10
VTEP1(config-vrf)# l3vni 1000
VTEP1(config-vrf)# exit
```

4. Enable OSPF for both global and VRF instances to exchange routing information and support BFD for neighbor monitoring:

```
VTEP1(config)# router ospf 1
VTEP1(config-router)# ospf router-id 1.1.1.1
VTEP1(config-router)# bfd all-interfaces
VTEP1(config-router)# redistribute connected
VTEP1(config-router)# network 10.1.31.0/24 area 0.0.0.0
VTEP1(config-router)# exit
VTEP1(config)# router ospf 2 L3_VRF1
VTEP1(config-router)# ospf router-id 1.1.1.1
VTEP1(config-router)# bfd all-interfaces
VTEP1(config-router)# redistribute bgp
VTEP1(config-router)# network 71.10.0.0/24 area 0.0.0.0
VTEP1(config-router)# exit
```

5. Enable IS-IS routing for both underlay and VRF contexts and enable BFD for fast convergence.

```
VTEP1(config)# router isis 1
VTEP1(config-router)# is-type level-1-2
VTEP1(config-router)# metric-style wide
VTEP1(config-router)# dynamic-hostname
VTEP1(config-router)# bfd all-interfaces
VTEP1(config-router)# net 49.0000.0000.0001.00
VTEP1(config-router)# redistribute connected
VTEP1(config-router)# exit

VTEP1(config)# router isis 2 L3_VRF3
VTEP1(config-router)# is-type level-1-2
VTEP1(config-router)# metric-style wide
VTEP1(config-router)# dynamic-hostname
VTEP1(config-router)# bfd all-interfaces
VTEP1(config-router)# net 49.0000.0000.0465.00
VTEP1(config-router)# redistribute bgp
VTEP1(config-router)# exit
```

6. Bind interfaces to IS-IS.

```
VTEP1(config)# interface lo
VTEP1(config-if)# ip router isis 1
VTEP1(config-if)# exit

VTEP1(config)# interface xe46
VTEP1(config-if)# ip router isis 1
VTEP1(config-if)# exit

VTEP1(config)# interface irb1603
VTEP1(config-irb-if)# ip router isis 2
VTEP1(config-irb-if)# exit
```

7. Enable BGP for EVPN and VRF instances, configure peer relationships, and apply BFD to detect failures quickly.

```
VTEP1(config)# router bgp 100
VTEP1(config-router)# bgp router-id 1.1.1.1
VTEP1(config-router)# no bgp inbound-route-filter
VTEP1(config-router)# neighbor 3.3.3.3 remote-as 100
VTEP1(config-router)# neighbor 3.3.3.3 update-source lo
VTEP1(config-router)# neighbor 3.3.3.3 advertisement-interval 0

VTEP1(config-router)# address-family l2vpn evpn
```

```
VTEP1(config-router-af)# neighbor 3.3.3.3 activate
VTEP1(config-router-af)# exit-address-family
```

8. Configure BGP for each VRF:

   L3_VRF1

```
VTEP1(config-router)# address-family ipv4 vrf L3_VRF1
VTEP1(config-router-af)# network 71.10.0.0/24
VTEP1(config-router-af)# redistribute connected
VTEP1(config-router-af)# redistribute ospf
VTEP1(config-router-af)# exit-address-family
```

   L3_VRF2

```
VTEP1(config-router)# address-family ipv4 vrf L3_VRF2
VTEP1(config-router-af)# redistribute connected
VTEP1(config-router-af)# neighbor 91.10.0.1 remote-as 100
VTEP1(config-router-af)# neighbor 91.10.0.1 activate
VTEP1(config-router-af)# neighbor 91.10.0.1 update-source 91.10.0.2
VTEP1(config-router-af)# neighbor 91.10.0.1 fall-over bfd
VTEP1(config-router-af)# exit-address-family
```

   L3_VRF3

```
VTEP1(config-router)# address-family ipv4 vrf L3_VRF3
VTEP1(config-router-af)# network 61.10.0.0/24
VTEP1(config-router-af)# redistribute connected
VTEP1(config-router-af)# redistribute isis
VTEP1(config-router-af)# exit-address-family
```

9. Enable VXLAN tunnel services and configure EVPN for Layer 2 and IRB services with anycast gateway MAC.

```
VTEP1(config)# nvo vxlan enable
VTEP1(config)# nvo vxlan irb
VTEP1(config)# evpn vxlan multihoming enable
VTEP1(config)# evpn irb-forwarding anycast-gateway-mac 0000.0000.1111
VTEP1(config)# nvo vxlan vtep-ip-global 1.1.1.1
```

10. Configure L2VNI for each VRF:

    L2 VRF1

```
VTEP1(config)# nvo vxlan id 1601 ingress-replication inner-vid-disabled
VTEP1(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2_VRF1
VTEP1(config-nvo)# evpn irb1601
VTEP1(config-nvo)# vni-name VNI-1601
VTEP1(config-nvo)# exit
```

    L2 VRF2

```
VTEP1(config)# nvo vxlan id 1602 ingress-replication inner-vid-disabled
VTEP1(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2_VRF2
VTEP1(config-nvo)# evpn irb1602
VTEP1(config-nvo)# vni-name VNI-1602
VTEP1(config-nvo)# exit
```

    L2 VRF3

```
VTEP1(config)# nvo vxlan id 1603 ingress-replication inner-vid-disabled
VTEP1(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2_VRF3
VTEP1(config-nvo)# evpn irb1603
VTEP1(config-nvo)# vni-name VNI-1603
VTEP1(config-nvo)# exit
```

11. Assign each IRB to its corresponding VRF and configure IP addresses and BFD.

IRB1602

```
VTEP1(config)# interface irb1602
VTEP1(config-irb-if)# ip vrf forwarding L3_VRF2
VTEP1(config-irb-if)# ip address 11.0.11.11/24
VTEP1(config-irb-if)# ip address 91.10.0.2/24 secondary
VTEP1(config-irb-if)# ipv6 address 1200::1/48
VTEP1(config-irb-if)# bfd session software
VTEP1(config-irb-if)# exit
```

IRB1603

```
VTEP1(config)# interface irb1603
VTEP1(config-irb-if)# ip vrf forwarding L3_VRF3
VTEP1(config-irb-if)# evpn irb-if-forwarding anycast-gateway-mac
VTEP1(config-irb-if)# ip address 13.0.0.13/24 anycast
VTEP1(config-irb-if)# ip address 61.10.0.2/24 secondary
VTEP1(config-irb-if)# ipv6 address 1300::1/48
VTEP1(config-irb-if)# bfd session software
VTEP1(config-irb-if)# exit
```

12. Attach underlay interfaces for IS-IS, configure port-channel for multihoming, and set interface parameters.

```
VTEP1(config)# interface xe46
VTEP1(config-if)# load-interval 30
VTEP1(config-if)# ip address 10.1.31.1/24
VTEP1(config-if)# mtu 9000
VTEP1(config-if)# ip router isis 1
VTEP1(config-if)# exit

VTEP1(config)# interface xe51/1
VTEP1(config-if)# channel-group 1 mode active
VTEP1(config-if)# exit

VTEP1(config)# interface po1
VTEP1(config-if)# switchport
VTEP1(config-if)# load-interval 30
VTEP1(config-if)# evpn multi-homed system-mac 0000.aaaa.bbbb
VTEP1(config-if)# exit
```

Configuration Snapshot

```
!
vlan database
 vlan-reservation 4039-4094
 vlan 100 bridge 1
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn vxlan multihoming enable
!
mac vrf m1
 rd 2.2.2.2:1
 route-target both 10:10
```

```
!
ip vrf vxlan_l3
 rd 2222:300
 route-target both 300:300
 l3vni 5053
!
evpn irb-forwarding anycast-gateway-mac 0000.0000.1111
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 100 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp m1
 evpn irb200
!
interface po10
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.1122.3344
!
interface ce10
 ip address 46.1.1.2/24
 mtu 7000
!
interface ce15
 channel-group 10 mode active
!
interface ce17
 description leaf2-leaf3
 ip address 20.0.0.1/24
 mtu 7000
!
interface irb200
 ip vrf forwarding vxlan_l3
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 50.1.1.1/24 anycast
 ip address 200.2.2.2/24 secondary
 mtu 9216
 bfd session software
!
interface lo
 ip address 127.0.0.1/8
 ip address 2.2.2.2/32 secondary
 ipv6 address ::1/128
!
 exit
!
router ospf 1
 ospf router-id 2.2.2.2
 network 2.2.2.2/32 area 0.0.0.0
 network 20.0.0.0/24 area 0.0.0.0
```

```
 network 46.1.1.0/24 area 0.0.0.0
!
router bgp 1
 bgp router-id 2.2.2.2
 neighbor 1.1.1.1 remote-as 1
 neighbor 3.3.3.3 remote-as 1
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 !
 address-family ipv4 unicast
 redistribute connected
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 3.3.3.3 activate
 exit-address-family
 !
 address-family ipv4 vrf vxlan_l3
 max-paths ibgp 2
 redistribute connected
 neighbor 200.2.2.10 remote-as 1
 neighbor 200.2.2.10 activate
 neighbor 200.2.2.10 update-source 200.2.2.2
 neighbor 200.2.2.10 fall-over bfd
 exit-address-family
 !
 exit
!
line console 0
 exec-timeout 0
!
nvo vxlan access-if port-vlan po10 100
 map vnid 100
!
!
end
```

Implementation Examples

A data center is using EVPN-VXLAN to provide Layer 2 and Layer 3 segmentation across a spine-leaf fabric. IRB interfaces are configured on leaf switches to enable routed communication between VLANs. BGP is used as the control protocol for EVPN route exchange.

Enable software-based Single-hop BFD using the CLI `bfd session software` under the IRB interface or BGP neighbor configuration. This allows BFD to run entirely in the control plane.

New CLI Commands

`bfd session software` parameter is added under bfd session command. For more info refer to the *Bidirectional Forwarding Commands section* in the *Layer 3 Guide*.

Validation

```
show bfd session vrf all
BFD process for VRF: (DEFAULT VRF)
=========================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface      Down-Reason  Remote-Addr
Number of Sessions:    0
BFD process for VRF: vxlan_l3
=========================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface      Down-Reason  Remote-Addr
2049       2060         IPv4         Single-Hop  Up          00:23:16  irb200         NA           200.2.2.10/32
Number of Sessions:    1
```

Glossary

The following provides definitions for key terms used throughout this document.

Single Home VxLAN	This refers to a Virtual Extensible LAN (VxLAN) deployment where a single data center or network site is connected to a single external network (usually the internet) for connectivity.
IRB	A networking feature that enables the integration of Layer 3 IP routing and Layer 2 MAC address bridging within the same interface, simplifying network management and resource utilization.
OSPF	A dynamic and efficient link-state routing protocol used to determine the best path for data packets in an IP network. It is characterized by rapid convergence and adaptability, making it suitable for large and dynamic networks.
ISIS	A routing protocol designed for scalability and stability in computer networks, commonly used in large Service Provider networks. It provides a robust framework for routing information exchange.
Layer 3 Routing	Network routing operations at the Network Layer (Layer 3) of the OSI model, focusing on routing IP packets between different subnets or networks.
Layer 2 Bridging	Network bridging operations at the Data Link Layer (Layer 2) of the OSI model, handling the forwarding of data frames based on MAC addresses within the same network segment.
EVPN	Ethernet VPN, a technology that provides advanced and efficient methods for Layer 2 and Layer 3 services in Ethernet networks, often used in data centers and service provider environments.

VxLAN - EVPN for Service Provider Network

CHAPTER 1 VXLAN EVPN EVC Configuration

This chapter shows how to configure VXLAN EVPN Ethernet Virtual Circuit (EVC) which embeds the functionality of EVPN-VXLAN access ports to allow EVC frames across VTEPs. With this configuration, customers in the same VLAN can communicate even they are placed across distributed data centers.

Overview

An EVC represents a logical relationship between Ethernet User Network Interface (UNI) in a provider-based Ethernet service. An EVC represents the service offered and is carried through the provider network. Each EVC is configured by a unique name across the provider network.

An EVC is an end–to–end representation of a single instance of a Layer 2 service that a service provider offers. An EVC embodies the different parameters based on which the service is offered. EVC prevents data transfer between sites that are not part of the same EVC.

EVC is an A–Z circuit that enables you to pass customer VLANs from one port on a node to another port on another node in the network. EVC represents a Carrier Ethernet service and is an entity that provides end–to–end connection between two or more customer end points.

Topology

Figure 1-6: VXLAN EVPN EVC

RTR1/VTEP1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode for loopback.
`(config-if)#ip address 33.33.33.0/31 secondary`	Assign secondary IP address.
`(config-if)#exit`	Exit interface mode.
`(config)#mac vrf vrf1`	Create mac routing/forwarding instance with vrf1 name and enter into VRF mode
`(config-vrf)# rd 100:11`	Assign RD value

(config-vrf)# route-target export 200:11	Assign route-target value for export
(config-vrf)# route-target import 400:11	Assign route-target value for import
(config-vrf)#exit	Exit VRF configuration mode
(config)#interface xe37	Enter interface mode for xe37
(config-if)#ip address 11.11.11.0/31	Assign IP address in /31 mask.
(config-if)#exit	Exit interface mode.
(config)#interface xe15	Enter interface mode for xe15
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit interface mode.
(config)#router bgp 100	Enter BGP router mode
(config-router)# bgp router-id 1.1.1.1	Assign BGP router ID
(config-router)#neighbor 11.11.11.1 remote-as 200	Specify a neighbor router with peer IP address and remote-as defined
(config-router)#neighbor 11.11.11.1 fall-over bfd	Configure single-hop BFD session for its BGP peer
(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family
(config-router-af)#network 33.33.33.0/31	Advertise loopback network into BGP for VTEP ID reachability
(config-router-af)#exit-address-family	Exit ipv4 unicast address family mode
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 11.11.11.1 activate	Activate the peer into address family mode
(config-router-af)#exit-address-family	Exit l2vpn address family mode
(config-router)#exit	Exit BGP router mode
(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan vtep-ip-global 33.33.33.0	Configure Source vtep-ip-global configuration
(config)#nvo vxlan id 1 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1	Assign VRF for evpn-bgp to carry EVPN route
(config-nvo)#exit	Exit VXLAN tenant mode.
(config)#nvo vxlan access-if port-vlan xe15 1000 inner-vlan 2000	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) and inner-vlan (CVLAN) mapping
(config-nvo-acc-if)#map vnid 1	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)#exit	Exit VXLAN access-interface mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configuration mode

RTR2

#configure terminal	Enter configure mode.
(config)#interface xe13/3	Enter interface mode for xe13/3
(config-if)#ip address 11.11.11.1/31	Assign IP address in /31 mask.

(config-if)#exit	Exit interface mode.
(config)#interface xe13/1	Enter interface mode for xe13/1
(config-if)#ip address 12.12.12.1/31	Assign IP address in /31 mask.
(config-if)#exit	Exit interface mode
(config)#router bgp 200	Enter BGP router mode
(config-router)# bgp router-id 2.2.2.2	Assign BGP router ID
(config-router)#neighbor 11.11.11.0 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 11.11.11.0 fall-over bfd	Configure single-hop BFD session for its BGP peer
(config-router)#neighbor 12.12.12.0 remote-as 300	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 12.12.12.0 fall-over bfd	Configure single-hop BFD session for its BGP peer
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 11.11.11.0 activate	Activate the peer into address family mode
(config-router-af)#neighbor 12.12.12.0 activate	Activate the peer into address family mode
(config-router-af)#exit-address-family	Exit l2vpn address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration
(config-router)#exit	Exit BGP router mode

RTR3/VTEP2

#configure terminal	Enter configure mode.
(config)#interface lo	Enter interface mode for loopback.
(config-if)#ip address 34.34.34.0/31 secondary	Assign secondary IP address.
(config-if)#exit	Exit interface mode.
(config)#mac vrf vrf1	Create mac routing/forwarding instance with vrf1 name and enter into vrf mode
(config-vrf)#rd 300:11	Assign RD value
(config-vrf)#route-target export 400:11	Assign route-target value for export
(config-vrf)#route-target import 200:11	Assign route-target value for import
(config-vrf)#exit	Exit vrf configuration mode
(config)#interface xe13	Enter interface mode for xe13
(config-if)#ip address 12.12.12.0/31	Assign IP address in /31 mask.
(config-if)#exit	Exit interface mode.
(config)#interface xe6	Enter interface mode for xe6
(config-if)#switchport	Make it L2 interface
(config-if)#exit	Exit interface mode.
(config)#router bgp 300	Enter BGP router mode
(config-router)# bgp router-id 3.3.3.3	Assign BGP router ID

(config-router)#neighbor 12.12.12.1 remote-as 200	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 12.12.12.1 fall-over bfd	Configure single-hop BFD session for its BGP peer
(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family
(config-router-af)#network 34.34.34.0/31	Advertise loopback network into BGP for VTEP ID reachability
(config-router-af)#exit-address-family	Exit ipv4 unicast address family mode
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 12.12.12.1 activate	Activate the peer into address family mode
(config-router-af)#exit-address-family	Exit l2vpn address family mode
(config-router)#exit	Exit BGP router mode
(config)#nvo vxlan enable	Enable VXLAN
(config)#nvo vxlan vtep-ip-global 34.34.34.0	Configure Source vtep-ip-global configuration
(config)#nvo vxlan id 1 ingress-replication inner-vid-disabled	Configure VXLAN Network identifier with/without inner-vid-disabled configure and enter into VXLAN tenant mode
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1	Assign VRF for evpn-bgp to carry EVPN route
(config-nvo)#exit	Exit VXLAN tenant mode.
(config)#nvo vxlan access-if port-vlan xe6 1000 inner-vlan 2000	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) and inner-vlan (CVLAN) mapping
(config-nvo-acc-if)#map vnid 1	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)#exit	Exit VXLAN access-interface mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configuration mode

Validation

Retaining SVLAN and CVLAN tags across Data Centers

RTR1/VTEP1

```
VTEP1#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 33.33.33.0
!
nvo vxlan id 1 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrf1
!
nvo vxlan access-if port-vlan xe1/1 1000 inner-vlan 2000
 map vnid 1
!
!
```

```
VTEP1#show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 8
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Dow
n   State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
11.11.11.1               4   200   73         73       8      0      0  00:30:41
               2      0      1      1      0      0

Total number of neighbors 1

Total number of Established sessions 1

VTEP1#show bgp l2vpn evpn
BGP table version is 8, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path
 Peer          Encap

RD[100:11] VRF[vrf1]:
*    [2]:[0]:[1]:[48,0000.339a.9397]:[0]:[1]
                       34.34.34.0           0         100           0       200
300 i  11.11.11.1      VXLAN
*>   [2]:[0]:[1]:[48,0000.339a.9abb]:[0]:[1]
                       33.33.33.0           0         100           32768    i  -
---------      VXLAN
*>   [3]:[1]:[32,33.33.33.1]
33.33.33.0           0        100     32768
i  ----------      VXLAN


*    [3]:[1]:[32,34.34.34.0]
                       34.34.34.0           0         100           0       200
300 i  11.11.11.1      VXLAN

RD[300:11]
*>   [2]:[0]:[1]:[48,0000.339a.9397]:[0]:[1]
                       34.34.34.0           0         100           0       200
```

```
300 i  11.11.11.1      VXLAN
*>   [3]:[1]:[32,34.34.34.0]
                       34.34.34.0           0         100           0       200
300 i  11.11.11.1      VXLAN

Total number of prefixes 6
VTEP1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     11.11.11.0/31 is directly connected, xe10/1, 00:36:00
C     33.33.33.0/31 is directly connected, lo, 00:37:33
B     34.34.34.0/31 [20/0] via 11.11.11.1, xe10/1, 00:27:03
C     127.0.0.0/8 is directly connected, lo, 23:14:51

Gateway of last resort is not set
VTEP1


VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
           AC - Access Port
          (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface     ESI
VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________
_____________________________________________
1        ----         L2       NW   ----          ------
---- ----      33.33.33.0       34.34.34.0
1        ----         --       AC   xe1/1         --- Single Homed Port ---
1000 ----      ----              ----
 Total number of entries are 2


VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
====
33.33.33.0       34.34.34.0       Installed     00:26:27     00:26:27
```

```
Total number of entries are 1
VTEP1#show nvo vxlan mac-table

================================================================================
================================================================
                                                          VXLAN MAC Entries
================================================================================
================================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
       Type                Status             AccessPortDesc
________________________________________________________________________________
________________________________________________________________

1         xe1/1      1000   2000         0000.339a.9abb 33.33.33.0
       Dynamic Local        -------             -------
1         ----       ----   ----         0000.339a.9397 34.34.34.0
       Dynamic Remote       -------             -------

Total number of entries are : 2
VTEP1#
```

RTR3/VTEP2

```
#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 34.34.34.0
!
nvo vxlan id 1 ingress-replication  inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrf1
!
nvo vxlan access-if port-vlan xe6 1000 inner-vlan 2000
 map vnid 1
!
VTEP2#show bgp l2vpn evpn summary
BGP router identifier 3.3.3.3, local AS number 300
BGP table version is 7
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Dow
n   State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
12.12.12.1               4   200   63         64       7      0      0  00:26:54
               2      0      1      1      0      0

Total number of neighbors 1

Total number of Established sessions 1
```

```
VTEP2#show bgp l2vpn evpn
BGP table version is 7, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path
 Peer          Encap

RD[100:11]
*>   [2]:[0]:[1]:[48,0000.339a.9abb]:[0]:[1]
                       33.33.33.0           0         100           0       200
100 i  12.12.12.1      VXLAN
*>   [3]:[1]:[32,33.33.33.0]
                       33.33.33.0           0         100           0       200
100 i  12.12.12.1      VXLAN


*>   [2]:[0]:[1]:[48,0000.339a.9397]:[0]:[1]
                       34.34.34.0           0         100           32768    i  -
---------      VXLAN
*    [2]:[0]:[1]:[48,0000.339a.9abb]:[0]:[1]
                       33.33.33.0           0         100           0       200
100 i  12.12.12.1      VXLAN
*    [3]:[1]:[32,33.33.33.0]
                       33.33.33.0           0         100           0       200
100 i  12.12.12.1      VXLAN
*>   [3]:[1]:[32,34.34.34.0]
                       34.34.34.0           0         100           32768    i  -
---------      VXLAN

Total number of prefixes 6

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default
```

```
IP Route Table for VRF "default"
C     12.12.12.0/31 is directly connected, xe13, 00:28:41
B     33.33.33.0/31 [20/0] via 12.12.12.1, xe13, 00:26:56
C     34.34.34.0/31 is directly connected, lo, 00:29:36
C     127.0.0.0/8 is directly connected, lo, 00:52:46

Gateway of last resort is not set

VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface     ESI
VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________
_____________________________________________
1        ----         L2       NW   ----          ------
---- ----      34.34.34.0       33.33.33.0
1        ----         --       AC   xe1/1         --- Single Homed Port ---
1000 ----      ----              ----
 Total number of entries are 2


VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
=============================================================================
====
33.33.33.0       34.34.34.0       Installed     00:26:27     00:26:27
Total number of entries are 1


VTEP2#show nvo vxlan mac-table
=============================================================================
=============================================================
                                                        VXLAN MAC Entries
=============================================================================
=============================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
       Type                Status              AccessPortDesc
_____________________________________________________________________________
_____________________________________________________________

1         ----       ----   ----         0000.339a.9abb 33.33.33.0
       Dynamic Remote       -------             -------
1         xe6      1000   2000         0000.339a.9397 34.34.34.0
       Dynamic Local        -------             -------
```

```
Total number of entries are : 2
#
```

Popping SVLAN and CVLAN Tag

Use the previous configuration on VTEP1 and perform the configuration below on VTEP2.

VTEP2

(config)#nvo vxlan access-if port xe6	Enable port-only mapping for access port
(config-nvo-acc-if)#map vnid 1	Map VXLAN Identified to access-port for VXLAN
(config-nvo-acc-if)#exit	Exit VXLAN access-interface mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configuration mode

RTR3/VTEP2

```
#show nvo vxlan
VXLAN Information
=================
Codes: NW - Network Port
AC - Access Port
(u) - Untagged
VNID VNI-Name VNI-Type Type Interface ESI
VLAN DF-Status Src-Addr Dst-Addr
________________________________________________________________________________
1 ---- L2 NW ---- ------
---- ---- 34.34.34.0 33.33.33.0
1 ---- -- AC xe6 --- Single Homed Port ---
---- ---- ---- ----
Total number of entries are 2

VTEP2#show nvo vxlan mac-table
================================================================================
=====================================================================
                                                         VXLAN MAC Entries
================================================================================
=====================================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
       Type                Status              AccessPortDesc
________________________________________________________________________________
_____________________________________________________________________

1         ----       ----   ----         0000.339a.9abb 33.33.33.0
       Dynamic Remote      -------             -------

Total number of entries are : 1
#
```

```
VTEP2#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 34.34.34.0
!
nvo vxlan id 1 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrf1
!
nvo vxlan access-if port-vlan xe1/1 3000 inner-vlan 2000
 map vnid 1

VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source            Destination      Status          Up/Down         Update
================================================================================
====
33.33.33.0       34.34.34.0       Installed     00:26:27     00:26:27
Total number of entries are 1
```

Popping and Later Pushing SVLAN Tag

Use the previous configuration on VTEP1 and perform the configuration below on VTEP2.

`(config)#nvo vxlan access-if port-vlan xe6 3000 inner-vlan 2000`	Enable port-vlan mapping i.e. access port to outer-vlan (SVLAN) and inner-vlan (CVLAN) mapping
`(config-nvo-acc-if)#map vnid 1`	Map VXLAN Identified to access-port for VXLAN
`(config-nvo-acc-if)#exit`	Exit VXLAN access-interface mode
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configuration mode

RTR3/VTEP2

```
#show running-config nvo vxlan
!
nvo vxlan enable
!
nvo vxlan vtep-ip-global 34.34.34.0
!
nvo vxlan id 1 ingress-replication  inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrf1
!
nvo vxlan access-if port-vlan xe1/1 3000 inner-vlan 2000
 map vnid 1
!
#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
```

```
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface     ESI
VLAN DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________
1        ----         L2       NW   ----          ------
---- ----      34.34.34.0       33.33.33.0
1        ----         --       AC   xe6         --- Single Homed Port ---
3000 ----      ----              ----

Total number of entries are 2

#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
=============================================================================
34.34.34.0       33.33.33.0       Installed     00:06:48     00:06:48
Total number of entries are 1

VTEP2#show nvo vxlan mac-table
=============================================================================
=================================================================
                              VXLAN MAC Entries
=============================================================================
=================================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
       Type                  Status              AccessPortDesc
_____________________________________________________________________________
_________________________________________________________________

1         ----       ----   ----         0000.339a.9abb 33.33.33.0
       Dynamic Remote       -------             -------

Total number of entries are : 1

#
```

CHAPTER 2 EVPN VXLAN E-Tree

Overview

Ethernet VPN Ethernet-Tree (EVPN E-Tree), is a networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, especially in scenarios where specific services or devices need controlled communication. It categorizes network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them.

Feature Characteristics

Implemented Scenario 1 of the EVPN E-Tree solution, as defined by RFC-8317, designates each Provider Edge (PE) node as either a Leaf or a Root site per Virtual Private Network (VPN) for VXLAN EVPN in OcNOS.

Scenario 1: Leaf or Root Site(s) per PE

Scenario 1 involves a topology with three PE nodes: PE-1, PE-2, and PE-3. PE-1 and PE-2 are Multi-Homed nodes (MH-1 and MH-2), with PE-3 acting as the Root node. PE-1 and PE-2 function as Leaf nodes and are part of a single home access interface (SH1 and SH2).

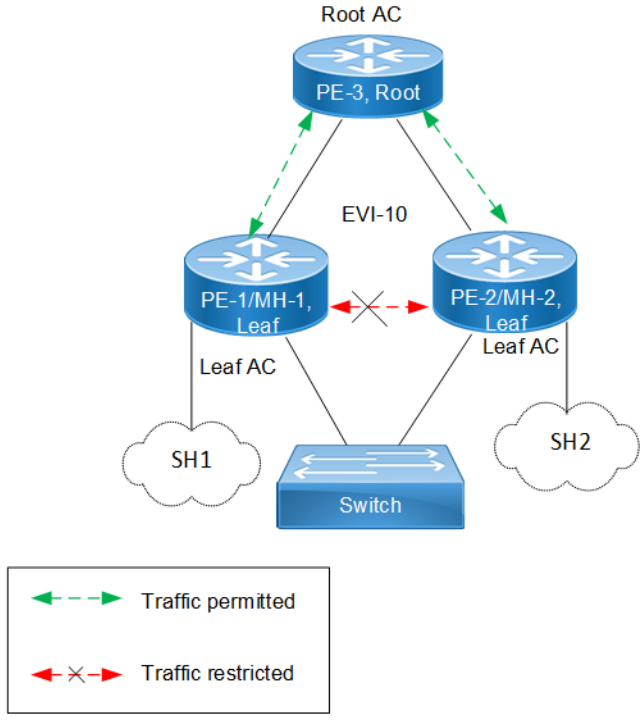

EVPN E-Tree

The classification ensures that communication follows specific rules:

- Communication between Leaf hosts is restricted, as indicated by red dotted lines with a cross mark (X) in the topology diagram. However, communication between Leaf and Root nodes, as well as between Root nodes, is permitted, marked by green dotted lines.

- Leaf nodes within PE-1 and PE-2 are isolated from each other, preventing intra-PE communication.

The scenario 1 is achieved through two main concepts:

1. **Inter-PE Communication**
   - The inter-PE Route Target (RT) Constraint Method is applicable only to Single-Homing (SH) devices. Two RTs per broadcast domain are utilized, with Leaf PEs exporting Leaf RTs and Root nodes exporting Root RTs. Leaf nodes import only Root RTs, allowing communication with Root PEs while preventing communication with other Leaf nodes. RT constraints limit the import of specific EVPN routes (MAC-IP and IMET routes) to designated paths for inter-PE communication.
   - IPI employs a proprietary method to support inter-PE connectivity for both SH and MH devices, using BGP extended community to advertise Leaf Indication in BGP routes and influence traffic flow for both Unicast and BUM traffic. This method enables implementation of ARP or ND cache suppression and MAC mobility sub-features specified in RFC-7432.
2. **Intra-PE communication:** Local Split Horizon controls intra-PE communication between Attachment Circuits (ACs) within Leaf PE nodes, ensuring that traffic between ACs does not egress to other Leaf ACs.

   Note: This functionality depends on hardware capabilities.

Benefits

EVPN E-Tree offers benefits in networking environments by providing efficient traffic control, enhanced security, scalability, and improved performance.

Efficient Traffic Control: EVPN E-Tree allows for efficient control over traffic within network broadcast domains. By segregating nodes into Leaf and Root categories, it enables precise management of communication flows, ensuring the traffic is directed only where needed.

Enhanced Security: The isolation of Leaf hosts from each other adds a layer of security to the network. This prevents unauthorized communication between devices within the same broadcast domain, reducing the risk of data breaches and unauthorized access.

Scalability: EVPN E-Tree is scalable, making it suitable for networks of various sizes and complexities. Whether deploying in small-scale environments or large enterprise networks, EVPN E-Tree offers flexibility and scalability to meet evolving business needs.

Improved Performance: By controlling communication paths and optimizing traffic flows, EVPN E-Tree can improve network performance. This ensures that critical data packets are delivered efficiently, reducing latency and enhancing overall network performance.

Prerequisites

In setting up a VXLAN EVPN network, certain prerequisites are essential to ensure proper functionality and connectivity.

Ensure VXLAN EVPN Configuration: Confirm that VXLAN, EVPN VXLAN, and VXLAN filtering are already enabled in the network as they are required for VXLAN EVPN Multihoming.

Define Interfaces and Loopback Addresses: Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific system MAC addresses (Ethernet Segment Identifier (ESI) values) for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

Configure OSPF and BGP for Dynamic Routing: Enable OSPF to facilitate dynamic routing within the network. Define OSPF router IDs to match loopback IP addresses and add network segments to OSPF areas for proper route

distribution. Additionally, establish BGP sessions to advertise routes between different nodes. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

Leaf Node

1. Enable VXLAN and EVPN MH

   Enable features like VXLAN and EVPN Multihoming, VXLAN filtering, and quality of service (QoS) capabilities on all Leaf nodes.

   ```
   !
   nvo vxlan enable
   !
   evpn vxlan multihoming enable
   !
   qos enable
   !
   ```

2. Configure Interfaces and Loopback

   Define a port channel interface (`po1`) as an L2 interface and assign the system MAC (`0000.0000.1111`) as the ESI value. Designate an interface (`xe7`) as a member port of `po1`. Assign the loopback IP address (`1.1.1.1`) to Leaf node, and set IP addresses (`10.10.10.1` and `10.10.11.1`) to interfaces (`xe45` and `xe49/2`), respectively, for connectivity with Spine nodes.

   ```
   !
   interface po1
    switchport
    evpn multi-homed system-mac 0000.0000.1111
   !
   interface lo
    ip address 1.1.1.1/32 secondary
   !
   interface xe7
    channel-group 1 mode active
   !
   interface xe45
    ip address 10.10.10.1/24
   !
   interface xe49/2
    ip address 10.10.11.1/24
    exit
   !
   ```

3. Configure OSPF

   In OSPF router mode, set the router ID (`1.1.1.1`), to match the loopback IP address. Add the loopback network (`1.1.1.1/32`) and networks (`10.10.10.0/24 and 10.10.11.0/24`) connected to Spine nodes in OSPF area `0`. Enable Bidirectional Forwarding Detection (BFD) on all OSPF interfaces for faster convergence.

   ```
   !
   router ospf 100
    ospf router-id 1.1.1.1
    bfd all-interfaces
    network 1.1.1.1/32 area 0.0.0.0
    network 10.10.10.0/24 area 0.0.0.0
    network 10.10.11.0/24 area 0.0.0.0
   !
   ```

4. Configure BGP

   In BGP router mode, set the router ID (`1.1.1.1`) to match the loopback IP address. Specify the loopback IP address of each Leaf node as neighbors with their respective remote AS numbers. Configure the loopback as the update source for each neighbor and set the advertisement interval (`0`) for rapid convergence. In L2VPN EVPN address family mode, activate each Leaf node (`2.2.2.2, 3.3.3.3, 4.4.4.4`) to establish connections within the EVPN address family.

```
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 2.2.2.2 remote-as 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 2.2.2.2 activate
 neighbor 3.3.3.3 activate
 neighbor 4.4.4.4 activate
 exit-address-family
 !
 exit
!
```

5. Configure VRF

   In VRF mode, create a MAC routing or forwarding instance (`VRF1`). Assign the Route Distinguisher (RD) value (`1.1.1.1:100`) and set both import and export route-target value (`100:100`) . Ensure that the same route-target value is configured on all Leaf nodes for MAC VRF to maintain consistency.

```
!
mac vrf VRF1
 rd 1.1.1.1:100
 route-target both 100:100
!
```

Spine Node

1. Configure Interfaces and Loopback

   Enable QoS and assign specific IP addresses to loopback interfaces. Configure IP addresses for interfaces connected to each Leaf node.

```
!
qos enable
!
interface ce1/2
 ip address 40.40.40.2/24
!
interface ce1/4
 ip address 10.10.10.2/24
!
interface ce24/1
 ip address 30.30.30.2/24
```

```
!
interface ce27/1
 ip address 20.20.20.2/24
!
interface lo
 ip address 5.5.5.5/32 secondary
!
```

2. Configure OSPF

   In OSPF router mode, set the router ID (`5.5.5.5`), to match the loopback IP address. Add the loopback network (`5.5.5.5/32`) and networks (`10.10.10.0/24, 20.20.20.0/24, 30.30.30.0/24, and 40.40.40.0/24`) connected to Leaf nodes in OSPF area `0`. Enable BFD on all OSPF interfaces for faster convergence.

```
!
router ospf 100
 ospf router-id 5.5.5.5
 bfd all-interfaces
 network 5.5.5.5/32 area 0.0.0.0
 network 10.10.10.0/24 area 0.0.0.0
 network 20.20.20.0/24 area 0.0.0.0
 network 30.30.30.0/24 area 0.0.0.0
 network 40.40.40.0/24 area 0.0.0.0
!
```

Configure Switch

Set up an IEEE VLAN bridge, enabling VLANs and associating them with bridge `1`. Configure interfaces (`xe57, po1, xe46, xe47`) to be part of bridge `1`, setting them as hybrid ports with VLAN (`1000`) allowed and egress-tagged enabled. Designate interfaces connected to Leaf nodes (`xe46 and xe47`) as member ports of `po1`.

```
!
bridge 1 protocol ieee vlan-bridge
!
vlan database
 vlan-reservation 4000-4094
 vlan 1000 bridge 1 state enable
!
interface po1
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe46
 channel-group 1 mode active
!
interface xe47
 channel-group 1 mode active
!
interface xe57
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
```

Configuration

Configure various nodes within the topology to set up a VXLAN EVPN E-Tree network.

Topology

The sample topology includes Leaf Nodes (VTEP1, VTEP2, VTEP3, and VTEP4), Spine Nodes (SPINE1 and SPINE2), and Switches (SWITCH1 and SWITCH2).

VTEP1 and VTEP2 belong to Multi-homed group 1 (MH1) with po1, while VTEP3 and VTEP4 are in Multi-homed group 2 (MH2) with po2. VTEP2 and VTEP4 connect to single home access ports SH1 and SH2, respectively. All VTEPs link to Spine nodes SPINE1 and SPINE2. SWITCH1 is multi-homed to VTEP1 and VTEP2, and SWITCH2 connects to VTEP3 and VTEP4.

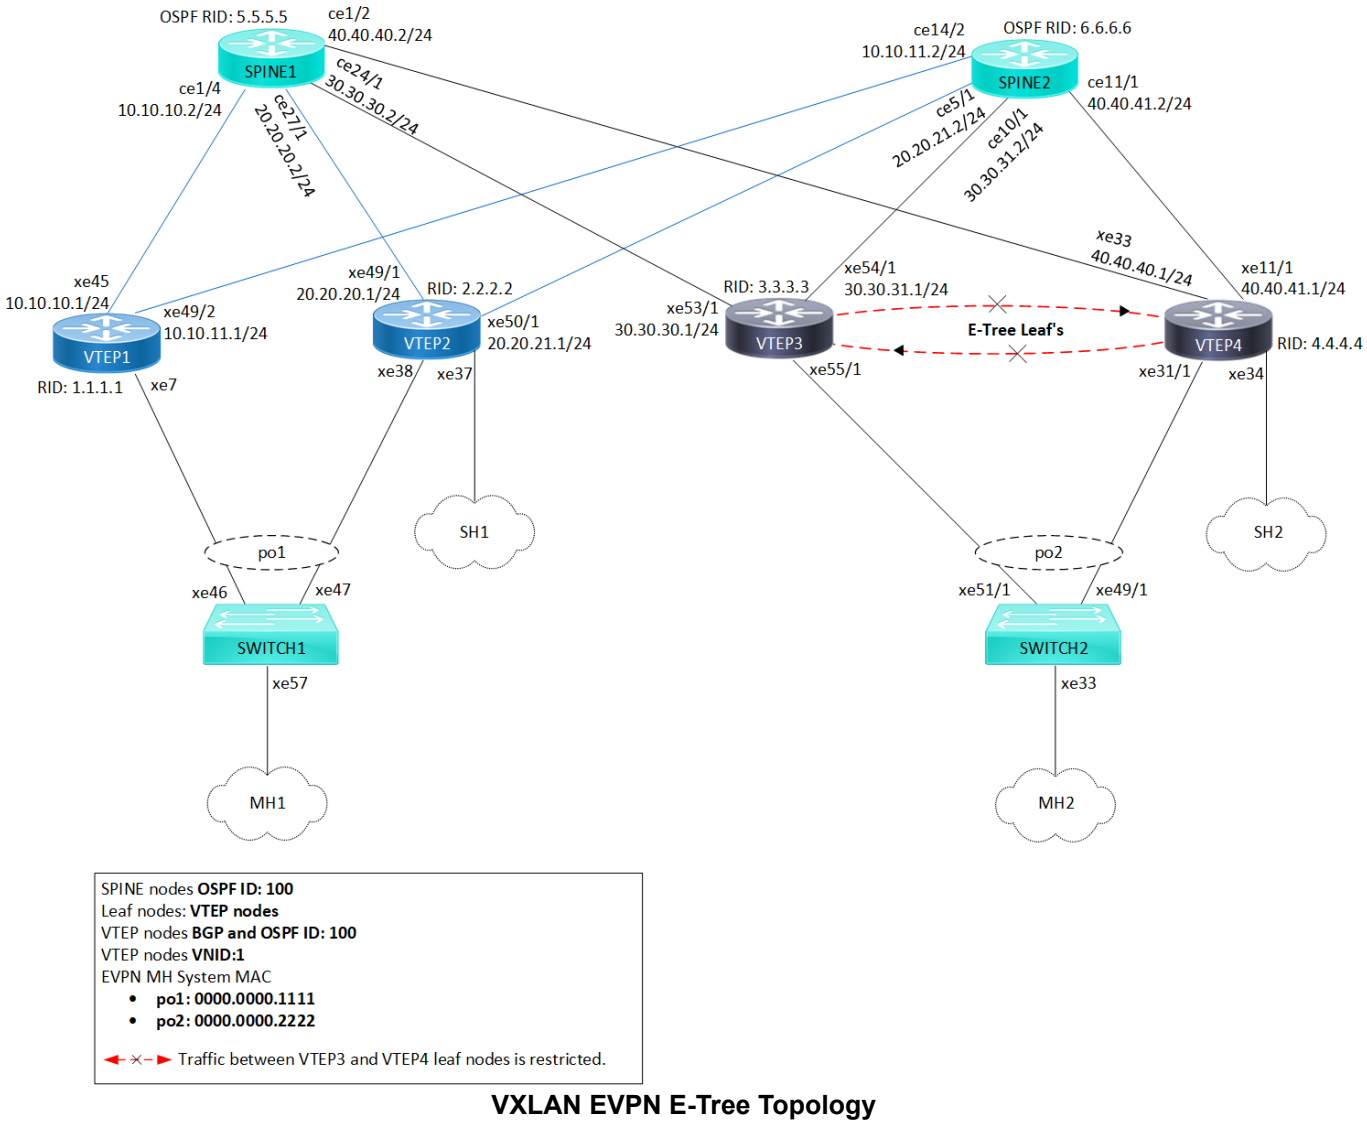

VXLAN EVPN E-Tree Topology

Note: Before configuring E-Tree, meet all Prerequisites for the following nodes:

- Leaf nodes: VTEP1, VTEP2, VTEP3, and VTEP4

- Spine nodes: SPINE1 and SPINE2
- Switches: SWITCH1 and SWITCH2

Enable EVPN E-Tree

The following E-Tree configurations applies to the VTEP nodes within the VXLAN network.

1. Enable EVPN E-Tree on VTEP3 and VTEP4 nodes, allowing them to participate in E-Tree functionality within the VXLAN network, controlling traffic and establishing hierarchical connections between Leaf nodes in the network architecture.

   ```
   (config)#evpn etree enable
   ```

2. Set the ESI hold time (`90 seconds`) on all VTEP nodes to allow the tunnel to establish during VXLAN initialization before bringing up the ESI. Configure the source VTEP IP address (`3.3.3.3`) which serves as the global identifier for VXLAN encapsulation and decapsulation within the network, facilitating proper communication and tunnel establishment.

   ```
   (config)#evpn esi hold-time 90
   (config)#nvo vxlan vtep-ip-global 3.3.3.3
   ```

3. Define VXLAN identifier (`10`) with ingress replication and disabled inner VLAN ID (VID) for **E-Tree leaf nodes** (VTEP3 and VTEP4) to support hierarchical connectivity and traffic control within the VXLAN network. This configuration allows for efficient replication of traffic at the ingress point and ensures that inner VLAN IDs are disabled, optimizing the functionality of E-Tree leaf nodes within the network architecture. On the VXLAN tenant node, assign VRF (`VRF1`) to EVPN-BGP for carrying EVPN routes within the VXLAN network.

   ```
   (config)#nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
   (config-nvo)#vxlan host-reachability- protocol evpn-bgp VRF1
   (config-nvo)#exit
   ```

4. Enable port-VLAN mapping (`po2`) with VLAN ID (`1000`) to facilitate multi-homed access on all VTEP nodes. Map VXLAN identifier (`10`) to the access port for VXLAN connectivity.

   ```
   (config)#interface po2.1000 switchport
   (config-if)#encapsulation dot1q 1000
   (config-if)#access-if-evpn
   (config-nvo-acc-if)#map vpn-id 10
   (config-nvo-acc-if)#exit
   (config)#commit
   ```

Validation

Use the show commands described in this section to verify the network for proper VXLAN EVPN E-Tree configuration.

Verify OSPF sessions between the VTEP nodes and the SPINEs within the VXLAN network using the `show ip ospf neighbor` command. This command displays OSPF neighbor details, including the state of the OSPF neighbor relationship. A `State` of `Full/DR` indicates a fully adjacent and operational state between the routers, confirming proper OSPF connectivity within the network.

```
VTEP1#show ip ospf  neighbor

Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID     Pri   State    Dead Time   Address      Interface           Instance ID
5.5.5.5           1   Full/DR  00:00:32    10.10.10.2      xe45                   0
6.6.6.6           1    Full/DR  00:00:30   10.10.11.2      xe49/2                 0
```

Verify the BGP session status between VTEPs, using the `show bgp l2vpn evpn summary` command output. The `Up/Down` field indicates the duration for which the BGP session has been up or down.

```
VTEP1#show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 9
1 BGP AS-PATH entries
0 BGP community entries

Neighbor    V    AS    MsgRcv     MsgSen  TblVer  InQ   OutQ Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
2.2.2.2     4   100    34         28       7      0      0   00:07:37  9                3      4      1      1          0
3.3.3.3     4   100    30         33       8      0      0   00:07:34  6                3      2      1      0          0
4.4.4.4     4   100    31         28       7      0      0   00:07:37  8                3      4      1      0          0

Total number of neighbors 3

Total number of Established sessions 3
```

To validate the BGP L2VPN output on VTEPs and check MAC-IP routes and ESI information, use the `show bgp l2vpn evpn` command output. This command verifies routes with status code `i (internal) and EVPN route types 2 and 4`, displaying detailed information for each VTEP nodes.

```
VTEP1#show bgp l2vpn evpn
BGP table version is 9, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[1.1.1.1:100] VRF[VRF1]:
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*> [2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1000]:[32,100.100.100.1]:[10]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*> [2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1001]:[128,1000::1][10]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[2]:[0]:[10]:[48,0000:2000:2000]:[32,200.200.200.1]:[10]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[2]:[0]:[10]:[48,0000:2000:2001]:[128,2000::1][10]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[2]:[0]:[10]:[48,0000:4000:4000]:[32,104.104.104.1]:[10]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[2]:[0]:[10]:[48,0000:4000:4001]:[128,1004::1][10]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*> [3]:[10]:[32,1.1.1.1]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i[3]:[10]:[32,2.2.2.2]
```

```
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[3]:[10]:[32,3.3.3.3]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i[3]:[10]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[1.1.1.1:64512] VRF[evpn-gvrf-1]:
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
*> [4]:[00:00:00:00:00:11:11:00:00:00]:[32,1.1.1.1]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i[4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:100]
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1000]:[32,100.100.100.1]:[10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1001]:[128,1000::1][10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[2]:[0]:[10]:[48,0000:2000:2000]:[32,200.200.200.1]:[10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[2]:[0]:[10]:[48,0000:2000:2001]:[128,2000::1][10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[3]:[10]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:64512]
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[3.3.3.3:100]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
*>i[3]:[10]:[32,3.3.3.3]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN

RD[3.3.3.3:64512]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN

RD[4.4.4.4:100]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[0]:[10]:[48,0000:4000:4000]:[32,104.104.104.1]:[10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[0]:[10]:[48,0000:4000:4001]:[128,1004::1][10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[3]:[10]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[4.4.4.4:64512]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

Total number of prefixes 42
```

Validate the LAG interfaces (`po1` and `po2`) are up for MH1 and MH2 by reviewing the `show etherchannel summary` output. Check the `Link` and `sync` fields, where link displays the port channel interface and ID number, and sync indicates whether MAC address synchronization is enabled to forward Layer 3 packets arriving on these interfaces.

```
VTEP1#show etherchannel summary
  Aggregator po1 100001
  Aggregator Type: Layer2
  Admin Key: 0001 - Oper Key 0001
     Link: xe7 (5005) sync: 1
```

Validate the status of NVO VXLAN on VTEPs by examining the output of the `show nvo vxlan` command. The `DF-Status` field displays the forwarding status of VXLAN tunnels as a Designated Forwarder (DF) or Non-Designated Forwarder (Non-DF).

```
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
10       ----         L2       NW   ----      ----                            ----      ----      1.1.1.1          4.4.4.4
10       ----         L2       NW   ----      ----                            ----      ----      1.1.1.1          3.3.3.3
10       ----         L2       NW   ----      ----                            ----      ----      1.1.1.1          2.2.2.2
10       ----         --       AC   po1       00:00:00:00:00:11:11:00:00:00 1000      DF        ----             ----

Total number of entries are 4

VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
10       ----         L2       NW   ----      ----                            ----      ----      2.2.2.2          4.4.4.4
10       ----         L2       NW   ----      ----                            ----      ----      2.2.2.2          1.1.1.1
10       ----         L2       NW   ----      ----                            ----      ----      2.2.2.2          3.3.3.3
10       ----         --       AC   xe37      --- Single Homed Port ---     1000      ----      ----             ----
10       ----         --       AC   po1       00:00:00:00:00:11:11:00:00:00 1000      NON-DF    ----             ----

Total number of entries are 5

VTEP3#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
10       ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          2.2.2.2
10       ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          1.1.1.1
10       ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4
10       ----         --       AC   po2       00:00:00:00:00:22:22:00:00:00 1000      DF        ----             ----

Total number of entries are 4

VTEP4#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
```

```
         AC - Access Port
        (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
10       ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          2.2.2.2
10       ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3
10       ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          1.1.1.1
10       ----         --       AC   xe34      --- Single Homed Port ---       1000      ----      ----             ----
10       ----         --       AC   po2       00:00:00:00:00:22:22:00:00:00   1000      NON-DF    ----             ----

Total number of entries are 5
```

Validate the NVO VXLAN tunnel status on VTEPs by reviewing the output of the `show nvo vxlan tunnel` command. The `Status` field indicates the current status of each tunnel. In this case, all three tunnels between VTEPs and their respective destinations are marked as `Installed`, confirming that these tunnels are successfully established and operating.

```
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status        Up/Down       Update
================================================================================
1.1.1.1           4.4.4.4          Installed     00:02:26      00:01:58
1.1.1.1           3.3.3.3          Installed     00:02:26      00:01:55
1.1.1.1           2.2.2.2          Installed     00:02:25      00:01:55

Total number of entries are 3
```

Validate the VXLAN access interface status on VTEPs by examining the output of the `show nvo vxlan access-if brief` command. The `up admin and link status` confirms that the access port associated with VXLAN is active and functioning properly on the VTEP nodes.

```
VTEP1#show nvo vxlan access-if brief

                      Inner                          Admin   Link
Interface  Vlan       vlan  Ifindex    Vnid         status  status
-------------------------------------------------------------------
po1        1000       ---   0x7a120    10           up      up

Total number of entries are 1
```

Static MAC-IP Advertisement

Configure static MAC-IP advertisement through SH and MH VTEPs from Root and Leaf nodes. Advertise static MAC addresses for IPv4 and IPv6 from MH1, MH2, SH1, and SH2 VTEPs. Ensure that VTEP1 and VTEP2 in MH1 have the same MAC addresses configured under the port-channel access port. Symmetrical configurations between MH VTEPs should be maintained.

Configure MH1 and MH2 VTEPs

Configure static MAC addresses for IPv4 (`100.100.100.1`) and IPv6 (`1000::1`) under the VXLAN MH access-port (`po1`) with VLAN ID (`1000`). Ensure that identical MAC addresses are set up within the MH1-VTEPs for advertisement. Apply similar configurations to MH2-VTEPs for static MAC-IP advertisement.

```
!
nvo vxlan access-if port-vlan po1 1000
 map vnid 10
 mac 0000.1000.1000 ip 100.100.100.1
 mac 0000.1000.1001 ipv6 1000::1
```

```
!
```

Configure SH1 and SH2 VTEPs

Configure static MAC addresses for IPv4 (`200.200.200.1`) and IPv6 (`2000::1`) under the VXLAN SH access-port (`xe37`) with VLAN ID (`1000`) on SH1 (VTEP2). This setup ensures that SH1 advertises these static MAC addresses over the specified VXLAN access-port. Repeat similar configurations for SH2 (VTEP4) using different static MAC addresses for both IPv4 and IPv6.

```
!
nvo vxlan access-if port-vlan xe37 1000
 map vnid 10
 mac 0000.2000.2000 ip 200.200.200.1
 mac 0000.2000.2001 ipv6 2000::1
!
```

Validation

Verify the MAC table entries on MH VTEPs (MH1 and MH2) and the SH VTEPs (VTEP2 and VTEP4). The MAC addresses are advertised using the ESI values from VTEP1 and VTEP2 for MH1, and from VTEP3 and VTEP4 for MH2. Additionally, verify the VTEP IP addresses associated with SH VTEP2 and VTEP4 for MAC advertisement.

In the output of the `show nvo vxlan mac-table` command on all VTEP nodes, the MAC entries advertised from Leaf VTEPs will have the `LeafFlag` field status `set`.

Note:

- MAC IPv4 or IPv6 configured under SH Leaf VTEP access port will be advertised to the Root VTEP and other Leaf VTEPs.
- MAC IPv4 or IPv6 configured under an MH Leaf VTEP access port must be symmetric and will be advertised to both the Root VTEP and other leaf VTEPs.
- MAC IPv4 or IPv6 configured under either SH or MH Root VTEP will be advertised to both the Root VTEP and the Leaf VTEPs.
- The Leaf-to-Leaf communication will display MAC status and tunnel status per VNI as Leaf type. The MAC will be in the discard state in the BCM shell.

```
VTEP1#show nvo vxlan mac-table
================================================================================================================================
                                                VXLAN MAC Entries
================================================================================================================================
VNID Interface VlanId  In-VlanId Mac-Addr       VTEP-Ip/ESI                   Type   Status  MAC move AccessPortDesc LeafFlag
--------------------------------------------------------------------------------------------------------------------------------
10   po1       1000    ----     0000.1000.1000  00:00:00:00:00:11:11:00:00:00 Static Local  ------- 0 -------         ----
10   po1       1000    ----     0000.1000.1001 00:00:00:00:00:11:11:00:00:00  Static Local ------- 0 -------          ----
10   ----      ----    ----     0000.2000.2000 2.2.2.2                        Static Remote   -------  0  -------     ----
10   ----      ----    ----     0000.2000.2001 2.2.2.2                        Static Remote   -------    0  -------   ----
10   ----      ----    ----     0000.3000.3000 00:00:00:00:00:22:22:00:00:00  Static Remote   -------    0  -------   set
10   ----      ----    ----     0000.3000.3001 00:00:00:00:00:22:22:00:00:00  Static Remote   -------    0   -------  set
10   ----      ----    ----     0000.4000.4000 4.4.4.4                        Static Remote   -------    0  -------   set
10   ----      ----    ----     0000.4000.4001 4.4.4.4                        Static Remote   -------    0 -------    set

Total number of entries are : 8

VTEP3#show nvo vxlan mac-table
================================================================================================================================
                                                VXLAN MAC Entries
================================================================================================================================
VNID Interface VlanId In-VlanId Mac-Addr       VTEP-Ip/ESI                 Type  Status  MAC move AccessPortDesc LeafFlag
--------------------------------------------------------------------------------------------------------------------------------


10  ----      ----    ----      0000.1000.1000 00:00:00:00:00:11:11:00:00:00  Static Remote  -------  0  -------    ----
10  ----      ----    ----      0000.1000.1001 00:00:00:00:00:11:11:00:00:00  Static Remote   ------- 0  -------    ----
10  ----      ----    ----      0000.2000.2000 2.2.2.2                        Static Remote   ------- 0   -------   ----
10  ----      ----    ----      0000.2000.2001 2.2.2.2                        Static Remote   -------  0 -------     ----
```

```
10  po2       1000    ----      0000.3000.3000 00:00:00:00:00:22:22:00:00:00  Static Local    -------  0 -------     set
10  po2      1000     ----      0000.3000.3001 00:00:00:00:00:22:22:00:00:00  Static Local    -------  0  -------    set
10  ----      ----    ----      0000.4000.4000 4.4.4.4                        Static Remote   -------   0  -------   set
10  ----      ----    ----      0000.4000.4001 4.4.4.4                        Static Remote   -------   0   ------- set

Total number of entries are : 8
```

Use the `show nvo vxlan arp-cache` command to verify the Address Resolution Protocol (ARP) cache information on all VTEP nodes. This command displays entries that map IPv4 addresses to MAC addresses within the specified VXLAN VNID network.

```
VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________
10       100.100.100.1    0000.1000.1000 Static      Local      ----
10       103.103.103.1    0000.3000.3000 Static      Remote     ----
10       104.104.104.1    0000.4000.4000 Static    Remote     ----
10       200.200.200.1    0000.2000.2000 Static    Remote     ----
Total number of entries are 4

VTEP3#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________
10       100.100.100.1    0000.1000.1000 Static Remote  ----
10       103.103.103.1    0000.3000.3000 Static Local   ----
10       104.104.104.1    0000.4000.4000 Static Remote  ----
10       200.200.200.1    0000.2000.2000 Static Remote  ----
Total number of entries are 4
```

Use the `show nvo vxlan nd-cache` command to verify the Neighbor Discovery (ND) cache information on all VTEP nodes. This command displays entries that map IPv6 addresses to MAC addresses within the specified VXLAN VNID network.

```
VTEP1#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr         Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________
10       1000::1         0000.1000.1001 Static Local     ----
10       1003::1         0000.3000.3001 Static Remote    ----
10       1004::1         0000.4000.4001 Static Remote    ----
10       2000::1         0000.2000.2001 Static Remote    ----
Total number of entries are 4

VTEP3#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr                                  Mac-Addr       Type         Age-Out
Retries-Left
_____________________________________________________________________________
_____________
10       1000::1                        0000.1000.1001 Static Remote    ----
```

```
10       1003::1                          0000.3000.3001 Static Local     ----
10       1004::1                          0000.4000.4001 Static Remote    ----
10       2000::1                          0000.2000.2001 Static Remote    ----
Total number of entries are 4
```

Network Topology Snippet Configurations

Here are the snippet configurations for all nodes in the given network topology.

VTEP1

```
    !
    nvo vxlan enable
    !
    evpn esi hold-time 90
    !
    evpn vxlan multihoming enable
    !
    mac vrf VRF1
     rd 1.1.1.1:100
     route-target both 100:100
    !
    nvo vxlan vtep-ip-global 1.1.1.1
    !
    nvo vxlan id 10 ingress-replication inner-vid-disabled
     vxlan host-reachability-protocol evpn-bgp VRF1
    !
    qos enable
    !
    interface po1
     switchport
     evpn multi-homed system-mac 0000.0000.1111
    !
    interface lo
     ip address 1.1.1.1/32 secondary
    !
    interface xe7
     channel-group 1 mode active
    !
    interface xe45
     ip address 10.10.10.1/24
    !
    interface xe49/2
     ip address 10.10.11.1/24
    !
     exit
    !

    router ospf 100
     ospf router-id 1.1.1.1
     bfd all-interfaces
     network 1.1.1.1/32 area 0.0.0.0
     network 10.10.10.0/24 area 0.0.0.0
     network 10.10.11.0/24 area 0.0.0.0
    !
    router bgp 100
     bgp router-id 1.1.1.1
```

```
 neighbor 2.2.2.2 remote-as 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 2.2.2.2 activate
 neighbor 3.3.3.3 activate
 neighbor 4.4.4.4 activate
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan po1 1000
 map vnid 10
 mac 0000.1000.1000 ip 100.100.100.1
 mac 0000.1000.1001 ipv6 1000::1
!
```

VTEP2

```
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf VRF1
 rd 2.2.2.2:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 10 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po1
 switchport
 evpn multi-homed system-mac 0000.0000.1111
!
interface lo
 ip address 2.2.2.2/32 secondary
!
interface xe38
 channel-group 1 mode active
!
interface xe49/1
 ip address 20.20.20.1/24
```

```
!
interface xe50/1
 ip address 20.20.21.1/24
!
 exit
!

router ospf 100
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0
 network 20.20.20.0/24 area 0.0.0.0
 network 20.20.21.0/24 area 0.0.0.0
!
router bgp 100
 bgp router-id 2.2.2.2
 neighbor 1.1.1.1 remote-as 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 3.3.3.3 activate
 neighbor 4.4.4.4 activate
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan xe37 1000
 map vnid 10
 mac 0000.2000.2000 ip 200.200.200.1
 mac 0000.2000.2001 ipv6 2000::1
!
nvo vxlan access-if port-vlan po1 1000
 map vnid 10
 mac 0000.1000.1000 ip 100.100.100.1
 mac 0000.1000.1001 ipv6 1000::1
!
```

VTEP3

```
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
evpn etree enable
!
```

```
mac vrf VRF1
 rd 3.3.3.3:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 3.3.3.3
!
nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
 vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po2
 switchport
 evpn multi-homed system-mac 0000.0000.2222
!
interface lo
 ip address 3.3.3.3/32 secondary
!
interface xe53/1
 ip address 30.30.30.1/24
!
interface xe54/1
 ip address 30.30.31.1/24
!
interface xe55/1
 channel-group 2 mode active
!
 exit
!
router ospf 100
 ospf router-id 3.3.3.3
 bfd all-interfaces
 network 3.3.3.3/32 area 0.0.0.0
 network 30.30.30.0/24 area 0.0.0.0
 network 30.30.31.0/24 area 0.0.0.0
!
router bgp 100
 bgp router-id 3.3.3.3
 neighbor 1.1.1.1 remote-as 100
 neighbor 2.2.2.2 remote-as 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 2.2.2.2 activate
 neighbor 4.4.4.4 activate
 exit-address-family
 !
 exit
!
!
```

```
nvo vxlan access-if port-vlan po2 1000
 map vnid 10
 mac 0000.3000.3000 ip 103.103.103.1
 mac 0000.3000.3001 ipv6 1003::1
!
```

VTEP4

```
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
evpn etree enable
!
mac vrf VRF1
 rd 4.4.4.4:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 4.4.4.4
!
nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
 vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po2
 switchport
 evpn multi-homed system-mac 0000.0000.2222
!
interface lo
 ip address 4.4.4.4/32 secondary
!
interface xe11/1
 ip address 40.40.41.1/24
!
interface xe31/1
 channel-group 2 mode active
!
interface xe33
 ip address 40.40.40.1/24
!
interface xe34
 switchport
!
 exit
!
router ospf 100
 ospf router-id 4.4.4.4
 bfd all-interfaces
 network 4.4.4.4/32 area 0.0.0.0
 network 40.40.40.0/24 area 0.0.0.0
 network 40.40.41.0/24 area 0.0.0.0
!
```

```
router bgp 100
 bgp router-id 4.4.4.4
 neighbor 1.1.1.1 remote-as 100
 neighbor 2.2.2.2 remote-as 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 2.2.2.2 activate
 neighbor 3.3.3.3 activate
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan xe34 1000
 map vnid 10
 mac 0000.4000.4000 ip 104.104.104.1
 mac 0000.4000.4001 ipv6 1004::1
!
nvo vxlan access-if port-vlan po2 1000
 map vnid 10
 mac 0000.3000.3000 ip 103.103.103.1
 mac 0000.3000.3001 ipv6 1003::1
!
```

SPINE1

```
!
qos enable
!
interface ce1/2
 ip address 40.40.40.2/24
!
interface ce1/4
 ip address 10.10.10.2/24
!
interface ce24/1
 ip address 30.30.30.2/24
!
interface ce27/1
 ip address 20.20.20.2/24
!
interface lo
 ip address 5.5.5.5/32 secondary
!
 exit
!
router ospf 100
 ospf router-id 5.5.5.5
 bfd all-interfaces
```

```
 network 5.5.5.5/32 area 0.0.0.0
 network 10.10.10.0/24 area 0.0.0.0
 network 20.20.20.0/24 area 0.0.0.0
 network 30.30.30.0/24 area 0.0.0.0
 network 40.40.40.0/24 area 0.0.0.0
!
```

SPINE2

```
!
qos enable
!
interface ce5/1
 ip address 20.20.21.2/24
!
interface ce10/1
 ip address 30.30.31.2/24
!
interface ce11/1
 ip address 40.40.41.2/24
!
interface ce14/2
 ip address 10.10.11.2/24
!
interface lo
 ip address 6.6.6.6/32 secondary
!
 exit
!
router ospf 100
 ospf router-id 6.6.6.6
 bfd all-interfaces
 network 6.6.6.6/32 area 0.0.0.0
 network 10.10.11.0/24 area 0.0.0.0
 network 20.20.21.0/24 area 0.0.0.0
 network 30.30.31.0/24 area 0.0.0.0
 network 40.40.41.0/24 area 0.0.0.0
!
```

SWITCH1

```
!
bridge 1 protocol ieee vlan-bridge
!
vlan database
 vlan-reservation 4000-4094
 vlan 1000 bridge 1 state enable
!
interface po1
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe46
```

```
 channel-group 1 mode active
!
interface xe47
 channel-group 1 mode active
!
interface xe57
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
 exit
!
```

SWITCH2

```
!
bridge 1 protocol ieee vlan-bridge
!
vlan database
 vlan-reservation 4000-4094
 vlan 1000 bridge 1 state enable
!
interface po2
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe33
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe49/1
 channel-group 2 mode active
!
interface xe51/1
 channel-group 2 mode active
!
 exit
!
```

Implementation Examples

Here is an example scenario and a solution for implementing EVPN E-Tree.

Scenario 1: Specific traffic isolation and control measures are essential in a network of EVPN L2VPN services or instances. Within a broadcast domain, services communicating with each other may result in flooding BUM traffic to all services within the domain. Moreover, hosts are learned and advertised between different sites/services.

Use Case 1: Implementing an EVPN E-Tree solution defines the network topology with distinct Root and Leaf classifications, BUM traffic flooding can be minimized, and traffic isolation can be achieved. This ensures efficient communication between services while preventing unnecessary traffic propagation and maintaining network integrity.

Scenario 2: An Internet Service Provider (ISP) provides services to multiple subscribers and aims to facilitate communication with them. However, the ISP needs to ensure that subscribers exclusively communicate with the ISP and not among themselves.

Use Case 2: Implementing EVPN E-Tree is essential to fulfill this requirement. By categorizing ISP services as Root and subscribers as Leaf, traffic isolation can be enforced. This configuration enables the ISP to communicate with subscribers while preventing inter-subscriber communication. As a result, network security is enhanced, and the ISP maintains control over communication within its network.

E-Tree CLI Commands

The EVPN E-Tree introduces the following configuration commands in OcNOS.

evpn etree

Use this command to enable E-Tree functionality within the EVPN configuration.

Command Syntax

```
evpn etree enable
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates how to activate E-Tree functionality for EVPN:

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable
```

Revised CLI Commands

The following is the revised command for configuring VXLAN EVPN E-Tree

nvo vxlan id

- The existing syntax now includes the newly added parameter for E-Tree, namely `etree-leaf`.
- The command `nvo vxlan id <VNID> ingress-replication inner-vid-disabled etree-leaf` allows users to tailor VXLAN behavior on a network device, specifying VXLAN parameters and indicating its participation as a leaf node in an E-Tree deployment. For more details, refer to the nvo vxlan id command in the VXLAN Commands chapter in the *OcNOS VXLAN Guide*.

Troubleshooting

1. When traffic, whether unicast (UC) or broadcast, is passed to the Intra Leaf site:
   - Check the sub-interface or physical interface counters to monitor traffic throughput and potential issues.
   - Verify the Leaf status of the corresponding VNI to ensure proper functionality.
   - Use packet sniffing tools to analyze packets in the egress direction for any anomalies or errors.
   - MAC entries learned via leaf access port should include the `set` keyword in the MAC table output.
2. If UC traffic is routed within inter-PE leaf sites:
   - Check the Leaf status of the VNI at both participating PE devices to confirm operational status.
   - Check if the advertised MAC is in discard or non-discard status using the `show mac table` command and `l2 show` in the BCM shell.
3. Investigate UC traffic drops from the Root to MH Leaf PE:
   - Check if MAC addresses are not installed in discard status within the MH peer's access port. This status could indicate issues with MAC learning or forwarding.
4. Evaluate traffic between Root and Leaf:
   - Confirm the establishment of both UC and BUM tunnels.
   - Ensure that unicast MAC addresses are not marked with a discard status in the MAC table.
5. Validate the exchange of routes between two BGP L2VPN peers:
   - Monitor BGP (Border Gateway Protocol) sessions to verify successful route exchange and propagation between the peers.
6. Convergence: Assess convergence by checking BFD configuration between BGP sessions.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Ethernet VPN Ethernet-Tree (EVPN E-Tree)	A networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, categorizing network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them.

Virtual Extensible LAN (VXLAN)	A technology that provides encapsulation techniques to create virtualized Layer 2 networks over Layer 3 infrastructure, facilitating scalable and flexible network designs.
Ethernet Virtual Private Network (EVPN)	A Layer 2 VPN technology that extends Ethernet services across data centers and wide-area networks using BGP.
Multi-homing (MH)	The ability of a device to connect to multiple network segments simultaneously to increase network availability and redundancy.
Provider Edge (PE) Node	A device at the edge of a service provider network that connects to customer premises equipment (CE) and participates in providing services to customers.
Leaf Node	In the context of EVPN E-Tree, a network node categorized to handle communication within specific broadcast domains and may connect to Root nodes.
Root Node	A network node within EVPN E-Tree that serves as the central point of communication and handles BUM traffic distribution.
Ethernet Segment Identifier (ESI)	A unique identifier used to identify Ethernet segments within a VXLAN network.

CHAPTER 3 VXLAN Tunnel Over SVI

This chapter contains the configurations for VXLAN Tunnel Over SVI.

Overview

VxLAN EVPN solution is envisioned to simplify the topology and configurations in Data Centers (DC). In Data Centers, CLOS topology was used, which makes network side pure L3 and uses EBGP as IGP.

VxLAN solution is required for Service Providers (SP) as well to run few of the services or all services in their network over VxLAN. When they choose to run few services over VxLAN, then on the network side there will be a need to run VxLAN over SVI.

Topology

The Topology shown below contains the 3 VTEPS i.e VTEP1 ,VTEP2 and VTEP3 and 3 core nodes P1 ,P2 and P3. Vxlan tunnel will be established between VTEPS over SVI interfaces. OSPF as IGP will be running between VTEPS and the core node to provide the end to end connectivity. Switch is connected between host and VTEP-1, VTEP-2 via dynamic LAG.

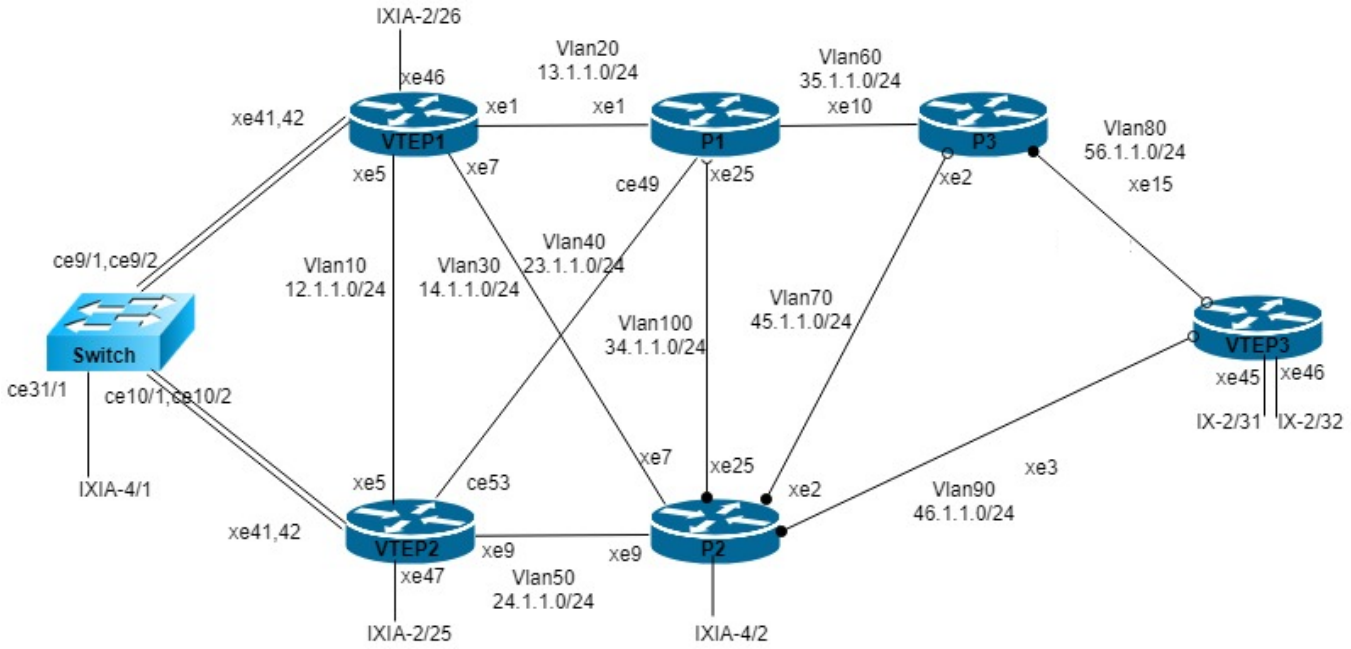

Figure 3-7: VXLAN over SVI

RTR1/VTEP1

`VTEP1#configure terminal`	Enter configuration terminal
`VTEP1(config)#hostname VTEP1`	Configure hostname
`VTEP1(config)#mac vrf vrf1`	Configure MAC vrf vrf1
`VTEP1(config-vrf)#rd 1.1.1.1:11`	Configure RD for vrf1

`VTEP1(config-vrf)#route-target both 10.10.10.10:100`	Configure RT for vrf1
`VTEP1(config-vrf)#mac vrf vrf2`	Configure MAC vrf vrf2
`VTEP1(config-vrf)#rd 1.1.1.1:12`	Configure RD for vrf2
`VTEP1(config-vrf)#route-target both 10.10.10.10:102`	Configure RT for vrf2
`VTEP1(config-vrf)#bfd interval 3 minrx 3 multiplier 3`	Configure BFD interval globally
`VTEP1(config)#evpn vxlan multihoming enable`	Enable EVPN VXLAN multihoming
`VTEP1(config)#nvo vxlan enable`	Enable VXLAN
`VTEP1(config)#qos enable`	Enable qos
`VTEP1(config)#qos statistics`	Enable qos statistics
`VTEP1(config)#bridge 1 protocol ieee vlan-bridge`	Configure IEEE VLAN bridge
`VTEP1(config)#no bridge 1 spanning-tree enable`	Disable spanning tree in bridge 1 globally
`VTEP1(config)#no igmp snooping`	Disable igmp snooping messages globally
`VTEP1(config)#vlan database`	Enter into the VLAN database
`VTEP1(config-vlan)#vlan 10-200 bridge 1 state enable`	Configure VLAN 10-200 and associate with bridge 1
`VTEP1(config-vlan)#vlan 4000 bridge 1 state enable`	Configure VLAN 4000 and associate with bridge 1
`VTEP1(config-vlan)#interface po1`	Enter interface mode
`VTEP1(config-if)#switchport`	Set the interface as Layer2 port
`VTEP1(config-if)#load-interval 30`	Configure load period in multiple of 30 seconds
`VTEP1(config-if)#evpn multi-homed system-mac aa22.3344.5566`	Configure EVPN system-MAC
`VTEP1(config-if)#interface lo`	Enter lo interface mode
`VTEP1(config-if)#ip address 1.1.1.1/32 secondary`	Assign loopback IP
`VTEP1(config-if)#interface vlan1.10`	Specify interface VLAN1.10 to be configured.
`VTEP1(config-if)#ip address 12.1.1.1/24`	Assign IP address
`VTEP1(config-if)#ip ospf cost 1`	Change OSPF cost of the link
`VTEP1(config-if)#interface vlan1.20`	Specify interface VLAN1.20 to be configured.
`VTEP1(config-if)#ip address 13.1.1.1/24`	Assign IP address
`VTEP1(config-if)#ip ospf cost 1`	Change OSPF cost of the link
`VTEP1(config-if)#interface vlan1.30`	Specify interface VLAN1.30 to be configured.
`VTEP1(config-if)#ip address 14.1.1.1/24`	Assign IP address
`VTEP1(config-if)#ip ospf cost 1`	Change OSPF cost of the link
`VTEP1(config-if)#interface xe1`	Enter interface mode
`VTEP1(config-if)#switchport`	Set the interface as Layer2 port
`VTEP1(config-if)#bridge-group 1 spanning-tree disable`	Associate the interface with bridge group 1 and disable spanning tree
`VTEP1(config-if)#switchport mode trunk`	Set the switching characteristics of this interface to trunk mode.

VTEP1(config-if)#switchport trunk allowed vlan add 20,29	Enable VLAN's allowed on this interface.
VTEP1(config-if)#switchport trunk native vlan 29	Configure native VLAN
VTEP1(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP1(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
VTEP1(config-if)#interface xe5	Enter interface mode
VTEP1(config-if)#switchport	Set the interface as Layer2 port
VTEP1(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
VTEP1(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
VTEP1(config-if)#switchport trunk allowed vlan add 10,19	Enable VLAN's allowed on this interface.
VTEP1(config-if)#switchport trunk native vlan 19	Configure native VLAN
VTEP1(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP1(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
VTEP1(config-if)#interface xe7	Enter interface mode
VTEP1(config-if)#switchport	Set the interface as Layer2 port
VTEP1(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
VTEP1(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
VTEP1(config-if)#switchport trunk allowed vlan add 30,39	Enable VLAN's allowed on this interface.
VTEP1(config-if)#switchport trunk native vlan 39	Configure native VLAN
VTEP1(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP1(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
VTEP1(config-if)#interface xe41	Enter interface mode
VTEP1(config-if)#channel-group 1 mode active	Map this interface to po1
VTEP1(config-if)#interface xe42	Enter interface mode
VTEP1(config-if)#channel-group 1 mode active	Map this interface to po1
VTEP1(config-if)#interface xe46	Enter interface mode
VTEP1(config-if)#switchport	Map this interface to po1
VTEP1(config-if)#router ospf 1	Enter OSPF configuration mode
VTEP1(config-router)#ospf router-id 1.1.1.1	Configure OSPF router id
VTEP1(config-router)#bfd all-interfaces	Enable BFD in all OSPF interfaces
VTEP1(config-router)#network 1.1.1.1/32 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
VTEP1(config-router)#network 12.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
VTEP1(config-router)#network 13.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.

`VTEP1(config-router)#network 14.1.1.0/24 area 0.0.0.0`	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
`VTEP1(config-router)#router bgp 100`	Enter Router BGP mode and define the AS number 100.
`VTEP1(config-router)#address-family ipv4 unicast`	Enter address-family IPv4 unicast mode
`VTEP1(config-router)#network 1.1.1.1/32`	Add the lo network to BGP route
`VTEP1(config-router-af)#neighbor 2.2.2.2 activate`	Activate neigbors
`VTEP1(config-router-af)#neighbor 6.6.6.6 activate`	Activate neigbors
`VTEP1(config-router-af)#exit-address-family`	Exit address-family mode.
`VTEP1(config-router)#neighbor 2.2.2.2 remote-as 100`	Configure BGP remote-as 100 with neighbor IP
`VTEP1(config-router)#neighbor 2.2.2.2 update-source lo`	Define BGP neighbors, to update the source routes with lo
`VTEP1(config-router)#neighbor 2.2.2.2 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP2
`VTEP1(config-router)#neighbor 6.6.6.6 remote-as 100`	Configure BGP remote-as 100 with neighbor IP
`VTEP1(config-router)#neighbor 6.6.6.6 update-source lo`	Define BGP neighbors, to update the source routes with lo
`VTEP1(config-router)#neighbor 6.6.6.6 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP3
`VTEP1(config-router)#address-family l2vpn evpn`	Enter in to BGP L2VPN EVPN address-family
`VTEP1(config-router-af)#neighbor 2.2.2.2 activate`	Activate neigbors
`VTEP1(config-router-af)#neighbor 6.6.6.6 activate`	Activate neigbors
`VTEP1(config-router-af)#exit-address-family`	Exit from BGP L2VPN EVPN address-family
`VTEP1(config-router)#nvo vxlan vtep-ip-global 1.1.1.1`	Configure VXLAN global IP
`VTEP1(config)#nvo vxlan id 1 ingress-replication inner-vid-disabled`	Create VNID 1
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1`	Associate VNID with EVPN and vrf1
`VTEP1(config-nvo)#nvo vxlan id 1000 ingress-replication inner-vid-disabled`	Create VNID 1000
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf2`	Associate VNID with EVPN and vrf2
`VTEP1(config-nvo-acc-if)#nvo vxlan access-if port-vlan po1 2000`	Create VXLAN access port port-vlan
`VTEP1(config-nvo-acc-if)#no shutdown`	Unshut the access interface
`VTEP1(config-nvo-acc-if)#map vnid 1000`	Map the VNID to access-if
`VTEP1(config-nvo)#nvo vxlan access-if port xe46`	Create VXLAN access port
`VTEP1(config-nvo-acc-if)#no shutdown`	Unshut the access interface
`VTEP1(config-nvo-acc-if)#map vnid 1`	Map the VNID to access-if

`VTEP1(config-nvo-acc-if)#nvo vxlan access-if port-vlan po1 1000`	Create VXLAN access port port-vlan
`VTEP1(config-nvo-acc-if)#no shutdown`	Unshut the access interface
`VTEP1(config-nvo-acc-if)#map vnid 1`	Map the VNID to access-if
`VTEP1(config-nvo-acc-if)#mac 0000.1111.1111 ip 100.1.1.100`	Configure static MAC IP
`VTEP1(config-nvo-acc-if)#mac 0000.1111.1112`	Configure static MAC
`VTEP1(config-nvo-acc-if)#mac 0000.1111.1113`	Configure static MAC
`VTEP1(config-nvo-acc-if)#mac 0000.1111.1114`	Configure static MAC
`VTEP1(config-nvo-acc-if)#mac 0000.1111.1115`	Configure static MAC
`VTEP1(config-nvo-acc-if)#commit`	Commit the candidate configuration to the running configuration

VTEP2

`VTEP2#configure terminal`	Enter configuration terminal
`VTEP2(config)#hostname VTEP2`	Configure hostname
`VTEP2(config)#mac vrf vrf1`	Configure MAC VRF vrf1
`VTEP2(config-vrf)#rd 2.2.2.2:11`	Configure RD for vrf1
`VTEP2(config-vrf)#route-target both 10.10.10.10:100`	Configure RT for vrf1
`VTEP2(config-vrf)#mac vrf vrf2`	Configure MAC VRF vrf2
`VTEP2(config-vrf)#rd 2.2.2.2:12`	Configure RD for vrf2
`VTEP2(config-vrf)#route-target both 10.10.10.10:102`	Configure RT for vrf2
`VTEP2(config)#nvo vxlan enable`	Enable VXLAN
`VTEP2(config)#evpn vxlan multihoming enable`	Enable EVPN VXLAN multihoming
`VTEP2(config)#qos enable`	Enable QOS
`VTEP2(config)#qos statistics`	Enable QOS statistics
`VTEP2(config)#bridge 1 protocol ieee vlan-bridge`	Configure IEEE VLAN bridge
`VTEP2(config)#no bridge 1 spanning-tree enable`	Disable spanning tree in bridge 1 globally
`VTEP2(config)#vlan database`	Enter into the VLAN database
`VTEP2(config-vlan)#vlan 10-200 bridge 1 state enable`	Configure VLAN 10-200 and associate with bridge 1
`VTEP2(config-vlan)#interface po1`	Enter interface mode
`VTEP2(config-if)#switchport`	Set the interface as Layer2 port
`VTEP2(config-if)#load-interval 30`	Configure load period in multiple of 30 seconds
`VTEP2(config-if)#evpn multi-homed system-mac aa22.3344.5566`	Configure EVPN system-MAC
`VTEP2(config-if)#interface lo`	Enter interface mode
`VTEP2(config-if)#ip address 2.2.2.2/32 secondary`	Configure loopback IP
`VTEP2(config-if)#interface vlan1.10`	Specify interface VLAN1.10 to be configured.

VTEP2(config-if)#ip address 12.1.1.2/24	Assign IP address
VTEP2(config-if)#interface vlan1.40	Specify interface VLAN1.40 to be configured.
VTEP2(config-if)#ip address 23.1.1.1/24	Assign IP address
VTEP2(config-if)#ip ospf cost 1	Change OSPF cost of the link
VTEP2(config-if)#interface vlan1.50	Specify interface VLAN1.50 to be configured.
VTEP2(config-if)#ip address 24.1.1.1/24	Assign IP address
VTEP2(config-if)#interface ce53	Enter interface mode
VTEP2(config-if)#switchport	Set the interface as Layer2 port
VTEP2(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
VTEP2(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
VTEP2(config-if)#switchport trunk allowed vlan add 40,49	Enable VLAN's allowed on this interface.
VTEP2(config-if)#switchport trunk native vlan 49	Configure native VLAN
VTEP2(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP2(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
VTEP2(config-if)#interface xe5	Enter interface mode
VTEP2(config-if)#switchport	Set the interface as Layer2 port
VTEP2(config-if)#bridge-group 1 spanningtree disable	Associate the interface with bridge group 1 and disable spanning tree
VTEP2(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode
VTEP2(config-if)#switchport trunk allowed vlan add 10,19	Enable VLAN's allowed on this interface
VTEP2(config-if)#switchport trunk native vlan 19	Configure native VLAN
VTEP2(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP2(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
VTEP2(config-if)#interface xe9	Enter interface mode
VTEP2(config-if)#switchport	Set the interface as Layer2 port
VTEP2(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
VTEP2(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
VTEP2(config-if)#switchport trunk allowed vlan add 50,59	Enable VLAN's allowed on this interface.
VTEP2(config-if)#switchport trunk native vlan 59	Configure native VLAN
VTEP2(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP2(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
VTEP2(config-if)#interface xe41	Enter interface mode
VTEP2(config-if)#channel-group 1 mode active	Map the interface to po1
VTEP2(config-if)#interface xe42	Enter interface mode

`VTEP2(config-if)#channel-group 1 mode active`	Map the interface to po1
`VTEP2(config-if)#interface xe47`	Enter interface mode
`VTEP2(config-if)#switchport`	Set the interface as Layer2 port
`VTEP2(config-if)#router ospf 1`	Enter OSPF configuration mode
`VTEP2(config-router)#ospf router-id 2.2.2.2`	Configure OSPF router id
`VTEP2(config-router)#network 2.2.2.2/32 area 0.0.0.0`	Enable BFD in all ospf interfaces
`VTEP2(config-router)#network 12.1.1.0/24 area 0.0.0.0`	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
`VTEP2(config-router)#network 23.1.1.0/24 area 0.0.0.0`	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
`VTEP2(config-router)#network 24.1.1.0/24 area 0.0.0.0`	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
`VTEP2(config-router)#network 25.1.1.0/24 area 0.0.0.0`	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
`VTEP2(config-router)#router bgp 100`	Enter Router BGP mode and define the AS number 100.
`VTEP2(config-router)#address-family ipv4 unicast`	Enter address-family IPv4 unicast mode
`VTEP2(config-router)#network 2.2.2.2/32`	Add the lo network to BGP route
`VTEP2(config-router-af)#neighbor 1.1.1.1 activate`	Activate neigbors
`VTEP2(config-router-af)#neighbor 6.6.6.6 activate`	Activate neigbors
`VTEP2(config-router-af)#exit-address-family`	Exit address-family mode.
`VTEP2(config-router)#neighbor 1.1.1.1 remote-as 100`	Configure BGP remote-as 100 with neighbor IP
`(config-router)#neighbor 1.1.1.1 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP1
`VTEP2(config-router)#neighbor 1.1.1.1 update-source lo`	Define BGP neighbors, to update the source routes with lo
`VTEP2(config-router)#neighbor 6.6.6.6 remote-as 100`	Configure BGP remote-as 100 with neighbor IP
`(config-router)#neighbor 6.6.6.6 advertisement-interval 0`	Configure advertisement-interval as 0 for fast convergence for VTEP3
`VTEP2(config-router)#neighbor 6.6.6.6 update-source lo`	Define BGP neighbors, to update the source routes with lo
`VTEP2(config-router)#address-family l2vpn evpn`	Enter in to BGP L2VPN EVPN address-family
`VTEP2(config-router-af)#neighbor 1.1.1.1 activate`	Activate neigbors
`VTEP2(config-router-af)#neighbor 6.6.6.6 activate`	Activate neigbors
`VTEP2(config-router-af)#exit-address-family`	Exit from BGP L2VPN EVPN address-family
`VTEP2(config-router)#nvo vxlan vtep-ip-global 2.2.2.2`	Configure VXLAN global IP
`VTEP2(config)#nvo vxlan id 1 ingress-replication inner-vid-disabled`	Create VNID 1

`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1`	Associate VNID with EVPN and vrf1
`VTEP2(config-nvo)#nvo vxlan id 1000 ingress-replication inner-vid-disabled`	Create VNID 1000
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf2`	Associate VNID with EVPN and vrf2
`VTEP2(config-nvo)#nvo vxlan access-if port xe47`	Create VXLAN access port
`VTEP2(config-nvo-acc-if)#map vnid 1`	Map the VNID to access-if
`VTEP2(config-nvo-acc-if)#nvo vxlan access-if port-vlan po1 2001`	Create VXLAN access port
`VTEP2(config-nvo-acc-if)#map vnid 1000`	Map the VNID to access-if
`VTEP2(config-nvo-acc-if)#nvo vxlan access-if port-vlan po1 2000`	Create VXLAN access port
`VTEP2(config-nvo-acc-if)#map vnid 1000`	Map the VNID to access-if
`VTEP2(config-nvo-acc-if)#commit`	Commit the candidate configuration to the running configuration

P1

`P1#configure terminal`	Enter configuration terminal
`P1(config)#hostname P1`	Configure hostname
`P1(config)#bfd interval 3 minrx 3 multiplier 3`	Configure BFD interval globally
`P1(config)#qos enable`	Enable qos
`P1(config)#bridge 1 protocol ieee vlan-bridge`	Configure IEEE VLAN bridge
`P1(config)#no bridge 1 spanning-tree enable`	Disable spanning tree in bridge 1 globally
`P1(config)#no igmp snooping`	Disable igmp snooping messages globally
`P1(config)#vlan database`	Enter into the VLAN database
`P1(config-vlan)#vlan 10-200 bridge 1 state enable`	Configure VLAN 10-200 and associate with bridge 1
`P1(config-if)#interface lo`	Enter loopback interface mode
`P1(config-if)#ip address 3.3.3.3/32 secondary`	Assign loopback IP
`P1(config-if)#interface vlan1.20`	Specify interface VLAN1.20 to be configured.
`P1(config-if)#ip address 13.1.1.2/24`	Assign IP address
`P1(config-if)#ip ospf cost 1`	Change OSPF cost of the link
`P1(config-if)#interface vlan1.40`	Specify interface VLAN1.40 to be configured.
`P1(config-if)#ip address 23.1.1.2/24`	Assign IP address
`P1(config-if)#interface vlan1.60`	Specify interface VLAN1.60 to be configured.
`P1(config-if)#ip address 35.1.1.1/24`	Assign IP address
`P1(config-if)#ip ospf cost 1`	Change OSPF cost of the link
`P1(config-if)#interface vlan1.100`	Specify interface VLAN1.100 to be configured.
`P1(config-if)#ip address 34.1.1.1/24`	Assign IP address

P1(config-if)#ip ospf cost 1	Change OSPF cost of the link
P1(config-vlan)#interface ce49	Enter interface mode
P1(config-if)#switchport	Set the interface as Layer2 port
P1(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P1(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P1(config-if)#switchport trunk allowed vlan add 40,49	Enable VLAN's allowed on this interface.
P1(config-if)#switchport trunk native vlan 49	Configure native VLAN
P1(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P1(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P1(config-if)#interface xe1	Enter interface mode
P1(config-if)#switchport	Set the interface as Layer2 port
P1(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P1(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P1(config-if)#switchport trunk allowed vlan add 20,29	Enable VLAN's allowed on this interface.
P1(config-if)#switchport trunk native vlan 29	Configure native VLAN
P1(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P1(config-if)#mtu 1600	Change the interface mtu value
P1(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P1(config-if)#interface xe10	Enter interface mode
P1(config-if)#switchport	Set the interface as Layer2 port
P1(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P1(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P1(config-if)#switchport trunk allowed vlan add 60,69	Enable VLAN's allowed on this interface.
P1(config-if)#switchport trunk native vlan 69	Configure native VLAN
P1(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P1(config-if)#mtu 1600	Change the interface mtu value
P1(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P1(config-if)#interface xe25	Enter interface mode
P1(config-if)#switchport	Set the interface as Layer2 port
P1(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P1(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.

P1(config-if)#switchport trunk allowed vlan add 100,109	Enable VLAN's allowed on this interface.
P1(config-if)#switchport trunk native vlan 109	Configure native VLAN
P1(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P1(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P1(config-if)#router ospf 1	Enter OSPF configuration mode
P1(config-router)#ospf router-id 3.3.3.3	Configure OSPF router id
P1(config-router)#bfd all-interfaces	Enable BFD in all OSPF interfaces
P1(config-router)#network 3.3.3.3/32 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P1(config-router)#network 10.10.10.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P1(config-router)#network 13.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P1(config-router)#network 23.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P1(config-router)#network 34.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P1(config-router)#network 35.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P1(config-router)#commit	Commit the candidate configuration to the running configuration

P2

P2#configure terminal	Enter configuration terminal
P2(config)#bfd interval 3 minrx 3 multiplier 3	Configure BFD interval globally
P2(config)#qos enable	Enable qos
P2(config)#hostname P2	Configure hostname
P2(config)#bridge 1 protocol ieee vlan-bridge	Configure IEEE VLAN bridge
P2(config)#no bridge 1 spanning-tree enable	Disable spanning tree in bridge 1 globally
P2(config)#no igmp snooping	Disable igmp snooping messages globally
P2(config)#vlan database	Enter into the VLAN database
P2(config-vlan)#vlan 10-200 bridge 1 state enable	Configure VLAN 10-200 and associate with bridge 1
P2(config-vlan)#interface lo	Enter lo interface mode
P2(config-if)#ip address 4.4.4.4/32 secondary	Assign loopback IP
P2(config-if)#interface vlan1.30	Specify interface VLAN1.30 to be configured.
P2(config-if)#ip address 14.1.1.2/24	Assign IP address
P2(config-if)#ip ospf cost 1	Change OSPF cost of the link
P2(config-if)#interface vlan1.50	Specify interface VLAN1.50 to be configured.
P2(config-if)#ip address 24.1.1.2/24	Assign IP address

P2(config-if)#interface vlan1.70	Specify interface VLAN1.70 to be configured.
P2(config-if)#ip address 45.1.1.1/24	Assign IP address
P2(config-if)#ip ospf cost 1	Change OSPF cost of the link
P2(config-if)#interface vlan1.90	Specify interface VLAN1.90 to be configured.
P2(config-if)#ip address 46.1.1.1/24	Assign IP address
P2(config-if)#ip ospf cost 1	Change OSPF cost of the link
P2(config-if)#interface vlan1.100	Specify interface VLAN1.100 to be configured.
P2(config-if)#ip address 34.1.1.2/24	Assign IP address
P2(config-if)#ip ospf cost 1	Change OSPF cost of the link
P2(config-if)#interface xe2	Enter interface mode
P2(config-if)#switchport	Set the interface as Layer2 port
P2(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P2(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P2(config-if)#switchport trunk allowed vlan add 70,79	Enable VLAN's allowed on this interface.
P2(config-if)#switchport trunk native vlan 79	Configure native VLAN
P2(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P2(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P2(config-if)#interface xe3	Enter interface mode
P2(config-if)#switchport	Set the interface as Layer2 port
P2(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P2(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P2(config-if)#switchport trunk allowed vlan add 90,99	Enable VLAN's allowed on this interface.
P2(config-if)#switchport trunk native vlan 99	Configure native VLAN
P2(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P2(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P2(config-if)#interface xe7	Enter interface mode
P2(config-if)#switchport	Set the interface as Layer2 port
P2(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P2(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P2(config-if)#switchport trunk allowed vlan add 30,39	Enable VLAN's allowed on this interface.
P2(config-if)#switchport trunk native vlan 39	Configure native VLAN
P2(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P2(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions

P2(config-if)#interface xe9	Enter interface mode
P2(config-if)#switchport	Set the interface as Layer2 port
P2(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P2(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P2(config-if)#switchport trunk allowed vlan add 50,59	Enable VLAN's allowed on this interface.
P2(config-if)#switchport trunk native vlan 59	Configure native VLAN
P2(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P2(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P2(config-if)#interface xe25	Enter interface mode
P2(config-if)#switchport	Set the interface as Layer2 port
P2(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P2(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P2(config-if)#switchport trunk allowed vlan add 100,109	Enable VLAN's allowed on this interface.
P2(config-if)#switchport trunk native vlan 109	Configure native VLAN
P2(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P2(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P2(config-if)#router ospf 1	Enter OSPF configuration mode
P2(config-router)#ospf router-id 4.4.4.4	Configure OSPF router id
P2(config-router)#bfd all-interfaces	Enable BFD in all OSPF interfaces
P2(config-router)#network 4.4.4.4/32 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P2(config-router)#network 14.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P2(config-router)#network 24.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P2(config-router)#network 34.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P2(config-router)#network 45.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P2(config-router)#network 46.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P2(config-router)#commit	Commit the candidate configuration to the running configuration

P3

P3(config)#hostname P3	Configure hostname
P3(config)#bfd interval 3 minrx 3 multiplier 3	Configure BFD interval globally

P3(config)#qos enable	Enable qos
P3(config)#bridge 1 protocol ieee vlan-bridge	Configure IEEE VLAN bridge
P3(config)#no bridge 1 spanning-tree enable	Disable spanning tree in bridge 1 globally
P3(config)#no igmp snooping	Disable igmp snooping messages globally
P3(config)#vlan database	Enter into the VLAN database
P3(config-vlan)#vlan 10-200 bridge 1 state enable	Configure VLAN 10-200 and associate with bridge 1
P3(config-vlan)#interface lo	Enter lo interface mode
P3(config-if)#ip address 5.5.5.5/32 secondary	Assign loopback IP
P3(config-if)#interface vlan1.60	Specify interface VLAN1.60 to be configured.
P3(config-if)#ip address 35.1.1.2/24	Assign IP address
P3(config-if)#ip ospf cost 1	Change OSPF cost of the link
P3(config-if)#interface vlan1.70	Specify interface VLAN1.70 to be configured.
P3(config-if)#ip address 45.1.1.2/24	Assign IP address
P3(config-if)#ip ospf cost 1	Change OSPF cost of the link
P3(config-if)#interface vlan1.80	Specify interface VLAN1.80 to be configured.
P3(config-if)#ip address 56.1.1.1/24	Assign IP address
P3(config-if)#ip ospf cost 1	Change OSPF cost of the link
P3(config-if)#interface vlan1.110	Specify interface VLAN1.1100 to be configured.
P3(config-if)#ip address 15.1.1.2/24	Assign IP address
P3(config-if)#ip ospf cost 1	Change OSPF cost of the link
P3(config-if)#interface vlan1.160	Specify interface VLAN1.160 to be configured.
P3(config-if)#ip address 25.1.1.2/24	Assign IP address
P3(config-if)#ip ospf cost 1	Change OSPF cost of the link
P3(config-if)#interface xe2	Enter interface mode
P3(config-if)#switchport	Set the interface as Layer2 port
P3(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P3(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P3(config-if)#switchport trunk allowed vlan add 70,79	Enable VLAN's allowed on this interface.
P3(config-if)#switchport trunk native vlan 79	Configure native VLAN
P3(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P3(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P3(config-if)#interface xe10	Enter interface mode
P3(config-if)#switchport	Set the interface as Layer2 port
P3(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P3(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.

P3(config-if)#switchport trunk allowed vlan add 60,69	Enable VLAN's allowed on this interface.
P3(config-if)#switchport trunk native vlan 69	Configure native VLAN
P3(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P3(config-if)#mtu 1600	Change interface mtu value
P3(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P3(config-if)#interface xe15	Enter interface mode
P3(config-if)#switchport	Set the interface as Layer2 port
P3(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
P3(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
P3(config-if)#switchport trunk allowed vlan add 80,89	Enable VLAN's allowed on this interface.
P3(config-if)#switchport trunk native vlan 89	Configure native VLAN
P3(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
P3(config-if)#mtu 1600	Change interface mtu value
P3(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
P3(config-if)#router ospf 1	Enter OSPF configuration mode
P3(config-router)#ospf router-id 5.5.5.5	Configure OSPF router id
P3(config-router)#bfd all-interfaces	Enable BFD in all ospf interfaces
P3(config-router)#network 5.5.5.5/32 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P3(config-router)#network 15.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P3(config-router)#network 25.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P3(config-router)#network 35.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P3(config-router)#network 45.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P3(config-router)#network 56.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
P3(config-router)#commit	Commit the candidate configuration to the running configuration

VTEP3

VTEP3#configure terminal	Enter configuration terminal
VTEP3(config)#hostname VTEP3	Configure hostname
VTEP3(config)#mac vrf vrf1	Configure MAC vrf vrf1
VTEP3(config-vrf)#rd 6.6.6.6:11	Configure RD for vrf1
VTEP3(config-vrf)#route-target both 10.10.10.10:100	Configure RT for vrf1

VTEP3(config-vrf)#mac vrf vrf2	Configure MAC vrf vrf2
VTEP3(config-vrf)#rd 6.6.6.6:12	Configure RD for vrf2
VTEP3(config-vrf)#route-target both 10.10.10.10:101	Configure RT for vrf2
VTEP3(config-vrf)#route-target both 10.10.10.10:102	Configure RT for vrf2
VTEP3(config-vrf)#bfd interval 3 minrx 3 multiplier 3	Enable EVPN VXLAN multihoming
VTEP3(config)#evpn vxlan multihoming enable	Enable statistics on VXLAN tunnel interface
VTEP3(config)#nvo vxlan enable	Enable VXLAN
VTEP3(config)#qos enable	Enable qos
VTEP3(config)#qos statistics	Enable qos statistics
VTEP3(config)#bridge 1 protocol ieee vlan-bridge	Configure IEEE VLAN bridge
VTEP3(config)#no bridge 1 spanning-tree enable	Disable spanning tree in bridge 1 globally
VTEP3(config)#no igmp snooping	Disable igmp snooping messages globally
VTEP3(config)#vlan database	Enter into the VLAN database
VTEP3(config-vlan)#vlan 10-200 bridge 1 state enable	Configure VLAN 10-200 and associate with bridge 1
VTEP3(config-vlan)#interface lo	Enter lo interface mode
VTEP3(config-if)#ip address 6.6.6.6/32 secondary	Assign loopback IP
VTEP3(config-if)#interface vlan1.80	Specify interface VLAN1.80 to be configured.
VTEP3(config-if)#ip address 56.1.1.2/24	Assign IP address
VTEP3(config-if)#ip ospf cost 1	Change OSPF cost of the link
VTEP3(config-if)#interface vlan1.90	Specify interface VLAN1.90 to be configured.
VTEP3(config-if)#ip address 46.1.1.2/24	Assign IP address
VTEP3(config-if)#ip ospf cost 1	Change OSPF cost of the link
VTEP3(config-if)#interface xe3	Enter interface mode
VTEP3(config-if)#switchport	Set the interface as Layer2 port
VTEP3(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree
VTEP3(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
VTEP3(config-if)#switchport trunk allowed vlan add 90,99	Enable VLAN's allowed on this interface.
VTEP3(config-if)#switchport trunk native vlan 99	Configure native VLAN
VTEP3(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP3(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
VTEP3(config-if)#interface xe15	Enter interface mode
VTEP3(config-if)#switchport	Set the interface as Layer2 port
VTEP3(config-if)#bridge-group 1 spanning-tree disable	Associate the interface with bridge group 1 and disable spanning tree

VTEP3(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode.
VTEP3(config-if)#switchport trunk allowed vlan add 80,89	Enable VLAN's allowed on this interface.
VTEP3(config-if)#switchport trunk native vlan 89	Configure native VLAN
VTEP3(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP3(config-if)#mtu 1600	Change interface mtu value
VTEP3(config-if)#spanning-tree edgeport	Set the port as an edge-port to enable rapid transitions
VTEP3(config-if)#interface xe45	Enter interface mode
VTEP3(config-if)#switchport	Set the interface as Layer2 port
VTEP3(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP3(config-if)#interface xe46	Enter interface mode
VTEP3(config-if)#switchport	Set the interface as Layer2 port
VTEP3(config-if)#load-interval 30	Configure load period in multiple of 30 seconds
VTEP3(config-if)#router ospf 1	Enter OSPF configuration mode
VTEP3(config-router)#ospf router-id 6.6.6.6	Configure OSPF router id
VTEP3(config-router)#bfd all-interfaces	Enable BFD in all OSPF interfaces
VTEP3(config-router)#network 6.6.6.6/32 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
VTEP3(config-router)#network 46.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
VTEP3(config-router)#network 56.1.1.0/24 area 0.0.0.0	Define the Network on which OSPF runs and associate the area ID (area 0) with the interface.
VTEP3(config-router)#router bgp 100	Enter address-family IPv4 unicast mode
VTEP3(config-router)#address-family ipv4 unicast	Add the lo network to BGP route
VTEP3(config-router)#network 6.6.6.6/32	Activate neighbors
VTEP3(config-router-af)#neighbor 1.1.1.1 activate	Activate neigbors
VTEP3(config-router-af)#neighbor 2.2.2.2 activate	Exit address-family mode.
VTEP3(config-router-af)#exit-address-family	Enter Router BGP mode and define the AS number 100.
VTEP3(config-router)#neighbor 1.1.1.1 remote-as 100	Configure BGP remote-as 100 with neighbor IP
VTEP3(config-router)#neighbor 1.1.1.1 update-source lo	Define BGP neighbors, to update the source routes with lo
VTEP3(config-router)#neighbor 2.2.2.2 remote-as 100	Configure BGP remote-as 100 with neighbor IP
VTEP3(config-router)#neighbor 2.2.2.2 update-source lo	Configure advertisement-interval as 0 for fast convergence for VTEP2
(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure advertisement-interval as 0 for fast convergence for VTEP1
(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Define BGP neighbors, to update the source routes with lo
VTEP3(config-router)#address-family l2vpn evpn	Enter in to BGP L2VPN EVPN address-family

`VTEP3(config-router-af)#neighbor 1.1.1.1 activate`	Activate neigbors
`VTEP3(config-router-af)#neighbor 2.2.2.2 activate`	Activate neigbors
`VTEP3(config-router-af)#exit-address-family`	Exit from BGP L2VPN EVPN address-family
`VTEP3(config-router)#nvo vxlan vtep-ip-global 6.6.6.6`	Configure VXLAN global IP
`VTEP3(config)#nvo vxlan id 1 ingress-replication inner-vid-disabled`	Create VNID 1
`VTEP3(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1`	Associate VNID with EVPN and vrf1
`VTEP3(config-nvo)#nvo vxlan id 1000 ingress-replication inner-vid-disabled`	Create VNID 1000
`VTEP3(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf2`	Associate VNID with EVPN and vrf2
`VTEP3(config-nvo)#nvo vxlan access-if port-vlan xe45 3001`	Create VXLAN access port port-VLAN
`VTEP3(config-nvo-acc-if)#map vnid 1`	Map the VNID to access-if
`VTEP3(config-nvo-acc-if)#nvo vxlan access-if port xe46`	Create VXLAN access port
`VTEP3(config-nvo-acc-if)#map vnid 1000`	Map the VNID to access-if
`VTEP3(config-nvo-acc-if)#commit`	Commit the candidate configuration to the running configuration

SWITCH

`SWITCH#configure terminal`	Enter configuration terminal
`SWITCH#(config)# bridge 1 protocol rstp vlan-bridge`	Configure rstp vlan bridge
`SWITCH#(config-vlan)#interface po1`	Enter interface mode
`SWITCH#(config-if)#switchport`	Set the interface as Layer2 port
`SWITCH#(config-if)#load-interval 30`	Configure load period in multiple of 30 seconds
`SWITCH#(config-if)#interface ce9/1`	Enter interface mode
`SWITCH#(config-if)#channel-group 1 mode active`	Map this interface to po1
`SWITCH#(config-if)#interface ce9/2`	Enter interface mode
`SWITCH#(config-if)#channel-group 1 mode active`	Map this interface to po1
`SWITCH#(config-if)#interface ce10/1`	Enter interface mode
`SWITCH#(config-if)#channel-group 1 mode active`	Map this interface to po1
`SWITCH#(config-if)#interface ce10/2`	Enter interface mode
`SWITCH#(config-if)#channel-group 1 mode active`	Map this interface to po1
`SWITCH#(config-if)#interface ce31/1`	Set the port as an edge-port to enable rapid transitions
`SWITCH#(config-if)#switchport`	Enter interface mode

SWITCH#(config-if)#bridge-group 1 spanning-tree disable	Set the interface as Layer2 port
SWITCH#(config-if)#switchport mode trunk	Associate the interface with bridge group 1 and disable spanning tree
SWITCH#(config-if)#switchport trunk allowed vlan all	Set the switching characteristics of this interface to trunk mode.
SWITCH#(config-if)#load-interval 30	Configure native VLAN
SWITCH#(config-if)#commit	Commit the candidate configuration to the running configuration

Validation

```
VTEP1#sh ip ospf neighbor

Total number of full neighbors: 3
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
2.2.2.2           1   Full/DR         00:00:30    12.1.1.2        vlan1.10              0
3.3.3.3           1   Full/DR         00:00:31    13.1.1.2        vlan1.20              0
4.4.4.4           1   Full/DR         00:00:33    14.1.1.2        vlan1.30              0


VTEP2#sh ip ospf neighbor

Total number of full neighbors: 3
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
1.1.1.1           1   Full/Backup     00:00:38    12.1.1.1        vlan1.10
0
3.3.3.3           1   Full/DR         00:00:39    23.1.1.2        vlan1.40              0
4.4.4.4           1   Full/DR         00:00:39    24.1.1.2        vlan1.50              0
P1#sh ip ospf neighbor

Total number of full neighbors: 4
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
1.1.1.1           1   Full/Backup     00:00:39    13.1.1.1        vlan1.20
0
2.2.2.2           1   Full/Backup     00:00:38    23.1.1.1        vlan1.40
0
4.4.4.4           1   Full/DR         00:00:40    34.1.1.2        vlan1.100             0
5.5.5.5           1   Full/DR         00:00:36    35.1.1.2        vlan1.60              0

P2#sh ip ospf neighbor

Total number of full neighbors: 5
OSPF process 1 VRF(default):
```

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
1.1.1.1           1   Full/Backup     00:00:30    14.1.1.1        vlan1.30
0
2.2.2.2           1   Full/Backup     00:00:38    24.1.1.1        vlan1.50
0
3.3.3.3           1   Full/Backup     00:00:33    34.1.1.1        vlan1.100
0
5.5.5.5           1   Full/DR         00:00:30    45.1.1.2        vlan1.70                0
6.6.6.6           1   Full/DR         00:00:34    46.1.1.2        vlan1.90                0
P2#
P3#sh ip ospf neighbor

Total number of full neighbors: 3
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
3.3.3.3           1   Full/Backup     00:00:34    35.1.1.1        vlan1.60
0
4.4.4.4           1   Full/Backup     00:00:34    45.1.1.1        vlan1.70
0
6.6.6.6           1   Full/DR         00:00:33    56.1.1.2        vlan1.80                0
P3#

VTEP1#sh bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 4
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
2.2.2.2                 4   100   23        22      4      0      0  00:07:34
2      0      0      2      0      0
6.6.6.6                 4   100   21        22      4      0      0  00:07:34
2      0      0      2      0      0

Total number of neighbors 2

Total number of Established sessions 2
VTEP1#
VTEP2#sh bgp l2vpn evpn summary
BGP router identifier 2.2.2.2, local AS number 100
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4   100   22        24      5      0      0  00:07:41
2      0      0      2      0      0
6.6.6.6                 4   100   24        27      5      0      0  00:08:51
2      0      0      2      0      0
```

```
Total number of neighbors 2

Total number of Established sessions 2
VTEP2#
VTEP3#sh bgp l2vpn evpn summary
BGP router identifier 6.6.6.6, local AS number 100
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
1.1.1.1                 4   100   23        21      5     0     0  00:07:44
2      0      0      2      0      0
2.2.2.2                 4   100   26        24      5     0     0  00:08:54
2      0      0      2      0      0

Total number of neighbors 2

Total number of Established sessions 2
VTEP3#

VTEP1#show nvo vxlan mac-table
================================================================================
================================================================
                                                          VXLAN MAC Entries
================================================================================
================================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc

--------------------------------------------------------------------------------
--------------------------------------------------------------------

1         po1        1000   ----         0000.1111.1111 00:aa:22:33:44:55:66:00:00:00
Static Local         -------             -------
1         po1        1000   ----         0000.1111.1112 00:aa:22:33:44:55:66:00:00:00
Static Local         -------             -------
1         po1        1000   ----         0000.1111.1113 00:aa:22:33:44:55:66:00:00:00
Static Local         -------             -------
1         po1        1000   ----         0000.1111.1114 00:aa:22:33:44:55:66:00:00:00
Static Local         -------             -------
1         po1        1000   ----         0000.1111.1115 00:aa:22:33:44:55:66:00:00:00
Static Local         -------             -------
1         po1        1000   ----         a82b.b57c.4470 00:aa:22:33:44:55:66:00:00:00
Dynamic Local        -------             -------
1000      ----       ----   ----         a82b.b57c.4476 00:aa:22:33:44:55:66:00:00:00
Dynamic Remote       -------             -------

Total number of entries are : 7

VTEP1#
VTEP2#sh nvo vxlan mac-table
```

```
==========================================================================================
================================================================
                                                           VXLAN MAC Entries
==========================================================================================
================================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc
__________________________________________________________________________________________
________________________________________________________________


1         ----       ----   ----         0000.1111.1111 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         0000.1111.1112 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         0000.1111.1113 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         0000.1111.1114 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         0000.1111.1115 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         a82b.b57c.4470 00:aa:22:33:44:55:66:00:00:00
Dynamic Remote       -------             -------
1000      po1        2000   ----         a82b.b57c.4476 00:aa:22:33:44:55:66:00:00:00
Dynamic Local        -------             -------

Total number of entries are : 7

VTEP2#
VTEP3#sh nvo vxlan mac-table
==========================================================================================
================================================================
                                                           VXLAN MAC Entries
==========================================================================================
================================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc
__________________________________________________________________________________________
________________________________________________________________


1         ----       ----   ----         0000.1111.1111 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         0000.1111.1112 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         0000.1111.1113 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         0000.1111.1114 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         0000.1111.1115 00:aa:22:33:44:55:66:00:00:00
Static Remote        -------             -------
1         ----       ----   ----         a82b.b57c.4470 00:aa:22:33:44:55:66:00:00:00
Dynamic Remote       -------             -------
1000      ----       ----   ----         a82b.b57c.4476 00:aa:22:33:44:55:66:00:00:00
Dynamic Remote       -------             -------

Total number of entries are : 7
```

```
VTEP3#
VTEP1#show nvo vxlan access-if brief

                 Inner                     Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
xe46       ---   ---   500000   1         up      up
po1        2000  ---   500001   1000      up      up
po1        2001  ---   500002   1000      up      up


Total number of entries are 2
VTEP1#
VTEP2#show nvo vxlan access-if brief

                 Inner                     Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
xe47       ---   ---   500000   1         up      up
po1        2001  ---   500001   1000      up      up
po1        2000  ---   500002   1000      up      up

Total number of entries are 3
VTEP2#
VTEP3#show nvo vxlan access-if brief

                 Inner                     Admin   Link
Interface  Vlan  vlan  Ifindex  Vnid      status  status
-----------------------------------------------------------
xe45       3001  ---   500000   1         up      up
xe46       ---   ---   500001   1000      up      up

Total number of entries are 2
VTEP3#
VTEP1#sh nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
1.1.1.1          6.6.6.6          Installed        00:03:59        00:03:59
1.1.1.1          2.2.2.2          Installed        00:03:59        00:03:59

Total number of entries are 2
VTEP1#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged
```

```
VNID     VNI-Name     VNI-Type Type Interface ESI                                VLAN DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
1        ----         L2       NW   ----      ----                                ---- --
--      1.1.1.1          6.6.6.6
1        ----         L2       NW   ----      ----                                ---- --
--      1.1.1.1          2.2.2.2
1000     ----       --      AC   po1      --- 00:aa:22:33:44:55:66:00:00:00 2000 DF
----              ----
1000     ----       --      AC   po1      --- 00:aa:22:33:44:55:66:00:00:00 2001 DF
----              ----
1000     ----         L2       NW   ----      ----                                ---- -
---      1.1.1.1          6.6.6.6
1000     ----         L2       NW   ----      ----                                ---- -
---      1.1.1.1          2.2.2.2
1000     ----         --       AC   xe46      --- Single Homed Port ---           ---- -
---      ----             ----

Total number of entries are 10
VTEP1#

VTEP2#sh nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
2.2.2.2          1.1.1.1          Installed       00:03:59        00:03:59
2.2.2.2          6.6.6.6          Installed       00:05:09        00:05:09

Total number of entries are 2
VTEP2#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                                VLAN DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
1        ----         L2       NW   ----      ----                                ---- --
--      2.2.2.2          1.1.1.1
1        ----         L2       NW   ----      ----                                ---- --
--      2.2.2.2          6.6.6.6
1        ----         --       AC   xe47      --- Single Homed Port ---           ---- --
--      ----             ----
1000     ----         L2       NW   ----      ----                                ---- --
--      2.2.2.2          1.1.1.1
1000     ----         L2       NW   ----      ----                                ---- --
--      2.2.2.2          6.6.6.6
1000     ----        --       AC   po1       --- 00:aa:22:33:44:55:66:00:00:00 2001 NON-
DF      ----             ----
```

```
1000     ----         --       AC   po1       --- 00:aa:22:33:44:55:66:00:00:00 2000 NON-
DF       ----              ----

Total number of entries are 11
VTEP2#

VTEP3#sh nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
6.6.6.6          1.1.1.1          Installed       00:03:58        00:03:58
6.6.6.6          2.2.2.2          Installed       00:05:08        00:04:03

Total number of entries are 2
VTEP3#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                          VLAN DF-Status
Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
1        ----         L2       NW   ----      ----                         ---- ----
6.6.6.6          1.1.1.1
1        ----         L2       NW   ----      ----                         ---- ----
6.6.6.6          2.2.2.2
1        ----         --       AC   xe45      --- Single Homed port ---    3001 ----
----             ----
1000     ----         L2       NW   ----      ----                         ---- ----
6.6.6.6          1.1.1.1
1000     ----         L2       NW   ----      ----                         ---- ----
6.6.6.6          2.2.2.2
1000     ----         --       AC   xe46      --- Single Homed Port ---     ---- ----
----             ----

Total number of entries are 10
VTEP3#
```

VxLAN - Artificial Intelligence Networking

CHAPTER 1 ECN and PFC Support for Lossless VxLAN Transport

This document describes the functionalities of ECN and PFC over VxLAN network for Artificial intelligence applications and other workloads that unitize RDMA over Converged Ethernet (RoCEv2) transport protocol.

Overview

Artificial intelligence and machine learning (AI/ML) workloads in modern data centers demand ultra-high bandwidth and low-latency network fabrics, typically utilizing RDMA over Converged Ethernet (RoCE) v2 transports to enable efficient Graphics Processing Unit (GPU)-to-GPU communication for distributed training. Because these applications are extremely sensitive to packet loss, even small amounts of dropped traffic can cause retransmissions, add latency, and significantly delay training completion.

To address this, data center fabrics are designed to be lossless, with:

- Priority Flow Control (PFC) - ensuring that traffic in critical queues is paused during congestion rather than dropped, while other traffic continues unaffected, and
- Explicit Congestion Notification (ECN) - that allows two endpoints in a network to signal congestion without dropping packets.

The combination of Explicit Congestion Notification (ECN) based congestion signaling and PFC is fundamental to achieving predictable, efficient, and scalable network performance required for AI/ML workload in large-scale data centers.

PFC

When the ingress buffer for a lossless priority on a port reaches its limit, the device immediately transmits an XOFF (pause) frame to the link partner. This action halts traffic transmission for the Classes of Service (CoS) mapped to that congested priority, preventing packet loss while in-flight packets are buffered.

OcNOS facilitates this process by supporting CoS-to-queue and DSCP-to-queue mapping profiles, which route incoming traffic to one of up to eight available lossless priority queues..

ECN - Marking

When an egress queue exceeds its configured buffering threshold, ECN manages traffic flow by marking packets with congestion bits instead of dropping them. This marking strategy helps reduce the occurrence of PFC pauses, thereby improving overall network throughput and latency.

Combining ECN and PFC

ECN and PFC work together to deliver lossless and efficient transport in the VxLAN network. While PFC provides a mechanism to pause incoming traffic when the buffering limit is reached, ECN complements it by signaling congestion before the buffers become full. Marking packets at early stages of congestion, allows endpoints to reduce their sending rate gracefully, rather than abruptly pauses traffic. This proactive approach lowers the frequency of PFC pause events, helping maintain higher throughput and reducing latency.

Feature Characteristics

OcNOS supports the following for ECN to work in L3 VxLAN transport:

- End-to-end ECN support, including:

    - ECN over VXLAN is supported for both Layer 2 and Layer 3 (IRB) VXLAN forwarding. ECN markings are propagated from ingress to egress - specifically, from the inner IP header to the outer IP header during VXLAN encapsulation, and back from the outer to the inner header upon decapsulation.
    - VTEP node that encapsulate packet with ECN value will copy the Inner IP ECN bits to the outer IP header.
    - Transit nodes in the middle of the tunnel carries the ECN value in the outer IP header.
    - VTEP that decapsulates the packet received from VxLAN and transfers the ECN marking to the customer packet.
- Triggers ECN marking or PFC pause frames based on queue occupancy when congestion occurs either in encapsulation or transit or decapsulation.

OcNOS supports the following for PFC to work in L2 and L3 VxLAN transport:

- CoS to queue and DSCP to queue mapping profiles can be configured using policy map. The mapping can be applied to tagged or untagged access ports.
- For L2 forwarded packet, queue is derived from the Class of Service (CoS) value.
- For L3 forwarded packet, queue is derived from the Differentiated Services Code Point (DSCP) value.

Benefits

This feature delivers lossless Ethernet fabric required for transport of AI/ML and other workloads that utilize RoCEv2 transport protocol.

- It enables multi-tenant and multi-cluster isolation using EVPN-VXLAN while successfully maintaining the lossless transport essential for RDMA (RoCEv2).
- For operators utilizing Trident3 (TR3) platforms and Trident4 (TR4) platforms, it guarantees predictable congestion management across large-scale fabrics.

Prerequisites

Ensure the IP underlay provides reachability between all VTEP loopback addresses. This connectivity, established via routing protocols like OSPF or BGP (both present in the base configurations), is fundamental for the VxLAN tunnels where PFC and ECN will function.

Configuration

This section describes the configuration procedures for Priority Flow Control (PFC) and Explicit Congestion Notification (ECN) over VxLAN. It involves applying PFC on underlay and access links, and enabling ECN within the Lossless queuing policy map applied to the hardware.

Topology

The following topology shows

- two Spine nodes (Spine1 - `Eth63/1, Eth1/1, Eth11/1` and Spine 2 - `cd27/1, cd3/1, cd9/1`) responsible for interconnectivity between VTEPs,
- three Leaf nodes (VTEP1- `cd3/1, cd4/1, xe0, cd8/1`, VTEP2- `ce1, xe34, ce31` and VTEP3 - `ce2, ce3, xe4/4`) that encapsulate and decapsulate traffic. They connect North to the Spines and South to the hosts/ access switch.

- A Switch (`Eth29/1, Eth1/1, Eth3/1`) that manages access layer connectivity for one multihoming host and three single homing hosts.

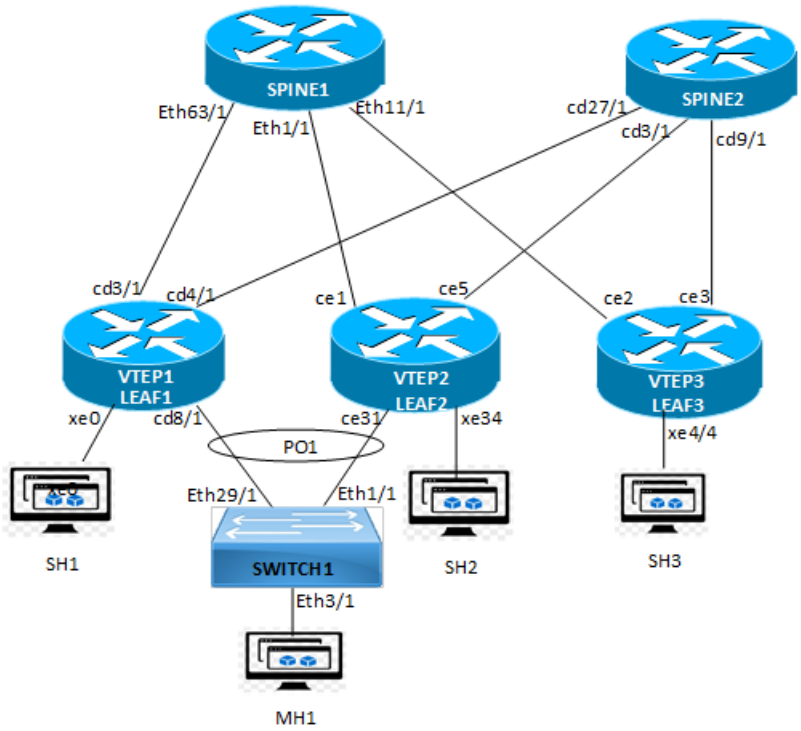

Figure 1-8: PFC and ECN over VxLAN

Configuring PFC and ECN over VxLAN

Configuring PFC and ECN over VxLAN involves the following steps:

1. Global ECN and Lossless Queuing Configuration on all devices.

   Note: The `#qos remark dscp` command is mandatory in the TR4 board, specifically to ensure ECN functionality operates correctly and that ECN-related counters increment as expected.

2. PFC configuration on Underlay and Access Interfaces.

   1. PFC on Spine Underlay Links (Layer 3 Routed)
   2. PFC on Leaf access switchport (VTEPs).
   3. PFC on Switch (Single home and Multihome).

3. Load-balancing configuration on all devices.

4. Configure a router to distribute connected routes within networks.

   Note: The below configuration procedure include setups for both OSPF and BGP to establish underlay reachability for VTEP loopbacks. However, only one active underlay protocol is required. Choose either OSPF or BGP based on your network requirements. Enable the chosen protocol (e.g., remove `shutdown` from the OSPF process or ensure BGP neighbors are correctly configured and active) and ensure the other protocol is disabled or unconfigured to avoid routing conflicts.

It includes the following sections:

- Configuration on VTEP1, VTEP2 and VTEP3

Configuration on VTEP1, VTEP2 and VTEP3

1. Enable the global QoS

```
qos enable
qos remark dscp
```

2. Configured PFC queuing policy for eight queues (q0–q7), all configured as priority lossless queues.

```
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
```

```
        lossless
        exit
       class type queuing default q4
        priority
        lossless
        exit
       class type queuing default q5
        priority
        lossless
        exit
       class type queuing default q6
        priority
        lossless
        exit
       class type queuing default q7
        priority
        lossless
        exit
      !
```

3. Configure the egress queuing policy map to define ECN behavior and specify priority and lossless attributes for all eight queues. Set each queue as lossless as required for RDMA or other PFC-enabled traffic. Enable Explicit Congestion Notification (ECN) marking based on WRED thresholds for green, yellow, and red color profiles on all lossless queues.

```
   policy-map type queuing default lossless_ecn_egress
    class type queuing default q0
     shape 10 gbps
     priority
     lossless
     random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
     exit
    class type queuing default q1
     shape 10 gbps
     priority
     lossless
     random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
     exit
    class type queuing default q2
     shape 10 gbps
     priority
     lossless
     random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
     exit
    class type queuing default q3
     shape 10 gbps
     priority
     lossless
     random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
```

```
        exit
       class type queuing default q4
        shape 10 gbps
        priority
        lossless
        random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
        exit
       class type queuing default q5
        shape 10 gbps
        priority
        lossless
        random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
        exit
       class type queuing default q6
        shape 10 gbps
        priority
        lossless
        random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
        exit
       class type queuing default q7
        shape 10 gbps
        priority
        lossless
        random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
        exit
      !
```

4. Configure PFC on the access and network ports of the fabric that is connected to hosts/servers and Spines respectively.

   On VTEP1

   Access Interfaces:

```
!
interface xe0
 description ***SH to SH ****
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 lldp-agent
  exit
!
interface po1
 description ****SW1-MH****
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
```

```
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 evpn multi-homed system-mac 0000.0000.3333
!
```

Network Interfaces:

```
interface cd3/1
 priority-flow-control mode auto
 ip address 30.1.15.1/24
 ipv6 address 3015::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!

interface cd4/1
 priority-flow-control mode auto
 ip address 30.1.16.1/24
 ipv6 address 3016::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
```

On VTEP2

Access Interfaces:

```
!
interface xe34
 description ***GPU-4 ***
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 lldp-agent
  exit
!!
interface po1
 description ****SW1-MH****
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 evpn multi-homed system-mac 0000.0000.3333
!
```

Network Interfaces:

```
interface ce1
 priority-flow-control mode auto
```

```
 ip address 10.1.17.1/24
 ipv6 address 1017::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit


interface ce5
 priority-flow-control mode auto
 ip address 10.1.18.1/24
 ipv6 address 1018::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
```

On VTEP3

Access Interfaces:

```
!
interface xe4/4
 description ***GPU-4 ***
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 lldp-agent
  exit
!
```

Network Interfaces:

```
interface ce2
 priority-flow-control mode auto


interface ce3
 priority-flow-control mode auto
```

5. Attach the configured ECN profile to the network ECMP interfaces on all the VTEPS/Spines in the fabric.

   On VTEP1

```
!
interface cd3/1
 priority-flow-control mode auto
 ip address 30.1.15.1/24
```

```
 ipv6 address 3015::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 service-policy type queuing output lossless_ecn_egress
 lldp-agent
  exit
!

interface cd4/1
 priority-flow-control mode auto
 ip address 30.1.16.1/24
 ipv6 address 3016::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 service-policy type queuing output lossless_ecn_egress
 lldp-agent
  exit
```

On VTEP2

```
interface ce1
 priority-flow-control mode auto
 ip address 10.1.17.1/24
 ipv6 address 1017::1/64
 service-policy type queuing output lossless_ecn_egress
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit


interface ce5
 priority-flow-control mode auto
 ip address 10.1.18.1/24
 ipv6 address 1018::1/64
 service-policy type queuing output lossless_ecn_egress
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
```

On VTEP2

```
Network Interface:
interface ce2
 priority-flow-control mode auto
```

```
 service-policy type queuing output lossless_ecn_egress

interface ce3
 priority-flow-control mode auto
 service-policy type queuing output lossless_ecn_egress
```

6. Configure load balancing using RTAG7 across multiple ECMP paths for both Layer 2 and Layer 3 traffic. Enable Layer 2 hashing based on source and destination MAC addresses, Ethertype, and VLAN. Configure IPv4 hashing to include Layer 4 source and destination ports and the protocol ID, which is required for VxLAN underlay routing. Enable support for up to 64 ECMP paths.

```
!
load-balance rtag7
load-balance rtag7 all-lag-if
load-balance rtag7 macro-flow
load-balance rtag7 l2 dest-mac src-mac ether-type vlan
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
rocev2-dest-qpairs
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port rocev2-dest-
qpairs next-hdr
!
maximum-paths 64
!
!
2 3 4 5 6
 load-interval 30
 trust dscp
!
```

7. Configure an IRB (Integrated Routing and Bridging) interface irb2 for each SMS VRF. By enabling evpn irb-if-forwarding anycast-gateway-mac, the interface functions as a distributed anycast gateway within an EVPN-VXLAN fabric, allowing all Leaf switches to advertise the same MAC address for this subnet. This setup provides tenant-specific routing isolation, seamless host mobility across the fabric, and distributed Layer 3 forwarding for devices in the network.

   On VTEP1

```
!
interface irb2
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 2.1.1.1/24 anycast
 ipv6 address 2::1/48 anycast
!
```

   On VTEP2

```
!
interface irb2
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 2.1.1.1/24 anycast
 ipv6 address 2::1/48 anycast
!
```

On VTEP3

```
!
interface irb2
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 2.1.1.1/24 anycast
 ipv6 address 2::1/48 anycast
```

8. Configure a OSPF router to distribute connected routes within networks in Area 0.0.0.0.

VTEP1

```
!
router ospf 100
 ospf router-id 3.3.3.3
 shutdown
 network 3.3.3.3/32 area 0.0.0.0
 network 10.1.15.0/24 area 0.0.0.0
 network 10.1.16.0/24 area 0.0.0.0
 network 20.1.15.0/24 area 0.0.0.0
 network 20.1.16.0/24 area 0.0.0.0
 network 30.1.15.0/24 area 0.0.0.0
 network 30.1.16.0/24 area 0.0.0.0
 network 103.1.1.0/24 area 0.0.0.1
 network 113.1.1.0/24 area 0.0.0.1
 network 173.1.1.0/24 area 0.0.0.1
 network 173.1.2.0/24 area 0.0.0.1
 network 173.1.3.0/24 area 0.0.0.1
 network 173.1.4.0/24 area 0.0.0.1
 network 173.1.5.0/24 area 0.0.0.1
!
router ipv6 ospf
 router-id 3.3.3.3
!
```

VTEP2

```
!
router ospf 100
 ospf router-id 4.4.4.4
 shutdown
 network 4.4.4.4/32 area 0.0.0.0
 network 10.1.17.0/24 area 0.0.0.0
 network 10.1.18.0/24 area 0.0.0.0
 network 104.1.1.0/24 area 0.0.0.1
 network 174.1.1.0/24 area 0.0.0.1
 network 174.1.2.0/24 area 0.0.0.1
 network 174.1.3.0/24 area 0.0.0.1
 network 174.1.4.0/24 area 0.0.0.1
 network 174.1.5.0/24 area 0.0.0.1
!
```

```
router ipv6 ospf
 router-id 4.4.4.4
!
```

9. Configure a BGP router.

VTEP1

```
!
router bgp 4294967203
 bgp router-id 3.3.3.3
 bgp bestpath as-path multipath-relax
 bgp log-neighbor-changes
 no bgp inbound-route-filter
 neighbor overlay peer-group
 neighbor overlay remote-as 4294967209
 neighbor overlay ebgp-multihop
 neighbor overlay update-source lo
 neighbor overlay advertisement-interval 0
 neighbor overlay fall-over bfd multihop
 neighbor underlay peer-group
 neighbor underlay remote-as 4294967209
 neighbor underlay authentication-key 0xb59db09d828b2528
 neighbor underlay as-origination-interval 1
 neighbor underlay advertisement-interval 0
 neighbor underlay fall-over bfd
 neighbor 7.7.7.7 peer-group overlay
 neighbor 8.8.8.8 peer-group overlay
 !
 bgp unnumbered-mode
 neighbor cd4/1 peergroup underlay
 neighbor cd3/1 peergroup underlay
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 10
 redistribute connected
 neighbor underlay activate
 !
 bgp v4-unnumbered-mode
 neighbor cd4/1 route-map HIG_MED in
 neighbor cd3/1 route-map HIG_MED in
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor overlay activate
 !
 bgp l2vpn-unnumbered-mode
 exit-l2vpn-unnumbered-mode
 !
```

```
exit-address-family
!
address-family ipv4 vrf l3_vrf_10
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_23
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_45
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_67
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_89
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_10
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_23
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_45
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_67
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_89
max-paths ebgp 50
redistribute connected
```

```
 exit-address-family
 !
 exit
```

VTEP2

```
!
router bgp 4294967204
 bgp router-id 4.4.4.4
 bgp bestpath as-path multipath-relax
 bgp log-neighbor-changes
 neighbor overlay peer-group
 neighbor overlay remote-as 4294967209
 neighbor overlay ebgp-multihop
 neighbor overlay update-source lo
 neighbor overlay advertisement-interval 0
 neighbor overlay fall-over bfd multihop
 neighbor underlay peer-group
 neighbor underlay remote-as 4294967209
 neighbor underlay authentication-key 0xb59db09d828b2528
 neighbor underlay as-origination-interval 1
 neighbor underlay advertisement-interval 0
 neighbor underlay fall-over bfd
 neighbor 7.7.7.7 peer-group overlay
 neighbor 8.8.8.8 peer-group overlay
 !
 bgp unnumbered-mode
 neighbor ce1 peergroup underlay
 neighbor ce5 peergroup underlay
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 10
 redistribute connected
 neighbor underlay activate
 !
 bgp v4-unnumbered-mode
 neighbor ce1 route-map HIG_MED in
 neighbor ce5 route-map HIG_MED in
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor overlay activate
 !
 bgp l2vpn-unnumbered-mode
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
```

VTEP3

```
!
router bgp 4294967205
 bgp router-id 5.5.5.5
 bgp bestpath as-path multipath-relax
 bgp log-neighbor-changes
 neighbor overlay peer-group
 neighbor overlay remote-as 4294967209
 neighbor overlay ebgp-multihop
 neighbor overlay update-source lo
 neighbor overlay advertisement-interval 0
 neighbor overlay fall-over bfd multihop
 neighbor underlay peer-group
 neighbor underlay remote-as 4294967209
 neighbor underlay authentication-key 0xb59db09d828b2528
 neighbor underlay as-origination-interval 1
 neighbor underlay advertisement-interval 0
 neighbor underlay fall-over bfd
 neighbor 7.7.7.7 peer-group overlay
 neighbor 8.8.8.8 peer-group overlay
 !
 bgp unnumbered-mode
 neighbor ce3 peergroup underlay
 neighbor ce2 peergroup underlay
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 10
 redistribute connected
 neighbor underlay activate
 !
 bgp v4-unnumbered-mode
 neighbor ce3 route-map HIG_MED in
 neighbor ce2 route-map HIG_MED in
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor overlay activate
 !
 bgp l2vpn-unnumbered-mode
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
```

10. Configure VxLAN access port to associate with VXLAN VNID for encapsulation or forward in the corresponding VXLAN overlay tunnel.

    On VTEP1

```
!
nvo vxlan access-if port-vlan xe0 2
 description ***VTEP3-SH-TENENT-2***
 map vnid 50002
!
nvo vxlan access-if port-vlan xe0 3
 description ***VTEP3-SH-TENENT-3***
 map vnid 50003
!
nvo vxlan access-if port-vlan xe0 4
 description ***VTEP3-SH-TENENT-4***
 map vnid 50004
!
nvo vxlan access-if port-vlan xe0 10
 description ***VTEP3-SH-TENENT-10***
 map vnid 50010
!
nvo vxlan access-if port-vlan po1 6
 description ***VTEP3-MH-TENENT-6***
 map vnid 50006
!
nvo vxlan access-if port-vlan po1 7
 description ***VTEP3-MH-TENENT-7***
 map vnid 50007
!
nvo vxlan access-if port-vlan po1 8
 description ***VTEP3-MH-TENENT-8***
 map vnid 50008
!
```

On VTEP2

```
!
nvo vxlan access-if port-vlan xe34 2
 description ***VTEP4-SH-TENENT-2***
 map vnid 50002
!
nvo vxlan access-if port-vlan xe34 3
 description ***VTEP4-SH-TENENT-3***
 map vnid 50003
!
nvo vxlan access-if port-vlan xe34 5
 description ***VTEP4-SH-TENENT-5***
 map vnid 50005
!
nvo vxlan access-if port-vlan xe34 10
 description ***VTEP4-SH-TENENT-10***
 map vnid 50010
!
nvo vxlan access-if port-vlan po1 6
 description ***VTEP4-MH-TENENT-6***
 map vnid 50006
```

```
!
nvo vxlan access-if port-vlan po1 7
 description ***VTEP4-MH-TENENT-7***
 map vnid 50007
!
nvo vxlan access-if port-vlan po1 8
 description ***VTEP4-MH-TENENT-8***
 map vnid 50008
!
```

On VTEP3

```
!
nvo vxlan access-if port-vlan xe4/4 6
 description ***VTEP5-SH-TENENT-6***
 map vnid 50006
!
nvo vxlan access-if port-vlan xe4/4 7
 description ***VTEP5-SH-TENENT-7***
 map vnid 50007
!
nvo vxlan access-if port-vlan xe4/4 9
 description ***VTEP5-SH-TENENT-9***
 map vnid 50009
!
```

Configuration on Spine 1 and Spine 2

1. Enable the global QoS

```
qos enable
qos remark dscp
```

2. Configured PFC queuing policy for eight queues (q0–q7), all configured as priority lossless queues.

```
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
```

```
        exit
       class type queuing default q4
        priority
        lossless
        exit
       class type queuing default q5
        priority
        lossless
        exit
       class type queuing default q6
        priority
        lossless
        exit
       class type queuing default q7
        priority
        lossless
        exit
      !
```

3. Configure the egress queuing policy map to define ECN behavior and specify priority and lossless attributes for all eight queues. Set each queue as lossless as required for RDMA or other PFC-enabled traffic. Enable Explicit Congestion Notification (ECN) marking based on WRED thresholds for green, yellow, and red color profiles on all lossless queues.

```
   policy-map type queuing default lossless_ecn_egress
    class type queuing default q0
     shape 10 gbps
     priority
     lossless
     random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
     exit
    class type queuing default q1
     shape 10 gbps
     priority
     lossless
     random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
     exit
    class type queuing default q2
     shape 10 gbps
     priority
     lossless
     random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
     exit
    class type queuing default q3
     shape 10 gbps
     priority
     lossless
     random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
     exit
```

```
       class type queuing default q4
        shape 10 gbps
        priority
        lossless
        random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
        exit
       class type queuing default q5
        shape 10 gbps
        priority
        lossless
        random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
        exit
       class type queuing default q6
        shape 10 gbps
        priority
        lossless
        random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
        exit
       class type queuing default q7
        shape 10 gbps
        priority
        lossless
        random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300
max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
        exit
      !
```

4. PFC is configured on the physical interfaces to handle Lossless Ethernet frames. Configure the physical interfaces with PFC enabled on priorities 0–7, IP address, and 30-second statistics interval.

   On Spine1

```
      !
      interface ethernet1/1
       priority-flow-control mode on
       priority-flow-control advertise-local-config
       priority-flow-control enable priority 0 1 2 3 4 5 6 7
       load-interval 30
       ip address 10.1.17.2/24
       ipv6 address 1017::2/64
       ip ospf cost 100
       ipv6 ospf cost 100 instance-id 0
       ipv6 router ospf area 0.0.0.0 instance-id 0
       lldp-agent
        exit
      !
      interface ethernet11/1
       priority-flow-control mode on
       priority-flow-control advertise-local-config
       priority-flow-control enable priority 0 1 2 3 4 5 6 7
```

```
 load-interval 30
 ip address 10.1.19.2/24
 ipv6 address 1019::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet63/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 30.1.15.2/24
 ipv6 address 3015::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
```

On Spine2

```
!
interface cd3/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 10.1.18.2/24
 ipv6 address 1018::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd27/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 30.1.16.2/24
 ipv6 address 3016::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd9/1
 priority-flow-control mode on
```

```
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 10.1.20.2/24
 ipv6 address 1020::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
```

5. Configure load balancing using RTAG7 across multiple ECMP paths for both Layer 2 and Layer 3 traffic. Enable Layer 2 hashing based on source and destination MAC addresses, Ethertype, and VLAN. Configure IPv4 hashing to include Layer 4 source and destination ports and the protocol ID, which is required for VxLAN underlay routing. Enable support for up to 64 ECMP paths.

```
!
load-balance rtag7
load-balance rtag7 all-lag-if
load-balance rtag7 macro-flow
load-balance rtag7 l2 dest-mac src-mac ether-type vlan
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
rocev2-dest-qpairs
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port rocev2-dest-
qpairs next-hdr
!
maximum-paths 64
```

6. Configure a OSPF router 1 to distribute connected routes within networks in Area 0.0.0.0

   On Spine 1

```
!
router ospf 100
 ospf router-id 7.7.7.7
 shutdown
 network 7.7.7.7/32 area 0.0.0.0
 network 10.1.11.0/24 area 0.0.0.0
 network 10.1.13.0/24 area 0.0.0.0
 network 10.1.15.0/24 area 0.0.0.0
 network 10.1.17.0/24 area 0.0.0.0
 network 10.1.19.0/24 area 0.0.0.0
 network 20.1.11.0/24 area 0.0.0.0
 network 20.1.13.0/24 area 0.0.0.0
 network 20.1.15.0/24 area 0.0.0.0
 network 20.1.17.0/24 area 0.0.0.0
 network 30.1.11.0/24 area 0.0.0.0
 network 30.1.13.0/24 area 0.0.0.0
 network 30.1.15.0/24 area 0.0.0.0
 network 30.1.17.0/24 area 0.0.0.0
 network 40.1.13.0/24 area 0.0.0.0
 network 50.1.13.0/24 area 0.0.0.0
 network 60.1.13.0/24 area 0.0.0.0
!
```

```
router ipv6 ospf
 router-id 7.7.7.7
!
```

On Spine 2

```
!
router ospf 100
 ospf router-id 8.8.8.8
 shutdown
 network 8.8.8.8/32 area 0.0.0.0
 network 10.1.12.0/24 area 0.0.0.0
 network 10.1.14.0/24 area 0.0.0.0
 network 10.1.16.0/24 area 0.0.0.0
 network 10.1.18.0/24 area 0.0.0.0
 network 10.1.20.0/24 area 0.0.0.0
 network 20.1.12.0/24 area 0.0.0.0
 network 20.1.14.0/24 area 0.0.0.0
 network 20.1.16.0/24 area 0.0.0.0
 network 20.1.18.0/24 area 0.0.0.0
 network 30.1.12.0/24 area 0.0.0.0
 network 30.1.14.0/24 area 0.0.0.0
 network 30.1.16.0/24 area 0.0.0.0
 network 30.1.18.0/24 area 0.0.0.0
!
router ipv6 ospf
 router-id 8.8.8.8
!
```

7. Create a local ERSPAN-sender monitor session with a source port and destination port to monitor session in both directions. Below is the sample configuration commands.

```
!
monitor session 1 type local
 source interface cd3/1 both
 destination interface xe34
 no shut
!
```

Configuration on Switch

This section describes the configuration of PFC on MH Access Switch.

1. Enable the global QoS

```
qos enable
qos remark dscp
```

2. Configure PFC on the L2 trunk interfaces connecting to the Multihomed VTEPs (Leaf1/Leaf2).

```
!
interface ethernet3/1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
```

```
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 lldp-agent
  exit
!
!
interface po1
 description ****VTEP1-VTEP2-MH****
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
 priority-flow-control mode auto
!
```

3. Configure load balancing using RTAG7 across multiple ECMP paths for both Layer 2 and Layer 3 traffic. Enable Layer 2 hashing based on source and destination MAC addresses, Ethertype, and VLAN. Configure IPv4 hashing to include Layer 4 source and destination ports and the protocol ID, which is required for VxLAN underlay routing. Enable support for up to 64 ECMP paths.

```
!
load-balance rtag7
load-balance rtag7 all-lag-if
load-balance rtag7 macro-flow
load-balance rtag7 l2 dest-mac src-mac ether-type vlan
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
rocev2-dest-qpairs
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port rocev2-dest-
qpairs next-hdr
!
```

Sample Running Configurations

VTEP1

```
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
!
service password-encryption
!
logging console 5
logging monitor 4
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps ospf
snmp-server enable traps bgp
!
bgp extended-asn-cap
```

```
!
load-balance rtag7
load-balance rtag7 all-lag-if
load-balance rtag7 macro-flow
load-balance rtag7 l2 dest-mac src-mac ether-type vlan
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id rocev2-
dest-qpairs
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port rocev2-dest-qpairs
next-hdr
!
bfd interval 100 minrx 100 multiplier 3
!
qos enable
!
hostname Leaf3
port cd3 breakout 1X100g serdes 25g
port cd8 breakout 1X100g serdes 25g
port cd4 breakout 1X100g serdes 25g
default-interface type eth-routed mtu 9200
default-interface type eth-switchport mtu 9200
default-interface type svi mtu 9200
default-interface type irb mtu 9200
default-interface type lag mtu 9200
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
default-interface load-interval 30
bridge 1 protocol ieee vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1
priority-flow-control enable bridge 1
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management

lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp tlv-select ieee-8021-org-specific data-center-bridging
set lldp msg-tx-hold 1
set lldp timer msg-tx-interval 5
set lldp tx-fast-init 1
```

```
lldp notification-interval 5
fault-management enable
maximum-paths 64
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
```

```
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
vlan database
 vlan-reservation 4061-4094
 vlan 311 bridge 1 name GPU-311 state enable
 vlan 312 bridge 1 name GPU-312 state enable
 vlan 313 bridge 1 name GPU-313 state enable
 vlan 314 bridge 1 name GPU-314 state enable
```

```
 vlan 315 bridge 1 name GPU-315 state enable
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn esi hold-time 100
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_10
 description ****SH-SH-REMOTE-ROUTING-SYMMETRIC-IRB-BETWEEN TENENT-4 AND TENENT-5****
 rd 3333:10
 route-target both 10:10
 l3vni 101010
!
ip vrf l3_vrf_23
 description ****SH-SH-REMOTE-ROUTING-ASYMMETRIC-IRB-BETWEEN TENENT-2 AND TENENT-3****
 rd 3333:23
 route-target both 23:23
!
ip vrf l3_vrf_45
 description ****SH-SH-REMOTE-ROUTING-SYMMETRIC-IRB-BETWEEN TENENT-4 AND TENENT-5****
 rd 3333:45
 route-target both 45:45
 l3vni 77777
!
ip vrf l3_vrf_67
 description ****SH-MH-REMOTE-ROUTING-ASYMMETRIC-IRB-BETWEEN TENENT-6 AND TENENT-7****
 rd 3333:67
 route-target both 67:67
!
ip vrf l3_vrf_89
 description ****SH-MH-REMOTE-ROUTING-SYMMETRIC-IRB-BETWEEN TENENT-8 AND TENENT-9****
 rd 3333:89
 route-target both 89:89
 l3vni 88888
!
ip vrf management
!
mac vrf l2_vrf_10
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-10****
 rd 3.3.3.3:10
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_2
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-2****
 rd 3.3.3.3:2
```

```
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_3
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-3****
 rd 3.3.3.3:3
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_4
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-4****
 rd 3.3.3.3:4
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_6
 evpn-vlan-service vlan-based
 description ****MH-SH-REMOTE-SWITCHING-TENENT-6****
 rd 3.3.3.3:6
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_7
 evpn-vlan-service vlan-based
 description ****MH-SH-REMOTE-SWITCHING-TENENT-7****
 rd 3.3.3.3:7
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_8
 evpn-vlan-service vlan-based
 description ****MH-SH-REMOTE-SWITCHING-TENENT-8****
 rd 3.3.3.3:8
 route-target both evpn-auto-rt
!
evpn irb-forwarding anycast-gateway-mac deed.beef.cafe
!
nvo vxlan vtep-ip-global 3.3.3.3
!
nvo vxlan id 50002 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_2
 evpn irb2
!
nvo vxlan id 50003 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_3
 evpn irb3
!
nvo vxlan id 50004 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_4
 evpn irb4
!
nvo vxlan id 50006 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_6
```

```
 evpn irb6
!
nvo vxlan id 50007 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_7
 evpn irb7
!
nvo vxlan id 50008 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_8
 evpn irb8
!
nvo vxlan id 50010 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_10
 evpn irb10
!
route-map HIG_MED permit 10
 set metric 55555
!
interface po1
 description ****SW1-MH****
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 evpn multi-homed system-mac 0000.0000.3333
!
interface cd3/1
 priority-flow-control mode auto
 ip address 30.1.15.1/24
 ipv6 address 3015::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd4/1
 priority-flow-control mode auto
 ip address 30.1.16.1/24
 ipv6 address 3016::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
interface cd8/1
 channel-group 1 mode active
 lldp-agent
  exit
!
interface eth0
```

```
 ip vrf forwarding management
 ip address dhcp
!
interface irb2
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 2.1.1.1/24 anycast
 ipv6 address 2::1/48 anycast
!
interface irb3
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 3.1.1.1/24 anycast
 ipv6 address 3::1/48 anycast
!
interface irb4
 ip vrf forwarding l3_vrf_45
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 4.1.1.1/24 anycast
 ipv6 address 4::1/48 anycast
!
interface irb6
 ip vrf forwarding l3_vrf_67
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 6.1.1.1/24 anycast
 ipv6 address 6::1/48 anycast
!
interface irb7
 ip vrf forwarding l3_vrf_67
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 7.1.1.1/24 anycast
 ipv6 address 7::1/48 anycast
!
interface irb8
 ip vrf forwarding l3_vrf_89
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 8.1.1.1/24 anycast
 ipv6 address 8::1/48 anycast
!
interface irb10
 ip vrf forwarding l3_vrf_10
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 10.1.1.1/24 anycast
 ipv6 address 10::1/48 anycast
!
interface irb102
 evpn irb-if-forwarding anycast-gateway-mac
!
interface lo
 ip address 127.0.0.1/8
```

```
 ip address 3.3.3.3/32 secondary
 ipv6 address ::1/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface vlan1.311
 ip address 173.1.1.1/24
 ipv6 address 173:1::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface vlan1.312
 ip address 173.1.2.1/24
 ipv6 address 173:2::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface vlan1.313
 ip address 173.1.3.1/24
 ipv6 address 173:3::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface vlan1.314
 ip address 173.1.4.1/24
 ipv6 address 173:4::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface vlan1.315
 ip address 173.1.5.1/24
 ipv6 address 173:5::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface xe0
 description ***SH to SH --> STC 2/17 ****
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 lldp-agent
  exit
!
router ospf 100
 ospf router-id 3.3.3.3
 shutdown
 network 3.3.3.3/32 area 0.0.0.0
 network 10.1.15.0/24 area 0.0.0.0
 network 10.1.16.0/24 area 0.0.0.0
 network 20.1.15.0/24 area 0.0.0.0
```

```
 network 20.1.16.0/24 area 0.0.0.0
 network 30.1.15.0/24 area 0.0.0.0
 network 30.1.16.0/24 area 0.0.0.0
 network 103.1.1.0/24 area 0.0.0.1
 network 113.1.1.0/24 area 0.0.0.1
 network 173.1.1.0/24 area 0.0.0.1
 network 173.1.2.0/24 area 0.0.0.1
 network 173.1.3.0/24 area 0.0.0.1
 network 173.1.4.0/24 area 0.0.0.1
 network 173.1.5.0/24 area 0.0.0.1
!
router ipv6 ospf
 router-id 3.3.3.3
!
router bgp 4294967203
 bgp router-id 3.3.3.3
 bgp bestpath as-path multipath-relax
 bgp log-neighbor-changes
 no bgp inbound-route-filter
 neighbor overlay peer-group
 neighbor overlay remote-as 4294967209
 neighbor overlay ebgp-multihop
 neighbor overlay update-source lo
 neighbor overlay advertisement-interval 0
 neighbor overlay fall-over bfd multihop
 neighbor underlay peer-group
 neighbor underlay remote-as 4294967209
 neighbor underlay authentication-key 0xb59db09d828b2528
 neighbor underlay as-origination-interval 1
 neighbor underlay advertisement-interval 0
 neighbor underlay fall-over bfd
 neighbor 7.7.7.7 peer-group overlay
 neighbor 8.8.8.8 peer-group overlay
 !
 bgp unnumbered-mode
 neighbor cd4/1 peergroup underlay
 neighbor cd3/1 peergroup underlay
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 10
 redistribute connected
 neighbor underlay activate
 !
 bgp v4-unnumbered-mode
 neighbor cd4/1 route-map HIG_MED in
 neighbor cd3/1 route-map HIG_MED in
 exit-v4-unnumbered-mode
 !
 exit-address-family
```

```
!
address-family l2vpn evpn
neighbor overlay activate
!
bgp l2vpn-unnumbered-mode
exit-l2vpn-unnumbered-mode
!
exit-address-family
!
address-family ipv4 vrf l3_vrf_10
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_23
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_45
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_67
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_89
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_10
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_23
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_45
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_67
```

```
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 address-family ipv6 vrf l3_vrf_89
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 exit
!
line console 0
 exec-timeout 0 0
line vty 0 16
 exec-timeout 0 0
!
nvo vxlan access-if port-vlan xe0 2
 description ***VTEP3-SH-TENENT-2***
 map vnid 50002
!
nvo vxlan access-if port-vlan xe0 3
 description ***VTEP3-SH-TENENT-3***
 map vnid 50003
!
nvo vxlan access-if port-vlan xe0 4
 description ***VTEP3-SH-TENENT-4***
 map vnid 50004
!
nvo vxlan access-if port-vlan xe0 10
 description ***VTEP3-SH-TENENT-10***
 map vnid 50010
!
nvo vxlan access-if port-vlan po1 6
 description ***VTEP3-MH-TENENT-6***
 map vnid 50006
!
nvo vxlan access-if port-vlan po1 7
 description ***VTEP3-MH-TENENT-7***
 map vnid 50007
!
nvo vxlan access-if port-vlan po1 8
 description ***VTEP3-MH-TENENT-8***
 map vnid 50008
!
!
end
```

VTEP2

```
!
```

```
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
!
service password-encryption
!
logging console 5
logging monitor 4
logging level all 2
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps ospf
snmp-server enable traps bgp
banner motd "NODE BELONGS TO L2 TEAM TESTING.PLEASE REFINE FROM USING IT"
!
bgp extended-asn-cap
!
load-balance rtag7
load-balance rtag7 all-lag-if
load-balance rtag7 macro-flow
load-balance rtag7 l2 dest-mac src-mac ether-type vlan
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port next-hdr
hardware-profile filter ingress-ipv6 disable
hardware-profile filter bfd-group enable
!
bfd interval 100 minrx 100 multiplier 3
!
qos enable
!
hostname Leaf4
default-interface type eth-routed mtu 9200
default-interface type eth-switchport mtu 9200
default-interface type svi mtu 9200
default-interface type irb mtu 9200
default-interface type lag mtu 9200
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
default-interface load-interval 30
bridge 1 protocol ieee vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1
priority-flow-control enable bridge 1
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
```

```
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
feature rsyslog
logging remote server 10.16.100.20 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp tlv-select ieee-8021-org-specific data-center-bridging
set lldp msg-tx-hold 1
set lldp timer msg-tx-interval 5
set lldp tx-fast-init 1
lldp notification-interval 1000
fault-management enable
maximum-paths 64
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
```

```
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 10 gbps
  priority
  lossless
```

```
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
vlan database
 vlan-reservation 4061-4094
 vlan 411 bridge 1 name GPU-411 state enable
 vlan 412 bridge 1 name GPU-412 state enable
 vlan 413 bridge 1 name GPU-413 state enable
 vlan 414 bridge 1 name GPU-414 state enable
 vlan 415 bridge 1 name GPU-415 state enable
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn esi hold-time 100
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_10
 description ****SH-SH-REMOTE-ROUTING-SYMMETRIC-IRB-BETWEEN TENENT-4 AND TENENT-5****
 rd 4444:10
 route-target both 10:10
 l3vni 101010
!
ip vrf l3_vrf_23
 description ****SH-SH-REMOTE-ROUTING-ASYMMETRIC-IRB-BETWEEN TENENT-2 AND TENENT-3****
 rd 4444:23
 route-target both 23:23
!
ip vrf l3_vrf_45
 description ****SH-SH-REMOTE-ROUTING-SYMMETRIC-IRB-BETWEEN TENENT-4 AND TENENT-5****
 rd 4444:45
 route-target both 45:45
 l3vni 77777
!
ip vrf l3_vrf_67
 description ****SH-MH-REMOTE-ROUTING-ASYMMETRIC-IRB-BETWEEN TENENT-6 AND TENENT-7****
 rd 4444:67
 route-target both 67:67
!
ip vrf l3_vrf_89
 description ****SH-MH-REMOTE-ROUTING-SYMMETRIC-IRB-BETWEEN TENENT-8 AND TENENT-9****
```

```
 rd 4444:89
 route-target both 89:89
 l3vni 88888
!
ip vrf management
!
mac vrf l2_vrf_10
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-10****
 rd 4.4.4.4:10
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_2
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-2****
 rd 4.4.4.4:2
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_3
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-3****
 rd 4.4.4.4:3
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_5
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-5****
 rd 4.4.4.4:5
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_6
 evpn-vlan-service vlan-based
 description ****MH-SH-REMOTE-SWITCHING-TENENT-6****
 rd 4.4.4.4:6
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_7
 evpn-vlan-service vlan-based
 description ****MH-SH-REMOTE-SWITCHING-TENENT-7****
 rd 4.4.4.4:7
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_8
 evpn-vlan-service vlan-based
 description ****MH-SH-REMOTE-SWITCHING-TENENT-8****
 rd 4.4.4.4:8
 route-target both evpn-auto-rt
!
evpn irb-forwarding anycast-gateway-mac deed.beef.cafe
!
```

```
nvo vxlan vtep-ip-global 4.4.4.4
!
nvo vxlan id 50002 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_2
 evpn irb2
!
nvo vxlan id 50003 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_3
 evpn irb3
!
nvo vxlan id 50005 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_5
 evpn irb5
!
nvo vxlan id 50006 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_6
 evpn irb6
!
nvo vxlan id 50007 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_7
 evpn irb7
!
nvo vxlan id 50008 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_8
 evpn irb8
!
nvo vxlan id 50010 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_10
 evpn irb10
!
route-map HIG_MED permit 10
 set metric 55555
!
interface po1
 description ****SW1-MH****
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 evpn multi-homed system-mac 0000.0000.3333
!
interface ce1
 priority-flow-control mode auto
 ip address 10.1.17.1/24
 ipv6 address 1017::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
```

```
!
interface ce3
 speed 100g
 shutdown
!
interface ce5
 priority-flow-control mode auto
 ip address 10.1.18.1/24
 ipv6 address 1018::1/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ce6
!
interface ce7
 description ***GPU-14 --> STC  ***
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 114.1.1.1/24
 ipv6 address 114::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
 lldp-agent
  exit
!
interface ce8
 description ***GPU-24.1 to 24.5 --> STC  ***
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 124.1.1.1/24
 ipv6 address 124::1/64
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ce31
 channel-group 1 mode active
 lldp-agent
  exit
!
interface ce32
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
```

```
interface irb2
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 2.1.1.1/24 anycast
 ipv6 address 2::1/48 anycast
!
interface irb3
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 3.1.1.1/24 anycast
 ipv6 address 3::1/48 anycast
!
interface irb5
 ip vrf forwarding l3_vrf_45
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 5.1.1.1/24 anycast
 ipv6 address 5::1/48 anycast
!
interface irb6
 ip vrf forwarding l3_vrf_67
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 6.1.1.1/24 anycast
 ipv6 address 6::1/48 anycast
!
interface irb7
 ip vrf forwarding l3_vrf_67
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 7.1.1.1/24 anycast
 ipv6 address 7::1/48 anycast
!
interface irb8
 ip vrf forwarding l3_vrf_89
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 8.1.1.1/24 anycast
 ipv6 address 8::1/48 anycast
!
interface irb10
 ip vrf forwarding l3_vrf_10
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 20.1.1.1/24 anycast
 ipv6 address 20::1/48 anycast
!
interface irb103
 evpn irb-if-forwarding anycast-gateway-mac
!
interface lo
 ip address 127.0.0.1/8
 ip address 4.4.4.4/32 secondary
 ipv6 address ::1/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
```

```
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface vlan1.411
 ip address 174.1.1.1/24
 ipv6 address 174:1::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface vlan1.412
 ip address 174.1.2.1/24
 ipv6 address 174:2::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface vlan1.413
 ip address 174.1.3.1/24
 ipv6 address 174:3::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface vlan1.414
 ip address 174.1.4.1/24
 ipv6 address 174:4::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface vlan1.415
 ip address 174.1.5.1/24
 ipv6 address 174:5::1/64
 ipv6 router ospf area 0.0.0.1 instance-id 0
!
interface xe33
!
interface xe34
 description ***GPU-4 --> STC 6/6 ***
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 lldp-agent
  exit
!
 exit
!
router ospf 100
 ospf router-id 4.4.4.4
 shutdown
 network 4.4.4.4/32 area 0.0.0.0
 network 10.1.17.0/24 area 0.0.0.0
 network 10.1.18.0/24 area 0.0.0.0
```

```
 network 104.1.1.0/24 area 0.0.0.1
 network 174.1.1.0/24 area 0.0.0.1
 network 174.1.2.0/24 area 0.0.0.1
 network 174.1.3.0/24 area 0.0.0.1
 network 174.1.4.0/24 area 0.0.0.1
 network 174.1.5.0/24 area 0.0.0.1
!
router ipv6 ospf
 router-id 4.4.4.4
!
router bgp 4294967204
 bgp router-id 4.4.4.4
 bgp bestpath as-path multipath-relax
 bgp log-neighbor-changes
 neighbor overlay peer-group
 neighbor overlay remote-as 4294967209
 neighbor overlay ebgp-multihop
 neighbor overlay update-source lo
 neighbor overlay advertisement-interval 0
 neighbor overlay fall-over bfd multihop
 neighbor underlay peer-group
 neighbor underlay remote-as 4294967209
 neighbor underlay authentication-key 0xb59db09d828b2528
 neighbor underlay as-origination-interval 1
 neighbor underlay advertisement-interval 0
 neighbor underlay fall-over bfd
 neighbor 7.7.7.7 peer-group overlay
 neighbor 8.8.8.8 peer-group overlay
 !
 bgp unnumbered-mode
 neighbor ce1 peergroup underlay
 neighbor ce5 peergroup underlay
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 10
 redistribute connected
 neighbor underlay activate
 !
 bgp v4-unnumbered-mode
 neighbor ce1 route-map HIG_MED in
 neighbor ce5 route-map HIG_MED in
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor overlay activate
 !
 bgp l2vpn-unnumbered-mode
```

```
exit-l2vpn-unnumbered-mode
!
exit-address-family
!
address-family ipv4 vrf l3_vrf_10
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_23
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_45
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_67
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv4 vrf l3_vrf_89
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_10
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_23
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_45
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_67
max-paths ebgp 50
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3_vrf_89
```

```
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 exit
!
line console 0
 exec-timeout 0 0
line vty 0 16
 exec-timeout 0 0
!
nvo vxlan access-if port-vlan xe34 2
 description ***VTEP4-SH-TENENT-2***
 map vnid 50002
!
nvo vxlan access-if port-vlan xe34 3
 description ***VTEP4-SH-TENENT-3***
 map vnid 50003
!
nvo vxlan access-if port-vlan xe34 5
 description ***VTEP4-SH-TENENT-5***
 map vnid 50005
!
nvo vxlan access-if port-vlan xe34 10
 description ***VTEP4-SH-TENENT-10***
 map vnid 50010
!
nvo vxlan access-if port-vlan po1 6
 description ***VTEP4-MH-TENENT-6***
 map vnid 50006
!
nvo vxlan access-if port-vlan po1 7
 description ***VTEP4-MH-TENENT-7***
 map vnid 50007
!
nvo vxlan access-if port-vlan po1 8
 description ***VTEP4-MH-TENENT-8***
 map vnid 50008
!
!
end
```

VTEP3

```
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
!
feature streaming-telemetry vrf management
```

```
!
telemetry maximum-subscribe-paths 140
!
service password-encryption
!
logging console 5
logging level all 2
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps ospf
snmp-server enable traps bgp
banner motd "NODE BELONGS TO L2 TEAM TESTING.PLEASE REFINE FROM USING IT"
!
bgp extended-asn-cap
!
load-balance rtag7
load-balance rtag7 macro-flow
load-balance rtag7 l2 dest-mac src-mac ether-type vlan
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port next-hdr
hardware-profile filter ingress-ipv6 disable
hardware-profile filter bfd-group enable
!
bfd interval 100 minrx 100 multiplier 3
!
qos enable
!
hostname Leaf5
port ce4 breakout 4X10g
default-interface type eth-routed mtu 9200
default-interface type eth-switchport mtu 9200
default-interface type svi mtu 9200
default-interface type irb mtu 9200
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
default-interface load-interval 30
bridge 1 protocol ieee vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1
priority-flow-control enable bridge 1
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
```

```
feature rsyslog
logging remote server 10.16.100.22 5 port 1514 vrf management
logging remote server 10.16.100.20 5 port 1514 vrf management
logging remote server 192.168.6.102 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp tlv-select ieee-8021-org-specific data-center-bridging
set lldp msg-tx-hold 1
set lldp timer msg-tx-interval 5
set lldp tx-fast-init 1
lldp notification-interval 1000
fault-management enable
maximum-paths 64
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
```

```
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
```

```
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
vlan database
 vlan-reservation 4001-4094
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn esi hold-time 100
!
evpn vxlan multihoming enable
!
ip vrf l3_vrf_23
 description ****SH-SH-REMOTE-ROUTING-ASYMMETRIC-IRB-BETWEEN TENENT-2 AND TENENT-3****
 rd 5555:23
 route-target both 23:23
!
ip vrf l3_vrf_67
 description ****SH-MH-REMOTE-ROUTING-ASYMMETRIC-IRB-BETWEEN TENENT-6 AND TENENT-7****
 rd 5555:67
 route-target both 67:67
!
ip vrf l3_vrf_89
 description ****SH-MH-REMOTE-ROUTING-SYMMETRIC-IRB-BETWEEN TENENT-8 AND TENENT-9****
 rd 5555:89
 route-target both 89:89
 l3vni 88888
!
ip vrf management
!
mac vrf l2_vrf_2
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-2****
 rd 5.5.5.5:2
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_3
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-3****
 rd 5.5.5.5:3
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_6
 evpn-vlan-service vlan-based
```

```
 description ****SH-MH-REMOTE-SWITCHING-TENENT-6****
 rd 5.5.5.5:6
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_7
 evpn-vlan-service vlan-based
 description ****SH-MH-REMOTE-SWITCHING-TENENT-7****
 rd 5.5.5.5:7
 route-target both evpn-auto-rt
!
mac vrf l2_vrf_9
 evpn-vlan-service vlan-based
 description ****SH-SH-REMOTE-SWITCHING-TENENT-9****
 rd 5.5.5.5:9
 route-target both evpn-auto-rt
!
evpn irb-forwarding anycast-gateway-mac deed.beef.cafe
!
nvo vxlan vtep-ip-global 5.5.5.5
!
nvo vxlan id 50002 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_2
 evpn irb2
!
nvo vxlan id 50003 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_3
 evpn irb3
!
nvo vxlan id 50006 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_6
 evpn irb6
!
nvo vxlan id 50007 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_7
 evpn irb7
!
nvo vxlan id 50009 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp l2_vrf_9
 evpn irb9
!
nvo vxlan id 50010 ingress-replication inner-vid-disabled
!
route-map HIG_MED permit 10
 set metric 55555
!
interface ce2
 priority-flow-control mode auto
!
interface ce3
 priority-flow-control mode auto
```

```
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface irb2
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 2.1.1.1/24 anycast
 ipv6 address 2::1/48 anycast
!
interface irb3
 ip vrf forwarding l3_vrf_23
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 3.1.1.1/24 anycast
 ipv6 address 3::1/48 anycast
!
interface irb5
 evpn irb-if-forwarding anycast-gateway-mac
 shutdown
!
interface irb6
 ip vrf forwarding l3_vrf_67
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 6.1.1.1/24 anycast
 ipv6 address 6::1/48 anycast
!
interface irb7
 ip vrf forwarding l3_vrf_67
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 7.1.1.1/24 anycast
 ipv6 address 7::1/48 anycast
!
interface irb9
 ip vrf forwarding l3_vrf_89
 evpn irb-if-forwarding anycast-gateway-mac
 ip address 9.1.1.1/24 anycast
 ipv6 address 9::1/48 anycast
!
interface irb10
 evpn irb-if-forwarding anycast-gateway-mac
 shutdown
!
interface lo
 ip address 127.0.0.1/8
 ip address 5.5.5.5/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
```

```
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe4/4
 description ***GPU-4 --> STC 6/6 ***
 switchport
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 lldp-agent
  exit
!
router bgp 4294967205
 bgp router-id 5.5.5.5
 bgp bestpath as-path multipath-relax
 bgp log-neighbor-changes
 neighbor overlay peer-group
 neighbor overlay remote-as 4294967209
 neighbor overlay ebgp-multihop
 neighbor overlay update-source lo
 neighbor overlay advertisement-interval 0
 neighbor overlay fall-over bfd multihop
 neighbor underlay peer-group
 neighbor underlay remote-as 4294967209
 neighbor underlay authentication-key 0xb59db09d828b2528
 neighbor underlay as-origination-interval 1
 neighbor underlay advertisement-interval 0
 neighbor underlay fall-over bfd
 neighbor 7.7.7.7 peer-group overlay
 neighbor 8.8.8.8 peer-group overlay
 !
 bgp unnumbered-mode
 neighbor ce3 peergroup underlay
 neighbor ce2 peergroup underlay
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 10
 redistribute connected
 neighbor underlay activate
 !
 bgp v4-unnumbered-mode
 neighbor ce3 route-map HIG_MED in
 neighbor ce2 route-map HIG_MED in
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor overlay activate
```

```
 !
 bgp l2vpn-unnumbered-mode
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
 address-family ipv4 vrf l3_vrf_23
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 address-family ipv4 vrf l3_vrf_67
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 address-family ipv4 vrf l3_vrf_89
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 address-family ipv6 vrf l3_vrf_23
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 address-family ipv6 vrf l3_vrf_67
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 address-family ipv6 vrf l3_vrf_89
 max-paths ebgp 50
 redistribute connected
 exit-address-family
 !
 exit
!
line console 0
 exec-timeout 0 0
line vty 0 16
 exec-timeout 0 0
!
nvo vxlan access-if port-vlan xe4/4 6
 description ***VTEP5-SH-TENENT-6***
 map vnid 50006
!
nvo vxlan access-if port-vlan xe4/4 7
 description ***VTEP5-SH-TENENT-7***
 map vnid 50007
```

```
!
nvo vxlan access-if port-vlan xe4/4 9
 description ***VTEP5-SH-TENENT-9***
 map vnid 50009
!
!
end
```

Spine 1

```
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
!
service password-encryption
!
logging monitor 4
logging level all 2
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps ospf
snmp-server enable traps bgp
banner motd "NODE BELONGS TO L2 TEAM TESTING.PLEASE REFINE FROM USING IT"
!
bgp extended-asn-cap
!
load-balance rtag7
load-balance rtag7 macro-flow
load-balance rtag7 l2 dest-mac src-mac ether-type vlan
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id rocev2-
dest-qpairs
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port rocev2-dest-qpairs
next-hdr
!
bfd interval 100 minrx 100 multiplier 3
!
qos enable
!
hostname Spine1
port ethernet63 breakout 1X100g serdes 25g
port ethernet11 breakout 1X100g serdes 25g
port ethernet1 breakout 1X100g serdes 25g
default-interface type eth-routed mtu 9200
default-interface type eth-switchport mtu 9200
default-interface type svi mtu 9200
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
default-interface load-interval 30
tfo Disable
errdisable cause stp-bpdu-guard
```

```
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
username aimldlb role network-admin password encrypted
$6$bJoWX2G1$aPj0g5d5ithuKPmXMvxjxtLgCBiasltAcOUfxRC6DGMfO99X5R7.JKUbznYl0vqiqYWTc64cl.M
fKNFXF1CSw.
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp tlv-select ieee-8021-org-specific data-center-bridging
set lldp msg-tx-hold 1
set lldp timer msg-tx-interval 5
set lldp tx-fast-init 1
lldp notification-interval 5
fault-management enable
maximum-paths 64
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
```

```
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 10 gbps
  priority
```

```
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
vlan database
 vlan-reservation 4001-4094
!
ip vrf management
!
route-map HIG_MED permit 10
 set metric 55555
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface ethernet1/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 10.1.17.2/24
 ipv6 address 1017::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet2/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 20.1.11.2/24
 ipv6 address 2011::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet3/1
```

```
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 10.1.11.2/24
 ipv6 address 1011::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet4
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 40.1.13.2/24
 ipv6 address 4013::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet5
!
interface ethernet6/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 lldp-agent
  exit
!
interface ethernet11/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 10.1.19.2/24
 ipv6 address 1019::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet25/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
```

```
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 50.1.13.2/24
 ipv6 address 5013::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet25/2
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 60.1.13.2/24
 ipv6 address 6013::2/64
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet30/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 20.1.13.2/24
 ipv6 address 2013::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet31/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 10.1.15.2/24
 ipv6 address 1015::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet32/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
```

```
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 20.1.15.2/24
 ipv6 address 2015::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet61/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 30.1.13.2/24
 ipv6 address 3013::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet62/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 30.1.11.2/24
 ipv6 address 3011::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet63/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 30.1.15.2/24
 ipv6 address 3015::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface ethernet64/1
```

```
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 ip address 10.1.13.2/24
 ipv6 address 1013::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface lo
 ip address 127.0.0.1/8
 ip address 7.7.7.7/32 secondary
 ipv6 address ::1/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface mgmt1
 ip vrf forwarding management
!
interface mgmt2
 ip vrf forwarding management
!
 exit
!
router ospf 100
 ospf router-id 7.7.7.7
 shutdown
 network 7.7.7.7/32 area 0.0.0.0
 network 10.1.11.0/24 area 0.0.0.0
 network 10.1.13.0/24 area 0.0.0.0
 network 10.1.15.0/24 area 0.0.0.0
 network 10.1.17.0/24 area 0.0.0.0
 network 10.1.19.0/24 area 0.0.0.0
 network 20.1.11.0/24 area 0.0.0.0
 network 20.1.13.0/24 area 0.0.0.0
 network 20.1.15.0/24 area 0.0.0.0
 network 20.1.17.0/24 area 0.0.0.0
 network 30.1.11.0/24 area 0.0.0.0
 network 30.1.13.0/24 area 0.0.0.0
 network 30.1.15.0/24 area 0.0.0.0
 network 30.1.17.0/24 area 0.0.0.0
 network 40.1.13.0/24 area 0.0.0.0
 network 50.1.13.0/24 area 0.0.0.0
```

```
 network 60.1.13.0/24 area 0.0.0.0
!
router ipv6 ospf
 router-id 7.7.7.7
!
router bgp 4294967209
 bgp router-id 7.7.7.7
 bgp bestpath as-path multipath-relax
 bgp log-neighbor-changes
 no bgp inbound-route-filter
 neighbor overlay peer-group
 neighbor overlay remote-as 12345
 neighbor overlay ebgp-multihop
 neighbor overlay update-source lo
 neighbor overlay advertisement-interval 0
 neighbor overlay fall-over bfd multihop
 neighbor underlay peer-group
 neighbor underlay remote-as 1
 neighbor underlay authentication-key 0xb59db09d828b2528
 neighbor underlay as-origination-interval 1
 neighbor underlay advertisement-interval 0
 neighbor underlay fall-over bfd
 neighbor 3.3.3.3 remote-as 4294967203
 neighbor 3.3.3.3 peer-group overlay
 neighbor 4.4.4.4 remote-as 4294967204
 neighbor 4.4.4.4 peer-group overlay
 neighbor 5.5.5.5 remote-as 4294967205
 neighbor 5.5.5.5 peer-group overlay
 !
 bgp unnumbered-mode
 neighbor ethernet6/1 remote-as external
 neighbor ethernet6/1 peergroup underlay
 neighbor ethernet1/1 remote-as external
 neighbor ethernet1/1 peergroup underlay
 neighbor ethernet31/1 remote-as external
 neighbor ethernet31/1 peergroup underlay
 neighbor ethernet63/1 remote-as external
 neighbor ethernet63/1 peergroup underlay
 neighbor ethernet32/1 remote-as external
 neighbor ethernet32/1 peergroup underlay
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 10
 redistribute connected
 neighbor underlay activate
 !
 bgp v4-unnumbered-mode
 neighbor ethernet6/1 route-map HIG_MED in
 neighbor ethernet1/1 route-map HIG_MED in
```

```
 neighbor ethernet31/1 route-map HIG_MED in
 neighbor ethernet63/1 route-map HIG_MED in
 neighbor ethernet32/1 route-map HIG_MED in
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor overlay activate
 !
 bgp l2vpn-unnumbered-mode
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
 exit
!
line console 0
 exec-timeout 0 0
line vty 0 16
 exec-timeout 0 0
!
!
end
```

Spine 2

```
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
!
service password-encryption
!
logging monitor 4
logging level all 2
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps ospf
snmp-server enable traps bgp
banner motd "NODE BELONGS TO L2 TEAM TESTING.PLEASE REFINE FROM USING IT"
!
bgp extended-asn-cap
!
load-balance rtag7
load-balance rtag7 macro-flow
load-balance rtag7 l2 dest-mac src-mac ether-type vlan
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id rocev2-
dest-qpairs
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port rocev2-dest-qpairs
next-hdr
!
bfd interval 100 minrx 100 multiplier 3
```

```
!
qos enable
!
hostname Spine2
port cd3 breakout 1X100g serdes 25g
port cd9 breakout 1X100g serdes 25g
port cd27 breakout 1X100g serdes 25g
default-interface type eth-routed mtu 9200
default-interface type eth-switchport mtu 9200
default-interface type svi mtu 9200
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
default-interface load-interval 30
tfo Disable
errdisable cause stp-bpdu-guard
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
username aimldlb role network-admin password encrypted $6$bJoWB2p/
$xK0l6QeSWRV2HmmXbPBRja8uM0ugLRt6o7V4wliZE4p2bNmWfBDY8Zo7iuLzki84KvThCkmerZEq6w09IvOzA1
feature rsyslog
logging remote server 10.16.100.20 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp tlv-select ieee-8021-org-specific data-center-bridging
set lldp msg-tx-hold 1
set lldp timer msg-tx-interval 5
set lldp tx-fast-init 1
lldp notification-interval 1000
fault-management enable
maximum-paths 64
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
```

```
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
```

```
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
vlan database
 vlan-reservation 4061-4094
!
ip vrf management
!
route-map HIG_MED permit 10
 set metric 55555
!
interface cd1/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 10.1.12.2/24
 ipv6 address 1012::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd2/1
 priority-flow-control mode on
```

```
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 lldp-agent
  exit
!
interface cd3/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 10.1.18.2/24
 ipv6 address 1018::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd4/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 20.1.12.2/24
 ipv6 address 2012::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd5
!
interface cd6
!
interface cd7/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 20.1.16.2/24
 ipv6 address 2016::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd8
!
interface cd9/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
```

```
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 10.1.20.2/24
 ipv6 address 1020::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd24
!
interface cd25/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 30.1.14.2/24
 ipv6 address 3014::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd26
!
interface cd27/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 30.1.16.2/24
 ipv6 address 3016::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd28
!
interface cd29/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 20.1.14.2/24
 ipv6 address 2014::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
```

```
!
interface cd30/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 10.1.14.2/24
 ipv6 address 1014::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd31/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 30.1.12.2/24
 ipv6 address 3012::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface cd32/1
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 ip address 10.1.16.2/24
 ipv6 address 1016::2/64
 ip ospf cost 100
 ipv6 ospf cost 100 instance-id 0
 ipv6 router ospf area 0.0.0.0 instance-id 0
 lldp-agent
  exit
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 8.8.8.8/32 secondary
 ipv6 address ::1/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
```

```
 ipv6 address ::1/128
!
interface xe33
 shutdown
!
interface xe34
 shutdown
!
 exit
!
router ospf 100
 ospf router-id 8.8.8.8
 shutdown
 network 8.8.8.8/32 area 0.0.0.0
 network 10.1.12.0/24 area 0.0.0.0
 network 10.1.14.0/24 area 0.0.0.0
 network 10.1.16.0/24 area 0.0.0.0
 network 10.1.18.0/24 area 0.0.0.0
 network 10.1.20.0/24 area 0.0.0.0
 network 20.1.12.0/24 area 0.0.0.0
 network 20.1.14.0/24 area 0.0.0.0
 network 20.1.16.0/24 area 0.0.0.0
 network 20.1.18.0/24 area 0.0.0.0
 network 30.1.12.0/24 area 0.0.0.0
 network 30.1.14.0/24 area 0.0.0.0
 network 30.1.16.0/24 area 0.0.0.0
 network 30.1.18.0/24 area 0.0.0.0
!
router ipv6 ospf
 router-id 8.8.8.8
!
router bgp 4294967209
 bgp router-id 8.8.8.8
 bgp bestpath as-path multipath-relax
 bgp log-neighbor-changes
 no bgp inbound-route-filter
 neighbor overlay peer-group
 neighbor overlay remote-as 12345
 neighbor overlay ebgp-multihop
 neighbor overlay update-source lo
 neighbor overlay advertisement-interval 0
 neighbor overlay fall-over bfd multihop
 neighbor underlay peer-group
 neighbor underlay remote-as 1
 neighbor underlay authentication-key 0xb59db09d828b2528
 neighbor underlay as-origination-interval 1
 neighbor underlay advertisement-interval 0
 neighbor underlay fall-over bfd
 neighbor 3.3.3.3 remote-as 4294967203
 neighbor 3.3.3.3 peer-group overlay
```

```
 neighbor 4.4.4.4 remote-as 4294967204
 neighbor 4.4.4.4 peer-group overlay
 neighbor 5.5.5.5 remote-as 4294967205
 neighbor 5.5.5.5 peer-group overlay
 !
 bgp unnumbered-mode
 neighbor cd2/1 remote-as external
 neighbor cd2/1 peergroup underlay
 neighbor cd3/1 remote-as external
 neighbor cd3/1 peergroup underlay
 neighbor cd32/1 remote-as external
 neighbor cd32/1 peergroup underlay
 neighbor cd27/1 remote-as external
 neighbor cd27/1 peergroup underlay
 neighbor cd7/1 remote-as external
 neighbor cd7/1 peergroup underlay
 exit-unnumbered-mode
 !
 address-family ipv4 unicast
 max-paths ebgp 10
 redistribute connected
 neighbor underlay activate
 !
 bgp v4-unnumbered-mode
 neighbor cd2/1 route-map HIG_MED in
 neighbor cd3/1 route-map HIG_MED in
 neighbor cd32/1 route-map HIG_MED in
 neighbor cd27/1 route-map HIG_MED in
 neighbor cd7/1 route-map HIG_MED in
 exit-v4-unnumbered-mode
 !
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor overlay activate
 !
 bgp l2vpn-unnumbered-mode
 exit-l2vpn-unnumbered-mode
 !
 exit-address-family
 !
 exit
!
line console 0
 exec-timeout 0 0
line vty 0 16
 exec-timeout 0 0
!
!
end
```

Switch

```
      !
      service password-encryption
      !
      logging level all 2
      snmp-server enable traps link linkDown
      snmp-server enable traps link linkUp
      snmp-server enable traps ospf
      snmp-server enable traps bgp
      !
      bgp extended-asn-cap
      !
      load-balance rtag7
      load-balance rtag7 all-lag-if
      load-balance rtag7 macro-flow
      load-balance rtag7 l2 dest-mac src-mac ether-type vlan
      load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id
rocev2-dest-qpairs
      load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port rocev2-dest-
qpairs
      !
      bfd interval 100 minrx 100 multiplier 3
      !
      qos enable
      !
      hostname SW1-MH
      port ethernet29 breakout 1X100g serdes 25g
      port ethernet3 breakout 1X100g serdes 25g
      port ethernet1 breakout 1X100g serdes 25g
      default-interface type eth-routed mtu 9200
      default-interface type eth-switchport mtu 9200
      default-interface type svi mtu 9200
      default-interface type irb mtu 9200
      default-interface type lag mtu 9200
      ip domain-lookup vrf management
      ip name-server vrf management 10.12.3.23
      default-interface load-interval 30
      bridge 1 protocol ieee vlan-bridge
      tfo Disable
      errdisable cause stp-bpdu-guard
      data-center-bridging enable bridge 1
      feature ssh vrf management
      no feature ssh
      feature dns relay
      ip dns relay
      ipv6 dns relay
      feature rsyslog
      logging remote server 10.16.100.20 5 port 1514 vrf management
      lldp run
      lldp tlv-select basic-mgmt port-description
      lldp tlv-select basic-mgmt system-name
```

```
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
vlan database
 vlan-reservation 4029-4094
 vlan 2-100 bridge 1
!
ip vrf management
!
interface po1
 description ****VTEP3-VTEP4-MH****
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
 priority-flow-control mode auto
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface ethernet1/1
 channel-group 1 mode active
!
interface ethernet3/1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
 priority-flow-control mode on
 priority-flow-control advertise-local-config
 priority-flow-control enable priority 0 1 2 3 4 5 6 7
 load-interval 30
 lldp-agent
  exit
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface mgmt64
!
```

```
       interface mgmt65
       !
        exit
       !
       line console 0
        exec-timeout 0 0
       line vty 0 16
        exec-timeout 0 0
       !
       !
end
```

PFC Validation

VTEP1

Before Congestion

This section shows the state of the VTEP1 switch under normal operation **before any congestion** was introduced.

Verify the status of established VxLAN tunnels. The output confirms that tunnels from the local VTEP1 (3.3.3.3) to remote VTEPs 5.5.5.5 (VTEP2) and 4.4.4.4 (VTEP3) are established.

```
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
3.3.3.3          5.5.5.5          Installed     02:02:14      02:02:02
3.3.3.3          4.4.4.4          Installed     02:02:13      02:02:02

Total number of entries are 2
```

Verify the detailed information about each configured VNID.

```
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
_____________________________________________________________________________________________________________________________________________
50002    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          5.5.5.5          --------------
50002    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          --------------
50002    ----         --       AC   xe0       --- Single Homed Port ---       2         ----      ----             ----
50003    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          5.5.5.5          --------------
50003    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          --------------
50003    ----         --       AC   xe0       --- Single Homed Port ---       3         ----      ----             ----
50004    ----         --       AC   xe0       --- Single Homed Port ---       4         ----      ----             ----
50006    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          5.5.5.5          --------------
50006    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          --------------
50006    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   6         DF        ----             ----
50007    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          5.5.5.5          --------------
50007    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          --------------
50007    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   7         NON-DF    ----             ----
50008    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          --------------
50008    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   8         DF        ----             ----
50010    ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          --------------
50010    ----         --       AC   xe0       --- Single Homed Port ---       10        ----      ----             ----
77777    ----         L3       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          34ef.b696.dba1
88888    ----         L3       NW   ----      ----                            ----      ----      3.3.3.3          5.5.5.5          5c07.5870.4300
88888    ----         L3       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          34ef.b696.dba1
101010   ----         L3       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4          34ef.b696.dba1

Total number of entries are 21
```

```
VTEP1#
```

Verify the status of BGP sessions used for underlay reachability.

```
VTEP1#show bgp summary
BGP router identifier 3.3.3.3, local AS number 4294967203
BGP table version is 4
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
cd4/1(fe80::5e17:83ff:feff:3072)        4 4294967209        352        352       4      0      0  02:28:46             7
cd3/1(fe80::d6dc:85ff:fe1a:30f1)        4 4294967209        351        354       4      0      0  02:28:33             6

Total number of neighbors 2

Total number of Established sessions 2
VTEP1#
```

Verify the BGP EVPN sessions used for overlay control plane (MAC/IP advertisement).

```
VTEP1#
VTEP1#show bgp l2vpn evpn summary
BGP router identifier 3.3.3.3, local AS number 4294967203
BGP table version is 196
4 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
7.7.7.7           4 4294967209        523        647     196      0      0  02:28:35           246      4    221     12      1      8
8.8.8.8           4 4294967209        522        584     195      0      0  02:28:49           246      4    221     12      1      8

Total number of neighbors 2

Total number of Established sessions 2
VTEP1#
```

Display the current transmit (Tx) and receive (Rx) traffic rates on interfaces.

```
VTEP1#show interface counters rate gbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx gbps   |    Rx pps   |    Tx gbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 cd3/1                0.00           6             5.62           76770
 cd4/1                0.00           16            2.41           32915
 xe0                  7.98           109663        0.00           0
VTEP1#
```

The above output indicates an active traffic flow through the VxLAN fabric.

Display the configured QoS settings.

```
VTEP1#show running-config qos
qos enable
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
```

```
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
```

Display both the configured (Admin) and current operational state of PFC on all interfaces.

```
VTEP1#show priority-flow-control details all

Admin Configuration
--------------------
interface           mode  advertise willing   cap   link delay  priorities
                                                     allowance
================================================================================
cd7                 on    on        off       8     0           0 1 2 3 4 5 6 7
xe0                 on    on        off       8     0           0 1 2 3 4 5 6 7
cd8/1               on    on        off       8     0           0 1 2 3 4 5 6 7
cd0/1               auto  on        on        8     0
cd2/1               auto  on        on        8     0
cd3/1               auto  on        on        8     0
cd4/1               auto  on        on        8     0
cd5/1               auto  on        on        8     0
cd6                 on    on        off       8     0           0 1 2 3 4 5 6 7
cd10/1              auto  on        on        8     0
po1                 on    on        off       8     0           0 1 2 3 4 5 6 7

--------------------------------------------------------------------------------

Operational Configuration
-------------------------
interface           state cap   link delay  priorities
                          allowance
================================================================================
cd7                 on    8     0           0 1 2 3 4 5 6 7
xe0                 on    8     0           0 1 2 3 4 5 6 7
cd8/1               on    8     0           0 1 2 3 4 5 6 7
cd3/1               on    8     0           0 1 2 3 4 5 6 7
cd4/1               on    8     0           0 1 2 3 4 5 6 7
cd6                 on    8     0           0 1 2 3 4 5 6 7
```

Verify the counters for PFC pause frames sent and received per priority on each interface.

```
VTEP1#show priority-flow-control statistics all
interface         pri   pause sent     pause received
=====================================================
cd7               0     0             0
cd7               1     0             0
cd7               2     0             0
cd7               3     0             0
cd7               4     0             0
cd7               5     0             0
cd7               6     0             0
cd7               7     0             0
xe0               0     0             0
xe0               1     0             0
```

```
xe0               2     0              0
xe0               3     0              0
xe0               4     0              0
xe0               5     0              0
xe0               6     0              0
xe0               7     0              0
cd8/1             0     0              0
cd8/1             1     0              0
cd8/1             2     0              0
cd8/1             3     0              0
cd8/1             4     0              0
cd8/1             5     0              0
cd8/1             6     0              0
cd8/1             7     0              0
cd0/1             0     0              0
cd0/1             1     0              0
cd0/1             2     0              0
cd0/1             3     0              0
cd0/1             4     0              0
cd0/1             5     0              0
cd0/1             6     0              0
cd0/1             7     0              0
VTEP1#
```

The above output indicates that no PFC pausing occurred before congestion as all counters are zero.

After Congestion

This section shows the state after congestion was induced by applying a 1 Gbps shaper on the egress network interface `cd3/1`.

Display the QoS configuration to verify the congestion applied on the interface `cd3/1`.

Note: Congestion is induced manually in the fabric using shaper on egress interface `cd3/1`.

```
=============================================================================================================

VTEP1#show running-config qos
qos enable
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
```

```
  lossless
  exit
!
interface cd3/1
 shape rate 1 gbps burst 10000
!
!
```

Verify the traffic is throttled on ingress as per egress shape configurations.

```
VTEP1#show interface counters rate gbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx gbps   |    Rx pps   |    Tx gbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 cd3/1                0.00           6             0.70           9601
 cd4/1                0.00           16            0.30           4128
 xe0                  1.00           13707         0.00           0
VTEP1#
```

The above output traffic rates reflect the congestion. The transmit rate on the shaped interface `cd3/1` is capped around 0.70 Gbps. Consequently, the receive rate on the access interface `xe0` has reduced to ~1.00 Gbps (down from ~7.98 Gbps) as the PFC throttles the ingress traffic. The rate on the other uplink `cd4/1` is also adjusted (~0.30 Gbps).

Verify the PFC pause frames counters on ingress interface `xe0` is incremented.

```
VTEP1#
VTEP1#show priority-flow-control statistics all
interface         pri   pause sent     pause received
=====================================================
cd7               0     0              0
cd7               1     0              0
cd7               2     0              0
cd7               3     0              0
cd7               4     0              0
cd7               5     0              0
cd7               6     0              0
cd7               7     0              0
xe0               0     0              0
xe0               1     0              0
xe0               2     0              0
xe0               3     0              0
xe0               4     0              0
xe0               5     0              0
xe0               6     0              0
xe0               7     177351         0
cd8/1             0     0              0
cd8/1             1     0              0
cd8/1             2     0              0
cd8/1             3     0              0
cd8/1             4     0              0
cd8/1             5     0              0
cd8/1             6     0              0
cd8/1             7     0              0
VTEP1#
```

The `pause sent` counter for interface `xe0` (the ingress access port receiving traffic) on `priority 7` is rapidly incrementing. This means VTEP1 is sending PFC pause frames out of interface `xe0` towards the connected host (SH1) for `priority 7` because the downstream path (specifically the shaped egress `cd3/1`) is congested for traffic mapped to that priority.

Verify that the PFC administrative and operational status remain unchanged (identical to before congestion).

```
VTEP1#show priority-flow-control details all

Admin Configuration
--------------------
interface           mode  advertise willing   cap   link delay  priorities
```

```
                                                    allowance
================================================================================
cd7                 on    on        off       8     0           0 1 2 3 4 5 6 7
xe0                 on    on        off       8     0           0 1 2 3 4 5 6 7
cd8/1               on    on        off       8     0           0 1 2 3 4 5 6 7
cd0/1               auto  on        on        8     0
cd2/1               auto  on        on        8     0
cd3/1               auto  on        on        8     0
cd4/1               auto  on        on        8     0
cd5/1               auto  on        on        8     0
cd6                 on    on        off       8     0           0 1 2 3 4 5 6 7
cd10/1              auto  on        on        8     0
po1                 on    on        off       8     0           0 1 2 3 4 5 6 7

Operational Configuration
-------------------------
interface           state cap   link delay  priorities
                          allowance
================================================================================
cd7                 on    8     0           0 1 2 3 4 5 6 7
xe0                 on    8     0           0 1 2 3 4 5 6 7
cd8/1               on    8     0           0 1 2 3 4 5 6 7
cd3/1               on    8     0           0 1 2 3 4 5 6 7
cd4/1               on    8     0           0 1 2 3 4 5 6 7
cd6                 on    8     0           0 1 2 3 4 5 6 7
```

VTEP2

```
VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source             Destination       Status         Up/Down        Update
================================================================================
4.4.4.4            3.3.3.3           Installed      00:03:53       00:03:04
4.4.4.4            5.5.5.5           Installed      00:03:53       00:03:04

Total number of entries are 2
VTEP2#show nvo vxlan
VXLAN Information
==================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
--------------------------------------------------------------------------------------------------------------------------------------------
50002    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50002    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          --------------
50002    ----         --       AC   xe34      --- Single Homed Port ---       2         ----      ----             ----
50003    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50003    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          --------------
50003    ----         --       AC   xe34      --- Single Homed Port ---       3         ----      ----             ----
50005    ----         --       AC   xe34      --- Single Homed Port ---       5         ----      ----             ----
50006    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50006    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          --------------
50006    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   6         NON-DF    ----             ----
50007    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50007    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          --------------
50007    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   7         DF        ----             ----
50008    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50008    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   8         NON-DF    ----             ----
50010    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50010    ----         --       AC   xe34      --- Single Homed Port ---       10        ----      ----             ----
77777    ----         L3       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          5c07.5827.cb5f
88888    ----         L3       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          5c07.5827.cb5f
88888    ----         L3       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          5c07.5870.4300
101010   ----         L3       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          5c07.5827.cb5f

Total number of entries are 21
```

```
VTEP2#
VTEP2#
VTEP2#show bgp summary
BGP router identifier 4.4.4.4, local AS number 4294967204
BGP table version is 6
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
ce5(fe80::5e17:83ff:feff:3060)          4 4294967209         15         14       6      0      0  00:04:06            7
ce1(fe80::d6dc:85ff:fe1a:30b0)          4 4294967209         12         14       6      0      0  00:03:13            6

Total number of neighbors 2

Total number of Established sessions 2
VTEP2#show bgp l2vpn evpn summary
BGP router identifier 4.4.4.4, local AS number 4294967204
BGP table version is 12
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
7.7.7.7          4 4294967209         46         61       6      0      0  00:03:16           56      4     33     10      1      8
8.8.8.8          4 4294967209         76         48      12      0      0  00:04:08           56      4     33     10      1      8

Total number of neighbors 2

Total number of Established sessions 2
VTEP2#


VTEP2#show interface counters rate gbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx gbps   |    Rx pps   |    Tx gbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 ce1                  4.01           54787         0.00           6
 ce5                  4.01           54787         0.00           6
 xe34                 0.00           0             7.98           109564
VTEP2#


VTEP2#show priority-flow-control details all

Admin Configuration
--------------------
interface           mode  advertise willing   cap   link delay  priorities
                                                     allowance
================================================================================
ce31                on    on        off       8     0           0 1 2 3 4 5 6 7
xe34                on    on        off       8     0           0 1 2 3 4 5 6 7
ce1                 auto  on        on        8     0
```

```
ce5                  auto  on        on        8     0
ce7                  on    on        off       8     0           0 1 2 3 4 5 6 7
ce8                  on    on        off       8     0           0 1 2 3 4 5 6 7
po1                  on    on        off       8     0           0 1 2 3 4 5 6 7

-------------------------------------------------------------------------------

Operational Configuration
-------------------------
interface            state cap   link delay  priorities
                                 allowance
================================================================================
ce31                 on    8     0           0 1 2 3 4 5 6 7
xe34                 on    8     0           0 1 2 3 4 5 6 7
ce1                  on    8     0           0 1 2 3 4 5 6 7
ce5                  on    8     0           0 1 2 3 4 5 6 7
ce7                  on    8     0           0 1 2 3 4 5 6 7
ce8                  on    8     0           0 1 2 3 4 5 6 7


VTEP2#show priority-flow-control statistics all
interface          pri   pause sent     pause received
=====================================================
ce31               0     0              0
ce31               1     0              0
ce31               2     0              0
ce31               3     0              0
ce31               4     0              0
ce31               5     0              0
ce31               6     0              0
ce31               7     0              0
xe34               0     0              0
xe34               1     0              0
xe34               2     0              0
xe34               3     0              0
xe34               4     0              0
xe34               5     0              0
xe34               6     0              0
xe34               7     0              0
ce1                0     0              0
ce1                1     0              0
ce1                2     0              0
ce1                3     0              0
ce1                4     0              0
ce1                5     0              0
ce1                6     0              0
ce1                7     0              0
ce5                0     0              0
ce5                1     0              0
ce5                2     0              0
```

```
ce5                 3       0               0
ce5                 4       0               0
ce5                 5       0               0
ce5                 6       0               0
ce5                 7       0               0
ce7                 0       0               0
ce7                 1       0               0
ce7                 2       0               0
ce7                 3       0               0
ce7                 4       0               0
ce7                 5       0               0
ce7                 6       0               0
ce7                 7       0               0
ce8                 0       0               0
ce8                 1       0               0
ce8                 2       0               0
ce8                 3       0               0
ce8                 4       0               0
ce8                 5       0               0
ce8                 6       0               0
ce8                 7       0               0
VTEP2#
```

VTEP3

```
VTEP3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination       Status         Up/Down        Update
================================================================================
5.5.5.5           3.3.3.3           Installed      00:08:23       00:07:34
5.5.5.5           4.4.4.4           Installed      00:08:23       00:07:34

Total number of entries are 2
VTEP3#
VTEP3#
VTEP3#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
_____________________________________________________________________________________________________________________________________________
50002    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          3.3.3.3          --------------
50002    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          4.4.4.4          --------------
50003    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          3.3.3.3          --------------
50003    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          4.4.4.4          --------------
50006    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          3.3.3.3          --------------
50006    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          4.4.4.4          --------------
50006    ----         --       AC   xe4/4     --- Single Homed Port ---       6         ----      ----             ----
50007    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          3.3.3.3          --------------
50007    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          4.4.4.4          --------------
50007    ----         --       AC   xe4/4     --- Single Homed Port ---       7         ----      ----             ----
50009    ----         --       AC   xe4/4     --- Single Homed Port ---       9         ----      ----             ----
88888    ----         L3       NW   ----      ----                            ----      ----      5.5.5.5          3.3.3.3          5c07.5827.cb5f
88888    ----         L3       NW   ----      ----                            ----      ----      5.5.5.5          4.4.4.4          34ef.b696.dba1
```

```
Total number of entries are 13
VTEP3#

VTEP3#show bgp summary
BGP router identifier 5.5.5.5, local AS number 4294967205
BGP table version is 4
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
ce3(fe80::d6dc:85ff:fe1a:30b8)          4 4294967209         23         25       4      0      0  00:07:42             7
ce2(fe80::5e17:83ff:feff:305e)          4 4294967209         27         23       4      0      0  00:08:36             8

Total number of neighbors 2

Total number of Established sessions 2
VTEP3#

VTEP3#show bgp l2vpn evpn summary
BGP router identifier 5.5.5.5, local AS number 4294967205
BGP table version is 8
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
7.7.7.7          4 4294967209         71         56       5      0      0  00:07:42           46      6     28      8      0      4
8.8.8.8          4 4294967209         73         40       8      0      0  00:08:36           46      6     28      8      0      4

Total number of neighbors 2

Total number of Established sessions 2
VTEP3#
VTEP3#show priority-flow-control details all

Admin Configuration
--------------------
interface           mode  advertise willing   cap   link delay  priorities
                                                     allowance
================================================================================
xe4/4               on    on        off       8     0           0 1 2 3 4 5 6 7
ce2                 auto  on        on        8     0
ce3                 auto  on        on        8     0

--------------------------------------------------------------------------------

Operational Configuration
-------------------------
interface           state cap   link delay  priorities
                          allowance
================================================================================
```

```
xe4/4                on    8     0           0 1 2 3 4 5 6 7
ce2                  on    8     0           0 1 2 3 4 5 6 7
ce3                  on    8     0           0 1 2 3 4 5 6 7


VTEP3#show priority-flow-control statistics all
interface         pri   pause sent     pause received
=====================================================
xe4/4             0     0               0
xe4/4             1     0               0
xe4/4             2     0               0
xe4/4             3     0               0
xe4/4             4     0               0
xe4/4             5     0               0
xe4/4             6     0               0
xe4/4             7     0               0
ce2               0     0               0
ce2               1     0               0
ce2               2     0               0
ce2               3     0               0
ce2               4     0               0
ce2               5     0               0
ce2               6     0               0
ce2               7     0               0
ce3               0     0               0
ce3               1     0               0
ce3               2     0               0
ce3               3     0               0
ce3               4     0               0
ce3               5     0               0
ce3               6     0               0
ce3               7     0               0
VTEP3#
```

Spine 1

```
SPINE-1#show bgp l2vpn evpn summary
BGP router identifier 7.7.7.7, local AS number 4294967209
BGP table version is 12
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
3.3.3.3          4 4294967203        106         94      12      0      0  00:08:40            42      4     24      7      1      6
4.4.4.4          4 4294967204         76         59      10      0      0  00:08:41            67      4     49      7      1      6
5.5.5.5          4 4294967205         58         74      12      0      0  00:08:40            22      0     15      5      0      2

Total number of neighbors 3

Total number of Established sessions 3
SPINE-1#show bgp summary
BGP router identifier 7.7.7.7, local AS number 4294967209
```

```
BGP table version is 4
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
ethernet32/1                            4     0          0          0       0      0      0     never          Idle
ethernet63/1(fe80::5e07:58ff:fe27:cb64)  4 4294967203     27         25       3      0      0  00:08:50            3
ethernet31/1                            4     0          0          0       0      0      0     never          Idle
ethernet1/1(fe80::36ef:b6ff:fe96:dba2)   4 4294967204     28         25       4      0      0  00:08:50            3
ethernet6/1(fe80::5e07:58ff:fe70:4305)   4 4294967205     28         25       4      0      0  00:08:50            1

Total number of neighbors 5

Total number of Established sessions 3
SPINE-1#

SPINE-1#show interface counters rate gbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx gbps   |    Rx pps   |    Tx gbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 ethernet1/1          0.00           6             4.01           54843
 ethernet6/1          0.00           6             0.00           6
 ethernet63/1         4.01           54844         0.00           6
SPINE-1#


SPINE-1#show priority-flow-control details all

Admin Configuration
--------------------
interface           mode  advertise willing   cap   link delay  priorities
                                                    allowance
================================================================================
ethernet1/1         on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet2/1         on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet4           on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet3/1         on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet6/1         on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet11/1        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet25/1        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet25/2        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet30/1        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet61/1        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet62/1        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet32/1        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet31/1        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet64/1        on    on        off       8     0           0 1 2 3 4 5 6 7
ethernet63/1        on    on        off       8     0           0 1 2 3 4 5 6 7

--------------------------------------------------------------------------------

Operational Configuration
```

```
-------------------------
interface           state cap   link delay  priorities
                              allowance
================================================================================
ethernet1/1         on    8     0           0 1 2 3 4 5 6 7
ethernet2/1         on    8     0           0 1 2 3 4 5 6 7
ethernet4           on    8     0           0 1 2 3 4 5 6 7
ethernet3/1         on    8     0           0 1 2 3 4 5 6 7
ethernet6/1         on    8     0           0 1 2 3 4 5 6 7
ethernet11/1        on    8     0           0 1 2 3 4 5 6 7
ethernet25/1        on    8     0           0 1 2 3 4 5 6 7
ethernet25/2        on    8     0           0 1 2 3 4 5 6 7
ethernet30/1        on    8     0           0 1 2 3 4 5 6 7
ethernet61/1        on    8     0           0 1 2 3 4 5 6 7
ethernet62/1        on    8     0           0 1 2 3 4 5 6 7
ethernet32/1        on    8     0           0 1 2 3 4 5 6 7
ethernet31/1        on    8     0           0 1 2 3 4 5 6 7
ethernet64/1        on    8     0           0 1 2 3 4 5 6 7
ethernet63/1        on    8     0           0 1 2 3 4 5 6 7


SPINE-1#show priority-flow-control statistics all
interface         pri   pause sent     pause received
=====================================================
ethernet1/1       0     0              0
ethernet1/1       1     0              0
ethernet1/1       2     0              0
ethernet1/1       3     0              0
ethernet1/1       4     0              0
ethernet1/1       5     0              0
ethernet1/1       6     0              0
ethernet1/1       7     0              0
ethernet2/1       0     0              0
ethernet2/1       1     0              0
ethernet2/1       2     0              0
ethernet2/1       3     0              0
ethernet2/1       4     0              0
ethernet2/1       5     0              0
ethernet2/1       6     0              0
ethernet2/1       7     0              0
ethernet4         0     0              0
ethernet4         1     0              0
ethernet4         2     0              0
ethernet4         3     0              0
ethernet4         4     0              0
ethernet4         5     0              0
ethernet4         6     0              0
ethernet4         7     0              0
ethernet3/1       0     0              0
ethernet3/1       1     0              0
```

```
ethernet3/1       2     0               0
ethernet3/1       3     0               0
ethernet3/1       4     0               0
ethernet3/1       5     0               0
ethernet3/1       6     0               0
ethernet3/1       7     0               0
ethernet6/1       0     0               0
ethernet6/1       1     0               0
ethernet6/1       2     0               0
ethernet6/1       3     0               0
ethernet6/1       4     0               0
ethernet6/1       5     0               0
ethernet6/1       6     0               0
ethernet6/1       7     0               0
ethernet11/1      0     0               0
ethernet11/1      1     0               0
ethernet11/1      2     0               0
ethernet11/1      3     0               0
ethernet11/1      4     0               0
ethernet11/1      5     0               0
ethernet11/1      6     0               0
ethernet11/1      7     0               0
ethernet25/1      0     0               0
ethernet25/1      1     0               0
ethernet25/1      2     0               0
ethernet25/1      3     0               0
ethernet25/1      4     0               0
ethernet25/1      5     0               0
ethernet25/1      6     0               0
ethernet25/1      7     0               0
ethernet25/2      0     0               0
ethernet25/2      1     0               0
ethernet25/2      2     0               0
ethernet25/2      3     0               0
ethernet25/2      4     0               0
ethernet25/2      5     0               0
ethernet25/2      6     0               0
ethernet25/2      7     0               0
ethernet30/1      0     0               0
ethernet30/1      1     0               0
ethernet30/1      2     0               0
ethernet30/1      3     0               0
ethernet30/1      4     0               0
ethernet30/1      5     0               0
ethernet30/1      6     0               0
ethernet30/1      7     0               0
ethernet61/1      0     0               0
ethernet61/1      1     0               0
ethernet61/1      2     0               0
ethernet61/1      3     0               0
```

```
ethernet61/1      4     0               0
ethernet61/1      5     0               0
ethernet61/1      6     0               0
ethernet61/1      7     0               0
ethernet62/1      0     0               0
ethernet62/1      1     0               0
ethernet62/1      2     0               0
ethernet62/1      3     0               0
ethernet62/1      4     0               0
ethernet62/1      5     0               0
ethernet62/1      6     0               0
ethernet62/1      7     0               0
ethernet32/1      0     0               0
ethernet32/1      1     0               0
ethernet32/1      2     0               0
ethernet32/1      3     0               0
ethernet32/1      4     0               0
ethernet32/1      5     0               0
ethernet32/1      6     0               0
ethernet32/1      7     0               0
ethernet31/1      0     0               0
ethernet31/1      1     0               0
ethernet31/1      2     0               0
ethernet31/1      3     0               0
ethernet31/1      4     0               0
ethernet31/1      5     0               0
ethernet31/1      6     0               0
ethernet31/1      7     0               0
ethernet64/1      0     0               0
ethernet64/1      1     0               0
ethernet64/1      2     0               0
ethernet64/1      3     0               0
ethernet64/1      4     0               0
ethernet64/1      5     0               0
ethernet64/1      6     0               0
ethernet64/1      7     0               0
ethernet63/1      0     0               0
ethernet63/1      1     0               0
ethernet63/1      2     0               0
ethernet63/1      3     0               0
ethernet63/1      4     0               0
ethernet63/1      5     0               0
ethernet63/1      6     0               0
ethernet63/1      7     0               0
SPINE-1#
```

Spine 2

```
SPINE-2#show bgp summary
BGP router identifier 8.8.8.8, local AS number 4294967209
```

```
BGP table version is 6
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
cd7/1                                   4     0          0          0       0      0      0     never          Idle
cd27/1(fe80::5e07:58ff:fe27:cb65)       4 4294967203     28         30       6      0      0  00:10:17             3
cd32/1                                  4     0          0          0       0      0      0     never          Idle
cd3/1(fe80::36ef:b6ff:fe96:dba6)        4 4294967204     29         30       5      0      0  00:10:17             3
cd2/1(fe80::5e07:58ff:fe70:4304)        4 4294967205     27         31       6      0      0  00:10:17             1

Total number of neighbors 5

Total number of Established sessions 3
SPINE-2#show bgp l2vpn evpn summary
BGP router identifier 8.8.8.8, local AS number 4294967209
BGP table version is 13
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
3.3.3.3          4 4294967203         59         97      13      0      0  00:10:21           42      4     24      7      1      6
4.4.4.4          4 4294967204         64         91      11      0      0  00:10:21           67      4     49      7      1      6
5.5.5.5          4 4294967205         42         77      13      0      0  00:10:21           22      0     15      5      0      2

Total number of neighbors 3

Total number of Established sessions 3
SPINE-2#

SPINE-2#show interface counters rate gbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx gbps   |    Rx pps   |    Tx gbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 cd2/1                0.00           6             0.00           6
 cd3/1                0.00           6             4.01           54837
 cd27/1               4.01           54848         0.00           16
SPINE-2#


SPINE-2#show priority-flow-control details all

Admin Configuration
--------------------
interface           mode  advertise willing   cap   link delay  priorities
                                                     allowance
================================================================================
cd9/1               on    on        off       8     0           0 1 2 3 4 5 6 7
cd2/1               on    on        off       8     0           0 1 2 3 4 5 6 7
cd1/1               on    on        off       8     0           0 1 2 3 4 5 6 7
cd3/1               on    on        off       8     0           0 1 2 3 4 5 6 7
cd4/1               on    on        off       8     0           0 1 2 3 4 5 6 7
cd7/1               on    on        off       8     0           0 1 2 3 4 5 6 7
```

```
cd25/1                on    on        off       8     0           0 1 2 3 4 5 6 7
cd27/1                on    on        off       8     0           0 1 2 3 4 5 6 7
cd30/1                on    on        off       8     0           0 1 2 3 4 5 6 7
cd29/1                on    on        off       8     0           0 1 2 3 4 5 6 7
cd32/1                on    on        off       8     0           0 1 2 3 4 5 6 7
cd31/1                on    on        off       8     0           0 1 2 3 4 5 6 7

------------------------------------------------------------------------------

Operational Configuration
-------------------------
interface           state cap   link delay  priorities
                            allowance
==============================================================================
cd9/1                 on    8     0           0 1 2 3 4 5 6 7
cd2/1                 on    8     0           0 1 2 3 4 5 6 7
cd1/1                 on    8     0           0 1 2 3 4 5 6 7
cd3/1                 on    8     0           0 1 2 3 4 5 6 7
cd4/1                 on    8     0           0 1 2 3 4 5 6 7
cd7/1                 on    8     0           0 1 2 3 4 5 6 7
cd25/1                on    8     0           0 1 2 3 4 5 6 7
cd27/1                on    8     0           0 1 2 3 4 5 6 7
cd30/1                on    8     0           0 1 2 3 4 5 6 7
cd29/1                on    8     0           0 1 2 3 4 5 6 7
cd32/1                on    8     0           0 1 2 3 4 5 6 7
cd31/1                on    8     0           0 1 2 3 4 5 6 7


SPINE-2#show priority-flow-control statistics all
interface         pri   pause sent     pause received
=====================================================
cd9/1             0     0              0
cd9/1             1     0              0
cd9/1             2     0              0
cd9/1             3     0              0
cd9/1             4     0              0
cd9/1             5     0              0
cd9/1             6     0              0
cd9/1             7     0              0
cd2/1             0     0              0
cd2/1             1     0              0
cd2/1             2     0              0
cd2/1             3     0              0
cd2/1             4     0              0
cd2/1             5     0              0
cd2/1             6     0              0
cd2/1             7     0              0
cd1/1             0     0              0
cd1/1             1     0              0
cd1/1             2     0              0
```

```
cd1/1               3      0               0
cd1/1               4      0               0
cd1/1               5      0               0
cd1/1               6      0               0
cd1/1               7      0               0
cd3/1               0      0               0
cd3/1               1      0               0
cd3/1               2      0               0
cd3/1               3      0               0
cd3/1               4      0               0
cd3/1               5      0               0
cd3/1               6      0               0
cd3/1               7      0               0
cd4/1               0      0               0
cd4/1               1      0               0
cd4/1               2      0               0
cd4/1               3      0               0
cd4/1               4      0               0
cd4/1               5      0               0
cd4/1               6      0               0
cd4/1               7      0               0
cd7/1               0      0               0
cd7/1               1      0               0
cd7/1               2      0               0
cd7/1               3      0               0
cd7/1               4      0               0
cd7/1               5      0               0
cd7/1               6      0               0
cd7/1               7      0               0
cd25/1              0      0               0
cd25/1              1      0               0
cd25/1              2      0               0
cd25/1              3      0               0
cd25/1              4      0               0
cd25/1              5      0               0
cd25/1              6      0               0
cd25/1              7      0               0
cd27/1              0      0               0
cd27/1              1      0               0
cd27/1              2      0               0
cd27/1              3      0               0
cd27/1              4      0               0
cd27/1              5      0               0
cd27/1              6      0               0
cd27/1              7      0               0
cd30/1              0      0               0
cd30/1              1      0               0
cd30/1              2      0               0
cd30/1              3      0               0
cd30/1              4      0               0
```

```
cd30/1              5      0               0
cd30/1              6      0               0
cd30/1              7      0               0
cd29/1              0      0               0
cd29/1              1      0               0
cd29/1              2      0               0
cd29/1              3      0               0
cd29/1              4      0               0
cd29/1              5      0               0
cd29/1              6      0               0
cd29/1              7      0               0
cd32/1              0      0               0
cd32/1              1      0               0
cd32/1              2      0               0
cd32/1              3      0               0
cd32/1              4      0               0
cd32/1              5      0               0
cd32/1              6      0               0
cd32/1              7      0               0
cd31/1              0      0               0
cd31/1              1      0               0
cd31/1              2      0               0
cd31/1              3      0               0
cd31/1              4      0               0
cd31/1              5      0               0
cd31/1              6      0               0
cd31/1              7      0               0
SPINE-2#
```

ECN Validation

VTEP1

Before Congestion

This section shows the state of the VTEP1 switch under normal operation **before any congestion** was introduced.

Verify the status of established VxLAN tunnels. The output confirms that tunnels from the local VTEP1 (3.3.3.3) to remote VTEPs 5.5.5.5 (VTEP2) and 4.4.4.4 (VTEP3) are established.

```
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
3.3.3.3          5.5.5.5          Installed     02:02:14      02:02:02
3.3.3.3          4.4.4.4          Installed     02:02:13      02:02:02

Total number of entries are 2
```

Verify the detailed information about each configured VNID.

```
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
_____________________________________________________________________________________________________________________________________________
50002    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          5.5.5.5          --------------
50002    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          --------------
50002    ----         --       AC   xe0       --- Single Homed Port ---     2         ----      ----             ----
50003    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          5.5.5.5          --------------
50003    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          --------------
50003    ----         --       AC   xe0       --- Single Homed Port ---     3         ----      ----             ----
50004    ----         --       AC   xe0       --- Single Homed Port ---     4         ----      ----             ----
50006    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          5.5.5.5          --------------
50006    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          --------------
50006    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00 6         DF        ----             ----
50007    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          5.5.5.5          --------------
50007    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          --------------
50007    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00 7         NON-DF    ----             ----
50008    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          --------------
50008    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00 8         DF        ----             ----
50010    ----         L2       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          --------------
50010    ----         --       AC   xe0       --- Single Homed Port ---     10        ----      ----             ----
77777    ----         L3       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          34ef.b696.dba1
88888    ----         L3       NW   ----      ----                          ----      ----      3.3.3.3          5.5.5.5          5c07.5870.4300
88888    ----         L3       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          34ef.b696.dba1
101010   ----         L3       NW   ----      ----                          ----      ----      3.3.3.3          4.4.4.4          34ef.b696.dba1

Total number of entries are 21
VTEP1#
```

Verify the status of BGP sessions used for underlay reachability.

```
VTEP1#show bgp summary
BGP router identifier 3.3.3.3, local AS number 4294967203
BGP table version is 4
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
cd4/1(fe80::5e17:83ff:feff:3072)        4 4294967209        352        352       4      0      0  02:28:46             7
cd3/1(fe80::d6dc:85ff:fe1a:30f1)        4 4294967209        351        354       4      0      0  02:28:33             6

Total number of neighbors 2

Total number of Established sessions 2
VTEP1#
VTEP1#
```

Verify the BGP EVPN sessions used for overlay control plane (MAC/IP advertisement).

```
VTEP1#show bgp l2vpn evpn summary
BGP router identifier 3.3.3.3, local AS number 4294967203
BGP table version is 196
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
7.7.7.7          4 4294967209        523        647     196      0      0  02:28:35           246      4    221     12      1      8
8.8.8.8          4 4294967209        522        584     195      0      0  02:28:49           246      4    221     12      1      8

Total number of neighbors 2
```

```
Total number of Established sessions 2
VTEP1#
```

Display the configured QoS settings.

```
VTEP1#show running-config qos
qos enable
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
```

```
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
```

Display the current transmit (Tx) and receive (Rx) traffic rates on interfaces.

```
VTEP1#show interface counters rate gbps
+------------------+-------------+-------------+-------------+-------------+
|     Interface    |    Rx gbps  |    Rx pps   |    Tx gbps  |    Tx pps   |
+------------------+-------------+-------------+-------------+-------------+
 cd3/1               0.00          5            5.62          76770
 cd4/1               0.00          16           2.41          32915
 xe0                 7.98          109663       0.00          0
VTEP1#
```

Verify ECN marked packet counters. The output is empty or shows zero, indicating no ECN marking has occurred.

```
VTEP1#show interface counters ecn
+-----------------+----------------------+
|  Interface      |  ECN marked packets  |
+-----------------+----------------------+
VTEP1#
```

Verify the policy map applied on interface `cd3/1`.

The output shows detailed packet and byte counters for each queue within the policy map before congestion. The non-zero output counters for q0 and q7 confirms that the traffic is flowing and no drop occurred.

```
VTEP1#show policy-map interface cd3/1

Interface cd3/1

Type Queuing policy-map : default-out-policy

Service-policy (queuing) output: default-out-policy
Interface Bandwidth 100000000 kbps
-----------------------------------
Class-map (queuing): q0
 priority
 lossless
      output              : 1342328 packets, 12276931888 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q1
 priority
 lossless
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q2
 priority
 lossless
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q3
 priority
```

```
 lossless
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q4
 priority
 lossless
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q5
 priority
 lossless
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q6
 priority
 lossless
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q7
 priority
 lossless
      output              : 82 packets, 5762 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): mc-q0
      output       : 0 packets, 0 bytes
      dropped      : 0 packets, 0 bytes

Class-map (queuing): mc-q1
      output       : 0 packets, 0 bytes
      dropped      : 0 packets, 0 bytes

Class-map (queuing): mc-q2
      output       : 0 packets, 0 bytes
      dropped      : 0 packets, 0 bytes

Class-map (queuing): mc-q3
      output       : 2 packets, 196 bytes
      dropped      : 0 packets, 0 bytes

Wred/Tail Drop Statistics :
-----------------------
 green  : 0 packets
 yellow : 0 packets
 red    : 0 packets
```

```
VTEP1#
```

After Congestion

This section shows the state after congestion was induced by applying the `lossless_ecn_egress` policy map (which includes a 1 Gbps shaper per queue) to the egress network interface cd3/1.

Display the QoS configuration to verify the congestion applied on the interface `cd3/1`.

Congestion is induced manually in the fabric using ECN policy on egress interface `cd3/1`.

```
VTEP1#show running-config qos
qos enable
!
qos remark dscp

policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 1 gbps
```

```
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 1 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 1 gbps
  priority
```

```
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 3
00 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
interface cd3/1
 service-policy type queuing output lossless_ecn_egress
!
!
```

Verify the traffic is throttled on ingress as per egress shape configurations. The transmit rate on `cd3/1` is now capped at exactly 1.00 Gbps. Ingress traffic on `xe0` remains high (~7.98 Gbps), indicating congestion is occurring at the egress `cd3/1`.

```
VTEP1#show interface counters rate gbps
+------------------+-------------+-------------+-------------+-------------+
|     Interface    |    Rx gbps  |    Rx pps   |    Tx gbps  |    Tx pps   |
+------------------+-------------+-------------+-------------+-------------+
 cd3/1               0.00          5             1.00          76770
 cd4/1               0.00          16            2.41           32915
 xe0                 7.98          109663        0.00           0
VTEP1#
```

Verify the ECN marked packet counters (Incrementing)

The output shows a significant number of ECN marked packets (7,158,293). This confirms that due to congestion (caused by the shaper) and the applied ECN policy, packets egressing the interface `cd3/1` are being marked.

```
VTEP1#show interface counters ecn
+-----------------+----------------------+
|  Interface      |  ECN marked packets  |
+-----------------+----------------------+
 cd3/1 : 7158293 packets
```

Verify the policy map statistics on interface cd3/1 to show the WRED/Tail drops statistics incremented.

The output shows significant traffic output on queues q0 and q3.

```
VTEP1#show policy-map interface cd3/1

Interface cd3/1

Type Queuing policy-map : lossless_ecn_egress

Service-policy (queuing) output: lossless_ecn_egress
Interface Bandwidth 100000000 kbps
-----------------------------------
Class-map (queuing): q0
 shape 1 gbps
 priority
 lossless
```

```
 random-detect green min-threshold 500 packets max-threshold 600 packets drop-
probability 80 yellow min-threshold 300 packets max-threshold 400 packets drop-
probability 80 red min-threshold 100 packets max-threshold 200 packets drop-probability
80 weight 8 ecn
      output              : 7202722 packets, 65876095412 bytes
      dropped             : 7158293 packets, 65140466300 bytes

Class-map (queuing): q1
 shape 1 gbps
 priority
 lossless
 random-detect green min-threshold 500 packets max-threshold 600 packets drop-
probability 80 yellow min-threshold 300 packets max-threshold 400 packets drop-
probability 80 red min-threshold 100 packets max-threshold 200 packets drop-probability
80 weight 8 ecn
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q2
 shape 1 gbps
 priority
 lossless
 random-detect green min-threshold 500 packets max-threshold 600 packets drop-
probability 80 yellow min-threshold 300 packets max-threshold 400 packets drop-
probability 80 red min-threshold 100 packets max-threshold 200 packets drop-probability
80 weight 8 ecn
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q3
 shape 1 gbps
 priority
 lossless
 random-detect green min-threshold 500 packets max-threshold 600 packets drop-
probability 80 yellow min-threshold 300 packets max-threshold 400 packets drop-
probability 80 red min-threshold 100 packets max-threshold 200 packets drop-probability
80 weight 8 ecn
      output              : 3383878 packets, 30948948188 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q4
 shape 1 gbps
 priority
 lossless
 random-detect green min-threshold 500 packets max-threshold 600 packets drop-
probability 80 yellow min-threshold 300 packets max-threshold 400 packets drop-
probability 80 red min-threshold 100 packets max-threshold 200 packets drop-probability
80 weight 8 ecn
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q5
 shape 1 gbps
 priority
```

```
 lossless
 random-detect green min-threshold 500 packets max-threshold 600 packets drop-
probability 80 yellow min-threshold 300 packets max-threshold 400 packets drop-
probability 80 red min-threshold 100 packets max-threshold 200 packets drop-probability
80 weight 8 ecn
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q6
 shape 1 gbps
 priority
 lossless
 random-detect green min-threshold 500 packets max-threshold 600 packets drop-
probability 80 yellow min-threshold 300 packets max-threshold 400 packets drop-
probability 80 red min-threshold 100 packets max-threshold 200 packets drop-probability
80 weight 8 ecn
      output              : 0 packets, 0 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): q7
 shape 1 gbps
 priority
 lossless
 random-detect green min-threshold 500 packets max-threshold 600 packets drop-
probability 80 yellow min-threshold 300 packets max-threshold 400 packets drop-
probability 80 red min-threshold 100 packets max-threshold 200 packets drop-probability
80 weight 8 ecn
      output              : 2296 packets, 188556 bytes
      dropped             : 0 packets, 0 bytes

Class-map (queuing): mc-q0
      output       : 0 packets, 0 bytes
      dropped      : 0 packets, 0 bytes

Class-map (queuing): mc-q1
      output       : 0 packets, 0 bytes
      dropped      : 0 packets, 0 bytes

Class-map (queuing): mc-q2
      output       : 0 packets, 0 bytes
      dropped      : 0 packets, 0 bytes

Class-map (queuing): mc-q3
      output       : 49 packets, 4818 bytes
      dropped      : 0 packets, 0 bytes

Wred/Tail Drop Statistics :
-----------------------
 green  : 7158293 packets
 yellow : 0 packets
 red    : 0 packets
VTEP1#
```

The above Wred/Tail Drop Statistics output shows that count exactly matches the ECN marked packets counter (counter from the previous command, cd3/1 : 7158293 packets), indicating these are packets marked due to crossing the green WRED/ECN threshold. This comparison helps to confirm that ECN counter and WRED statistics within the policy map under congestion successfully triggers ECM marking.

Note: It is important to capture and check the packets on egress side are marked with CE bits during congestion.

VTEP2

```
VTEP2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination       Status         Up/Down        Update
================================================================================
4.4.4.4           3.3.3.3           Installed      00:03:53       00:03:04
4.4.4.4           5.5.5.5           Installed      00:03:53       00:03:04

Total number of entries are 2
VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
_____________________________________________________________________________________________________________________________________
50002    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50002    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          --------------
50002    ----         --       AC   xe34      --- Single Homed Port ---       2         ----      ----             ----
50003    ----         L2       NW   ----      ----                          --VTEP3--     ----      4.4.4.4          3.3.3.3          --------------
50003    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          --------------
50003    ----         --       AC   xe34      --- Single Homed Port ---       3         ----      ----             ----
50005    ----         --       AC   xe34      --- Single Homed Port ---       5         ----      ----             ----
50006    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50006    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          --------------
50006    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   6         NON-DF    ----             ----
50007    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50007    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          --------------
50007    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   7         DF        ----             ----
50008    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50008    ----         --       AC   po1       00:00:00:00:00:33:33:00:00:00   8         NON-DF    ----             ----
50010    ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          --------------
50010    ----         --       AC   xe34      --- Single Homed Port ---       10        ----      ----             ----
77777    ----         L3       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          5c07.5827.cb5f
88888    ----         L3       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          5c07.5827.cb5f
88888    ----         L3       NW   ----      ----                            ----      ----      4.4.4.4          5.5.5.5          5c07.5870.4300
101010   ----         L3       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3          5c07.5827.cb5f

Total number of entries are 21

VTEP2#
VTEP2#show bgp summary
BGP router identifier 4.4.4.4, local AS number 4294967204
BGP table version is 6
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
ce5(fe80::5e17:83ff:feff:3060)          4 4294967209         15         14       6      0      0  00:04:06             7
ce1(fe80::d6dc:85ff:fe1a:30b0)          4 4294967209         12         14       6      0      0  00:03:13             6
```

```
Total number of neighbors 2

Total number of Established sessions 2
VTEP2#show bgp l2vpn evpn summary
BGP router identifier 4.4.4.4, local AS number 4294967204
BGP table version is 12
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
7.7.7.7          4 4294967209        46         61       6      0      0  00:03:16           56      4     33     10      1      8
8.8.8.8          4 4294967209        76         48      12      0      0  00:04:08           56      4     33     10      1      8

Total number of neighbors 2

Total number of Established sessions 2
VTEP2#

VTEP2#show interface counters rate gbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx gbps   |    Rx pps   |    Tx gbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 ce1                  4.01           54787         0.00           6
 ce5                  4.01           54787         0.00           6
 xe34                 0.00           0             7.98           109564
VTEP2#

VTEP2#show running-config qos
qos enable
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
```

```
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
```

```
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
!
VTEP2#
VTEP2#
VTEP2#show interface counters ecn
+------------------+----------------------+
|  Interface       |  ECN marked packets  |
+------------------+----------------------+
VTEP2#
```

VTEP3

```
VTEP3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
5.5.5.5          3.3.3.3          Installed     00:08:23      00:07:34
5.5.5.5          4.4.4.4          Installed     00:08:23      00:07:34

Total number of entries are 2
VTEP3#
VTEP3#
VTEP3#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr         Router-Mac
-------------------------------------------------------------------------------------------------------------------------------------------
50002    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          3.3.3.3          --------------
50002    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          4.4.4.4          --------------
50003    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          3.3.3.3          --------------
50003    ----         L2       NW   ----      ----                            ----      ----      5.5.5.5          4.4.4.4          --------------
```

```
50006    ----        L2      NW   ----     ----                        ----     ----     5.5.5.5         3.3.3.3         --------------
50006    ----        L2      NW   ----     ----                        ----     ----     5.5.5.5         4.4.4.4         --------------
50006    ----        --      AC   xe4/4    --- Single Homed Port ---    6        ----     ----            ----
50007    ----        L2      NW   ----     ----                        ----     ----     5.5.5.5         3.3.3.3         --------------
50007    ----        L2      NW   ----     ----                        ----     ----     5.5.5.5         4.4.4.4         --------------
50007    ----        --      AC   xe4/4    --- Single Homed Port ---    7        ----     ----            ----
50009    ----        --      AC   xe4/4    --- Single Homed Port ---    9        ----     ----            ----
88888    ----        L3      NW   ----     ----                        ----     ----     5.5.5.5         3.3.3.3         5c07.5827.cb5f
88888    ----        L3      NW   ----     ----                        ----     ----     5.5.5.5         4.4.4.4         34ef.b696.dba1

Total number of entries are 13
VTEP3#
VTEP3#
VTEP3#
VTEP3#show bgp summary
BGP router identifier 5.5.5.5, local AS number 4294967205
BGP table version is 4
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS      MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
ce3(fe80::d6dc:85ff:fe1a:30b8)           4 4294967209     23         25       4      0      0              00:07:42        7
ce2(fe80::5e17:83ff:feff:305e)           4 4294967209     27         23       4      0      0              00:08:36        8

Total number of neighbors 2

Total number of Established sessions 2
VTEP3#
VTEP3#
VTEP3#show bgp l2vpn evpn summary
BGP router identifier 5.5.5.5, local AS number 4294967205
BGP table version is 8
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down   State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
7.7.7.7          4 4294967209     71         56       5      0      0   00:07:42            46      6     28      8      0      4
8.8.8.8          4 4294967209     73         40       8      0      0   00:08:36            46      6     28      8      0      4

Total number of neighbors 2
Total number of Established sessions 2
VTEP3#


VTEP3#show running-config qos
qos enable
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
```

```
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
```

```
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
!
VTEP3#
VTEP3#
VTEP3#show interface counters ecn
+------------------+----------------------+
|  Interface       |  ECN marked packets  |
+------------------+----------------------+
VTEP3#
```

Spine1

```
SPINE-1#show bgp l2vpn evpn summary
BGP router identifier 7.7.7.7, local AS number 4294967209
BGP table version is 12
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS       MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down   State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
3.3.3.3          4 4294967203     106         94      12      0      0   00:08:40           42      4     24      7      1      6
4.4.4.4          4 4294967204      76         59      10      0      0   00:08:41           67      4     49      7      1      6
5.5.5.5          4 4294967205      58         74      12      0      0   00:08:40           22      0     15      5      0      2
```

```
Total number of neighbors 3

Total number of Established sessions 3
SPINE-1#show bgp summary
BGP router identifier 7.7.7.7, local AS number 4294967209
BGP table version is 4
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
ethernet32/1                            4     0          0          0       0      0      0     never          Idle
ethernet63/1(fe80::5e07:58ff:fe27:cb64)  4 4294967203     27         25       3      0      0  00:08:50             3
ethernet31/1                            4     0          0          0       0      0      0     never          Idle
ethernet1/1(fe80::36ef:b6ff:fe96:dba2)   4 4294967204     28         25       4      0      0  00:08:50             3
ethernet6/1(fe80::5e07:58ff:fe70:4305)   4 4294967205     28         25       4      0      0  00:08:50             1

Total number of neighbors 5

Total number of Established sessions 3
SPINE-1#

SPINE-1#show interface counters rate gbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx gbps   |    Rx pps   |    Tx gbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 ethernet1/1          0.00           6             4.01           54843
 ethernet6/1          0.00           6             0.00           6
 ethernet63/1         4.01           54844         0.00           6
SPINE-1#


SPINE-1#show running-config qos
qos enable
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
```

```
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
```

```
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q6
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
!

SPINE-1#show interface counters ecn
+------------------+----------------------+
|  Interface       |  ECN marked packets  |
+------------------+----------------------+
SPINE-1#
```

Spine2

```
SPINE-2#show bgp summary
BGP router identifier 8.8.8.8, local AS number 4294967209
BGP table version is 6
4 BGP AS-PATH entries
0 BGP community entries
10  Configured ebgp ECMP multipath: Currently set at 10

Neighbor                                V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
cd7/1                                   4     0          0          0       0      0      0     never          Idle
cd27/1(fe80::5e07:58ff:fe27:cb65)       4 4294967203     28         30       6      0      0  00:10:17             3
cd32/1                                  4     0          0          0       0      0      0     never          Idle
cd3/1(fe80::36ef:b6ff:fe96:dba6)        4 4294967204     29         30       5      0      0  00:10:17             3
cd2/1(fe80::5e07:58ff:fe70:4304)        4 4294967205     27         31       6      0      0  00:10:17             1

Total number of neighbors 5

Total number of Established sessions 3
SPINE-2#show bgp l2vpn evpn summary
BGP router identifier 8.8.8.8, local AS number 4294967209
```

```
BGP table version is 13
4 BGP AS-PATH entries
0 BGP community entries

Neighbor         V    AS       MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down   State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
3.3.3.3          4 4294967203       59         97      13      0      0   00:10:21           42      4     24      7      1      6
4.4.4.4          4 4294967204       64         91      11      0      0   00:10:21           67      4     49      7      1      6
5.5.5.5          4 4294967205       42         77      13      0      0   00:10:21           22      0     15      5      0      2

Total number of neighbors 3

Total number of Established sessions 3
SPINE-2#

SPINE-2#show interface counters rate gbps
+--------------------+--------------+-------------+--------------+-------------+
|     Interface      |    Rx gbps   |    Rx pps   |    Tx gbps   |    Tx pps   |
+--------------------+--------------+-------------+--------------+-------------+
 cd2/1                 0.00           6             0.00           6
 cd3/1                 0.00           6             4.01           54837
 cd27/1                4.01           54848         0.00           16
SPINE-2#


SPINE-2#show running-config qos
qos enable
!
policy-map type queuing default default-out-policy
 class type queuing default q0
  priority
  lossless
  exit
 class type queuing default q1
  priority
  lossless
  exit
 class type queuing default q2
  priority
  lossless
  exit
 class type queuing default q3
  priority
  lossless
  exit
 class type queuing default q4
  priority
  lossless
  exit
 class type queuing default q5
  priority
  lossless
  exit
```

```
 class type queuing default q6
  priority
  lossless
  exit
 class type queuing default q7
  priority
  lossless
  exit
!
policy-map type queuing default lossless_ecn_egress
 class type queuing default q0
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q1
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q2
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q3
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q4
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q5
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
```

```
 class type queuing default q6
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
 class type queuing default q7
  shape 10 gbps
  priority
  lossless
  random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-
threshold 400 red min-threshold 100 max-threshold 200 packets ecn
  exit
!
!

SPINE-2#show interface counters ecn
+------------------+----------------------+
|  Interface       |  ECN marked packets  |
+------------------+----------------------+
SPINE-2#
```

Implementation Examples

Here are examples illustrating how the configured features (PFC, ECN, QoS policies) would function in a real-time scenario within the described VxLAN fabric:

Example 1: PFC Handling Congestion

Scenario: A GPU server (SH1) connected to Leaf1's access port (`xe0`) sends a burst of high-priority RDMA traffic (e.g., DSCP mapped to queue 3) towards another server connected to Leaf2 (VTEP2). This traffic traverses Spine1.

Congestion Point: The egress port on Spine1 towards Leaf2 (e.g., `ethernet31/1)` experiences temporary congestion for queue 3.

Real-time Actions:

1. Buffer Threshold: As packets build up in queue 3 on `ethernet31/1`, the buffer occupancy crosses the configured PFC threshold for that priority.
2. PFC Pause: Spine1 generates an IEEE 802.3X PFC pause frame specifically for priority 3 and sends it back towards the source of the traffic (in this case, upstream towards Leaf1 via interface `ethernet63/1`).
3. Transmission Pause: Leaf1, upon receiving the PFC pause frame on its interface connected to Spine1 (e.g., cd4/1 - inferred connection based on IP), temporarily stops transmitting traffic only for priority 3 towards Spine1. Traffic for other priorities (0-2, 4-7) continues flowing if the link isn't paused for those priorities.
4. Buffering: Packets for priority 3 arriving at Spine1 during the pause are buffered in shared memory.
5. Resume: Once the congestion on Spine's egress port clears and buffer levels drop below the resume threshold, Spine1 stops sending pause frames (or sends a resume frame, depending on implementation). Leaf1 then resumes sending priority 3 traffic.

Outcome: Packet loss for the high-priority RDMA traffic is prevented, maintaining lossless transport critical for the AI/ML workload.

Example 2: ECN Marking and Propagation

Scenario: Similar to Example 1, RDMA traffic flows from SH1 (Leaf1) to SH2 (Leaf2) via Spine1. The traffic is marked with a DSCP value mapped to queue 3.

Congestion Point: The egress port queue 3 on Spine1 (ethernet31/1) starts experiencing congestion, but before the PFC threshold is hit.

Real-time Actions:

1. WRED Threshold: As packets queue up, the average queue depth crosses the configured random-detect ... min-threshold for ECN marking within the lossless_ecn_egress policy-map (assuming it's applied).
2. ECN Marking (Transit): Spine1 marks the Congestion Experienced (CE) codepoint in the outer IP header (DSCP field's ECN bits) of the VxLAN-encapsulated packets passing through ethernet31/1. It does not drop the packet.
3. ECN Propagation (Decapsulation): Leaf2 (the destination VTEP) receives the VxLAN packet marked with CE. During decapsulation, it copies the CE marking from the outer IP header to the inner IP header of the original RDMA packet.
4. Endpoint Notification: The destination server (SH2) receives the RDMA packet with the CE marking. Following RoCEv2 or TCP ECN standards, it signals back to the source server (SH1) via its next packet/acknowledgment that congestion was encountered.
5. Rate Reduction: The source server (SH1), upon receiving the congestion signal, reduces its transmission rate.

Outcome: Congestion is signaled before buffers become full enough to trigger a potentially disruptive PFC pause. The endpoints proactively reduce speed, alleviating congestion and improving overall throughput and latency compared to relying solely on PFC. ECN marking is consistently handled across the VxLAN tunnel from encapsulation to decapsulation.

Troubleshooting

Here are troubleshooting tips for PFC and ECN over VxLAN:

Verify Global QoS and Policy Map Configuration

- Is QoS Enabled?: Ensure `qos enable` is configured globally on all relevant devices (VTEPs, Spines).
- Is the Correct Policy Map Defined?: Confirm the `lossless_ecn_egress` policy map is defined correctly with `lossless` properties and `random-detect ...` packets `ecn` settings for the required queues (typically q0-q7).
- Is the Policy Map Applied?: Crucially, verify that the `lossless_ecn_egress` policy map is actually applied either globally or to the relevant interfaces (uplinks, access ports). Use `show running-config | include service-policy` or similar commands. If it's not applied, the defined ECN/lossless behavior will not be active.

Check DSCP/CoS Mappings and Trust Settings

- Is Traffic Mapped Correctly?: Verify the mapping from incoming DSCP values (for L3/IP traffic like RoCE) or CoS values (for L2 traffic) to the correct hardware Traffic Classes (TCs)/queues3. Ensure the TCs corresponding to lossless traffic (e.g., RDMA) are mapped to the queues configured as lossless in the policy map. Use commands like `show qos map dscp-tc`.
- Is Traffic Class Mapped to Priority Group (PG)?: Confirm that the relevant Traffic Classes are mapped to the appropriate Priority Groups (PGs) that PFC will act upon. Use `show qos map tc-pg`.
- Are Markings Trusted?: Ensure interfaces receiving prioritized traffic are configured to trust the incoming markings (e.g., `qos trust dscp` on routed ports or IRB interfaces, `qos trust cos` on L2 ports if applicable). Check relevant interface configurations.

Troubleshoot Priority Flow Control (PFC) Issues

- Is PFC Enabled Operationally?: Check the PFC status on interfaces where it's configured (`priority-flow-control mode on`). Use `show priority-flow-control interface <interface_name>`. Look for Operational Status: `On` and ensure the expected priorities (0-7 in the examples) are enabled.
- Are Pause Frames Being Sent/Received?: Check PFC counters using s`how priority-flow-control interface <interface_name> counters`.
  - Tx Pause: Indicates the local device is congested and sending pauses upstream.
  - Rx Pause: Indicates a downstream device is congested and sending pauses to this device. Analyze counters end-to-end to pinpoint congestion.
- Is LLDP Running?: PFC negotiation often relies on DCBX exchanged via LLDP. Ensure LLDP is globally enabled (lldp run) and active on the relevant interfaces (lldp-agent). While mode `on` might force PFC, LLDP issues can sometimes interfere.
- Check MTU: Ensure consistent MTU settings across the entire path (access, leaf, spine). Mismatched MTUs can cause drops, especially for large RDMA frames, potentially triggering excessive PFC.

Troubleshoot Explicit Congestion Notification (ECN) Issues

- Check ECN Thresholds Values: Verify the random-detect thresholds within the applied lossless_ecn_egress policy map. Ensure they are set appropriately relative to PFC thresholds to allow ECN marking before excessive pausing occurs.

- Check Queues Crossing Thresholds: During congestion, monitor queue depths (`show qos queue interface <if_name>`) to see if they exceed the configured ECN minimum thresholds.
- Is ECN Marking Happening?:
  - Use SPAN/ERSPAN to capture traffic after a potential congestion point (e.g., Spine egress). Check the outer IP header's DSCP/ECN field for the Congestion Experienced (CE) codepoint (binary 11).
- Is ECN Propagating Through VxLAN?:
  - Capture traffic on the egress VTEP before decapsulation. Verify the outer IP header has the CE marking.
  - Capture traffic after decapsulation on the egress VTEP. Verify the inner IP header now carries the CE marking.
- Are Endpoints Reacting?: Check server-side monitoring tools. Confirm the destination server is receiving CE-marked packets and signaling congestion back (e.g., via ECE bit in TCP ACKs or specific RoCEv2 mechanisms), and that the source server is reducing its sending rate.

Verify Underlay Network and VxLAN Health

- Check Basic Connectivity: Ensure VTEP loopbacks are reachable via the underlay (use standard `ping/traceroute`). Verify BGP EVPN and underlay routing protocol sessions are stable (`show ip bgp summary`, `show ip ospf neighbor`, etc.).
- Check Load Balancing: Ensure traffic is being distributed across ECMP paths as expected (`show load-balance rtag7 effective`). Uneven hashing could overload specific links, triggering PFC/ECN.

Glossary

The following table provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
AI/ML	Artificial Intelligence / Machine Learning. Workloads demanding high bandwidth and low latency, often using RDMA/RoCE.x‘x‘
CoS	Class of Service. Used for selecting queues in L2 forwarding.
DSCP	Differentiated Services Code Point. Used for selecting queues in L3 forwarding. OcNOS supports DSCP-to-Queue mapping profiles.
VxLAN	Virtual Extensible LAN → A tunneling protocol that encapsulates Layer 2 traffic into Layer 3 packets, enabling large-scale virtual networks.
ECN	Explicit Congestion Notification. Allows network endpoints to signal congestion without dropping packets. It marks packets with congestion bits when queue thresholds are crossed. Works in L3 VxLAN transport, requiring support at VTEPs and transit nodes
GPU	Graphics Processing Unit. Large clusters are interconnected for distributed AI/ML training
IRB	Integrated Routing and Bridging. Supported for L3 VxLAN forwarding.

PFC	Priority Flow Control. Pauses traffic in specific priority queues during congestion to prevent packet loss, allowing other traffic to continue. Maps packets to Priority Groups (PGs) and sends XOFF frames when shared memory limits are exceeded. Works in both L2 and L3 VxLAN transport.
PG	Priority Group. PFC maps packets to PGs based on shared memory limits. Buffers packets during PFC pauses. Up to eight profiles are supported.
RDMA	Remote Direct Memory Access. Used by AI/ML workloads for high bandwidth, low-latency interconnects.
RoCE	RDMA over Converged Ethernet. A protocol used by AI/ML workloads.
VTEP	VxLAN Tunnel Endpoint. Nodes that encapsulate/decapsulate VxLAN packets. ECN functionality is mandatory on VTEP nodes for L3 VxLAN
VxLAN	Virtual Extensible LAN. A tunneling protocol encapsulating L2 traffic into L3 packets for large-scale virtual networks.
WRED	Weighted Random Early Detection. Thresholds used by ECN to determine when to start marking packets for congestion.

.

VxLAN - Operation Administration Maintenance

CHAPTER 1 VxLAN Operation Administration Maintenance

Overview

In Data centers, where Leaf and Spine CLOS network topology is deployed, L2 and L3 VxLAN overlay packets gets forwarded over IP underlay. Traditional monitoring tools such as ping, trace-route could only provide information related to underlay path. For both reactive and proactive fault management of overlay networks, there arises a need to have Ethernet OAM like functions available for the overlay networks.

Like traditional Ethernet, the VxLAN networks need OAM (Operations, Administration, and Maintenance) services to verify connectivity, continuity, and fault isolation between Leaf switches. To achieve this, VxLAN OAM functionality is implemented in the OcNOS 7.0 release. It defines MEPs (Maintenance End Points) and MIPs (Maintenance Intermediate Points) within the VxLAN Leaf and Spine architecture.

Leafs are the nodes where the VxLAN tunnel endpoints (VTEPs) are logically located. It encapsulates standard Ethernet traffic from hosts into VxLAN packets for transport across the IP underlay and decapsulates them upon arrival. The document uses Leaf and VTEP terms interchangeably.

The Spine switch operates purely at the IP underlay level, utilizing routing protocols like OSPF and BGP to route encapsulated packets to the destination VTEPs. So, it is viewed as a transit node in the underlay path.

VxLAN Maintenance End Point (MEP)

A VxLAN MEP is a logical endpoint of an OAM Maintenance Entity (ME) located at a VTEP. It is the OAM anchor point at the VTEP edge of a VxLAN segment. It is responsible for generating, transmitting, and terminating VxLAN OAM frames (like continuity check message, ping and pathtrace). VxLAN MEPs are uniquely identified using the VxLAN Network Identifier (VNI) and remote VTEP IP. The OAM feature works on the Down MEP.

VxLAN Maintenance Intermediate Point (MIP)

A VxLAN MIP is a passive point inside the VxLAN underlay path (typically spine node). It does not originate OAM frames but can respond to certain OAM requests (like `pathtrace`). It is an intermediate transit point that allows operators to trace the underlay path hop by hop. The MIPs placement is conceptually inside the VxLAN tunnel between VTEPs.

Feature Characteristics

In Leaf and Spine CLOS network topology, Equal-Cost Multi-Path (ECMP) routes exist to and from each leaf and spine switch, ensuring multiple redundant paths for traffic forwarding. Thus the VxLAN data packets traverse through multiple ECMP paths. Identifying the exact point of failure in the path is often complex. The CLI commands introduced in VxLAN OAM feature are designed to overcome these challenges. This feature provides the following OAM services to verify loopback, pathtrace and continuity check message (CCM) exchanges between MEPs:

- **Fault Verification (`loopback/ping`) :** This function utilizes loopback messages (`ping`) to verify basic VTEP reachability and tunnel connectivity between leaf nodes. These packets traverse the spine or MIP nodes, which perform standard IP routing without engaging in OAM path discovery. The spine simply forwards the packet to the next hop.
- **Fault Isolation (`pathtrace`):** This function utilizes trace messages (`pathtrace`) to identify and isolate the exact path of a flow between leaf nodes. The spine or MIP node plays an active role in this process by uplifting the pathtrace packets having a TTL of 1 and responding, which allows operators to trace the underlay path hop-by-hop.
- **Loss of Continuity Detection:** Performs one-way connectivity checks to detect any loss of communication between endpoints. The interval for the CCM message is dependent on the CPU load and CCM message scale,

the minimum interval (resolution) supported is 10 ms, aligning with existing Ethernet OAM behavior in non-scaled environments. If ECMP paths exist to a remote leaf, all paths are probed, and their connectivity states are reported individually.

- OAM functionality, including CCM message handling, is implemented in the control plane (i.e., processed by the CPU), thus the CCM packet intervals may not be honoured as configured.
- The Round Trip Time (RTT) values reported in the `ping` and `pathtrace` output are measured by the control plane module, so it is not accurate link-level RTT.
- It supports both static or dynamic VxLAN tunnels for single homing or multi-homing setup and is agnostic to the underlay protocols used in establishing routing or tunnels.

Limitation:

- VxLAN for IRB or Layer 3 services are not supported.
- VxLAN over IPv6 is not supported.
- Only one active `ping` instance is allowed per MEP at a time.
- Only one active `pathtrace` instance is allowed per MEP at a time.
- The `ping` or `pathtrace` operations to the local VTEP IP address are not supported.
- It is supported on the Trident 3 switch: AS7726-32X, AS7326-56X, AS5835-54X only.
- It is supported on the Qumran 2a switch: AS7535-28XB only.

VxLAN OAM Packet Flow

In the VxLAN network topology, server nodes are connected to leaf switches on the access side, while the uplink ports of the leaf switches are connected to spine nodes.

Figure 1-9: VxLAN OAM Packet Flow

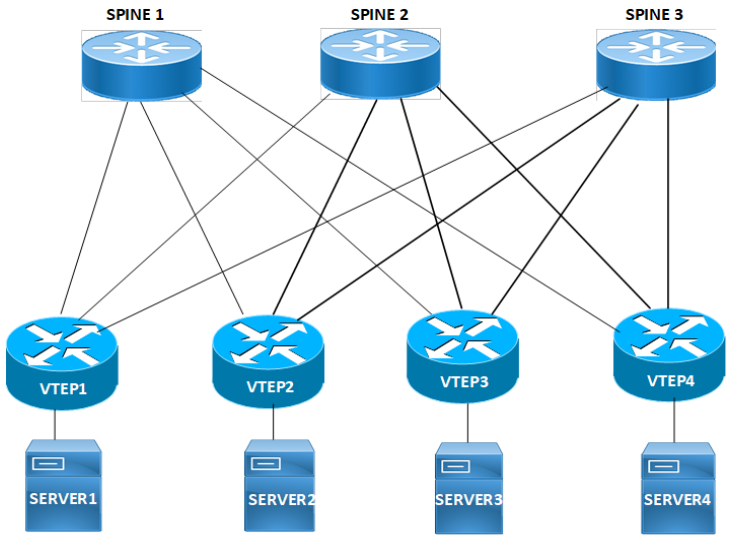

Following is the flow of VxLAN OAM Packet communication:

1. The source leaf generates a OAM packet targeting the destination VTEP.
2. The source leaf encapsulates the OAM packet into the full VxLAN tunnel header.

    1. Host side L2 parameters supplied through Ping and Pahtrace commands are encapsulated into the inner L2 header. And these parameters are not applicable to the Continuity Check Message operation.

3. The encapsulated OAM packet is forwarded across the underlay network. The Packet forwarding is based on the installed routes and the load-balancing configuration.

4. The target leaf receives the packet and decapsulates the VxLAN tunnel headers. It identifies the inner packet as a VxLAN OAM message, and sends it to the OAM process.

    1. For interactive checks like `ping` or `pathtrace`, the target leaf generates a corresponding OAM response packet.

    2. CCM messages are consumed by the process and resets the CCM timers.

5. This response is then encapsulated back into a VxLAN tunnel header and sent back through the underlay.

VxLAN OAM Modes

The CFM objects (Maintenance Domain (MD) or Maintenance Association (MA) and Maintenance Endpoint (MEP) / Maintenance Intermediate Point (MIP)) required for the VxLAN OAM operations are created by default or explicitly through the configurations. The default CFM objects are created when the feature is enabled. And the explicit objects are created as part of the CFM module configurations. Both the implicit and explicit configurations can co-exist in the system at the same time. However, it is the VxLAN OAM operations which decides the implicit or explicit mode functioning, based on the supplied CLI parameters.

Note: CCM operation is performed only on the explicit CFM objects.

Below Table 1-1 explains the VxLAN OAM mode, configurations and the allowed operations.

Table 1-1: VxLAN OAM Modes

	Implicit Mode	**Explicit Mode**

Leaf Configurations	Creates MA, MD, and MEP CFM objects implicitly
	`nvo vxlan-oam`
	
	For details, refer to nvo vxlan-oam CLI section.
		Creates MA, MD, and MIP CFM objects explicitly. Refer to ethernet cfm domain-type and ethernet cfm mep CLI sections for more information.
		
		`nvo vxlan-oam`
		`ethernet cfm domain-type character-string domain-name testmd level 3 mip-creation none`
		`service ma-type string ma-name testma`
		`evpn 2`
		`ethernet cfm mep down mpid 12 active true remote-vtep 1.1.1.1`

Spine Configurations	Creates MA, MD, and MIP CFM objects implicitly.
	
	`nvo vxlan-oam spine.`
	
	For details, refer to nvo vxlan-oam CLI section.
		Creates MA, MD, and MIP CFM objects explicitly. Refer to ethernet cfm domain-type and ethernet cfm mep CLI sections for more information.
		
		`nvo vxlan-oam spine`
		`ethernet cfm domain-type character-string domain-name testmd level 3 mip-creation static`
		`service ma-type string ma-name testma`

Table 1-1: VxLAN OAM Modes

	Implicit Mode	Explicit Mode
Ping and Pathtrace	It supports only `ping` and `pathtrace` operations and when the CFM object parameters are not supplied, these operations are performed in this mode • Below are the examples without CFM object parameters. • `ping ethernet evpn destination 1.1.1.1 2 flow-inport xe1 flow-dst 0010.7700.0001 flow-vlan 102 flowethtype 0x8100` • `pathtrace ethernet evpn destination 1.1.1.1 2 flow-inport xe1 flow-dst 0010.7700.0001 flow-vlan 102 flowethtype 0x8100` For details, refer to ping ethernet evpn and pathtrace ethernet evpn CLIs section.	It supports `ping` and `pathtrace` operations. The parameters provided in the `ping` and `pathtrace` CLI command chooses the mode. Below are the examples with CFM object parameters. • `ping ethernet evpn 12 domain testdomain ma testma destination 1.1.1.1 2 flow-inport xe1 flow-dst 0010.7700.0001 flow-vlan 102 flow-ethtype 0x8100` • `pathtrace ethernet evpn 12 domain testdomain ma testma destination 1.1.1.1 2 flow-inport xe1 flow-dst 0010.7700.0001 flow-vlan 102 flow-ethtype 0x8100` Above CLIs carry the MEP ID `12`, MA name `testma`, and the MD name `testdomain`.

CCM	Not applicable.	It supports CCM operation. Below is the example with MA/MD/MEP objects.
		`nvo vxlan-oam`
		`ethernet cfm domain-type characterstring domain-name testmd level 3 mipcreation none`
		`service ma-type string ma-name testma`
		`evpn 2`
		`mep crosscheck mpid 21`
		`ethernet cfm mep down mpid 12 active true remote-vtep 1.1.1.1`
		`cc multicast state enable`

Using Same Path for VxLAN OAM and Data Packets

For accurate packet path debugging, OAM packets (`ping` and `pathtrace`) must follow the same forwarding logic as host data packets. The OAM packets are generated directly on the VTEP node while VxLAN data typically originates from a host. Since the VTEP generates OAM packets locally, users can supply the flow parameters (such as ingress port, destination MAC, and VLAN) in the `ping` and `pathtrace` CLIs manually by copying the parameters of the host device to simulate a specific host's traffic.

Following are the example that shows the usage of flow parameters in `ping` and `pathtrace` CLIs.

`pathtrace ethernet evpn 12 domain testdomain ma testma destination 1.1.1.1 2 flow-inport xe1` **`flow-dst 0010.7700.0001 flow-vlan 102 flow-ethtype 0x8100`**

`ping ethernet evpn 12 domain testdomain ma testma destination 1.1.1.1 2` **`flow-inport xe1 flow-dst 0010.7700.0001 flow-vlan 102 flow-ethtype 0x8100`**

Benefits

It provides the capability of OAM operation for VxLAN networks.

Prerequisites

Ensure the following prerequisites are met before configuring VxLAN OAM functionality:

- Enable VxLAN hardware capability on Leaf and Spine nodes.

Configuration

This section details the systematic procedures required to establish the VxLAN OAM feature, starting with the Leaf-Spine (CLOS) network topology and providing step-by-step Command Line Interface (CLI) instructions to configure VTEPs, Spines, and access switches for OAM operation, BGP routing, and CFM maintenance points.

Topology

The VxLAN OAM feature is demonstrated using a standard 2-Tier Leaf-Spine (CLOS) topology. This topology ensures redundant paths and is fundamental to the EVPN architecture used for VxLAN control.

The topology shown in Figure 1-10 consists of:

- 2 Spine nodes: Spine1 and Spine2
- 4 Leaf nodes are explicitly named as VTEPs (VTEP1, VTEP2, VTEP3, and VTEP4)
- 2 Access Switches: One connecting VTEP1/VTEP2, another connecting VTEP3/VTEP4

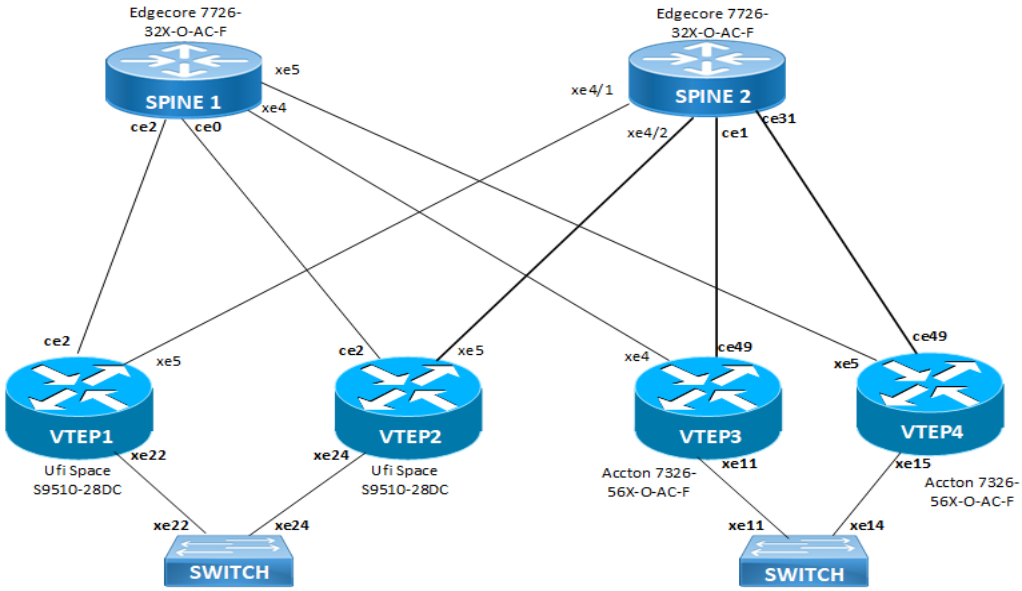

Figure 1-10: VxLAN OAM Topology

Procedure to configure VxLAN OAM

The configuration includes the steps listed in the below table.

Table 1-2: VxLAN OAM Configuration on Nodes

Nodes	VxLAN OAM Functionality	Relevant Configuration Steps
Leaf Nodes (VTEP1 - VTEP4)	Act as VTEPs (VxLAN Tunnel Endpoints) and are configured as MEPs (Maintenance End Points). Each VTEP uses a unique Loopback IP address (e.g., 1.1.1.1) as its VTEP-IP source/ destination. `nvo vxlan-oam`, explicit MEP creation (`ethernet cfm mep down`...), BGP configuration for VTEP-IP advertisement, EVPN instance mapping.	Configure VTEP1, VTEP2, VTEP3, and VTEP4 Sample running configuration on VTEP 1 Sample running configuration on VTEP 2 Sample running configuration on VTEP 3 Sample running configuration on VTEP 4 Validation VTEP1, VTEP2, VTEP3, VTEP4
Spine Nodes (Spine 1, Spine 2)	Act as the transit layer in the IP underlay and are configured as MIPs (Maintenance Intermediate Points). They also serve as BGP Route Reflectors for the EVPN overlay `nvo vxlan-oam spine`, explicit CFM domain configuration with static MIP creation (mip-creation static), BGP configuration for VTEP reachability and EVPN route reflection.	Configure the Spine 1 and Spine 2 Sample running configuration on Spine 1 Sample running configuration on Spine 2 Validation Spine 1, Spine 2
Access Switches	Provide the Layer 2 connectivity to the servers and connect to the VTEPs via trunk links (Port-Channels/LAG) that carry the traffic for the VLAN mapped to the VxLAN segment (VNI ID). Standard Layer 2 configuration with Rapid Spanning Tree Protocol (RSTP) bridge and VLAN trunking on the VTEP-facing links.	Configure Switch 1 and Switch 2 Sample running configuration on Switch 1

Configure VTEP1, VTEP2, VTEP3, and VTEP4

1. Login to global configuration mode. Enable VxLAN hardware filtering and multihoming on VTEP1 and VTEP2.

```
#configure terminal
(config)#hardware-profile filter vxlan enable
(config)#hardware-profile filter vxlan-mh enable
```

   on VTEP3 and VTEP4

```
#configure terminal
(config)#hardware-profile filter bfd-group enable
```

2. Configure host name and enable VxLAN. Network Virtualization Overlay (NVO) globally. Set the EVPN Ethernet Segment Identifier (ESI) hold-time and enable VxLAN multihoming for redundancy.

```
(config)#nvo vxlan enable
(config)#evpn esi hold-time 90
(config)#evpn vxlan multihoming enable
```

3. Configure an EVPN service (MAC VRF) named `vrf2`, map to a VLAN, RD (carries MAC/IP advertisements for VLAN2), and route target.

   on VTEP1

```
(config)#mac vrf vrf2
(config-vrf)#evpn-vlan-service vlan-based
(config-vrf)#rd 1.1.1.1:2
(config-vrf)#route-target both 100:2
```

on VTEP2

```
(config)#mac vrf vrf2
(config-vrf)#evpn-vlan-service vlan-based
(config-vrf)#rd 2.2.2.2:2
(config-vrf)#route-target both 100:2
```

on VTEP3

```
(config)#mac vrf vrf2
(config-vrf)#evpn-vlan-service vlan-based
(config-vrf)#rd 3.3.3.3:2
(config-vrf)#route-target both 100:2
```

on VTEP4

```
(config)#mac vrf vrf2
(config-vrf)#evpn-vlan-service vlan-based
(config-vrf)#rd 4.4.4.4:2
(config-vrf)#route-target both 100:2
```

4. Set the global VTEP IP address (loopback address) for VxLAN encapsulation and decapsulation. Configure NVO VxLAN ID 2, broadcast traffic replication at ingress and control plan as EVPN-BGP for MAC/IP learning. Associate the previously created EVPN VRF (`vrf2`) instance. This configuration maps the VLAN/EVPN instance (VLAN2) to a VxLAN segment (VNID 2).

   on VTEP1

   ```
   (config)#nvo vxlan vtep-ip-global 1.1.1.1
   (config)#nvo vxlan id 2 ingress-replication
   (config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf2
   ```

   on VTEP2

   ```
   (config)#nvo vxlan vtep-ip-global 2.2.2.2
   (config)#nvo vxlan id 2 ingress-replication
   (config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf2
   ```

   on VTEP3

   ```
   (config)#nvo vxlan vtep-ip-global 3.3.3.3
   (config)#nvo vxlan id 2 ingress-replication  bridge-vlan 2
   (config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf2
   ```

   on VTEP4

   ```
   (config)#nvo vxlan vtep-ip-global 4.4.4.4
   (config)#nvo vxlan id 2 ingress-replication  bridge-vlan 2
   (config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf2
   ```

5. Configure a multihomed EVPN access port interface with switchport mode, load interval, EVPN MAC holdtime, and multihomed shared system MAC address across multi-homed leafs (Ethernet Segment).

   on VTEP1 and VTEP2

   ```
   (config)#interface po12
   (config-if)#switchport
   ```

```
(config-if)#load-interval 30
(config-if)#evpn-mac-holdtime 180
(config-if)#evpn multi-homed system-mac 0000.0000.1212
```

on VTEP3

```
(config)#interface po34
(config-if)#switchport
(config-if)#bridge-group 1 spanning-tree disable
(config-if)#switchport trunk allowed vlan add 2
(config-if)#load-interval 30
(config-if)#evpn-mac-holdtime 180
(config-if)#access-if-vxlan
(config-if)#evpn multi-homed system-mac 0000.0000.3434
```

on VTEP4

6. Configure the access side subinterface for encapsulation, rewrite VLAN tag with POP, and EVPN mapping. This configuration connects VLAN 2 traffic from the access side (`po12.2`) to EVPN/VxLAN (VPN-ID 2) overlay.

   on VTEP1 and VTEP2

   ```
   (config)#interface po12.2 switchport
   (config-if)#encapsulation dot1q 2
   (config-if)#rewrite pop
   (config-if)#access-if-evpn
   (config-acc-if-evpn)#map vpn-id 2
   ```

7. Configure the underlay network connecting to both spines and access side with description and IP address on:

   VTEP1

   ```
   (config)#interface ce2
   (config-if)#description to_C1-SPINE1
   (config-if)#ip address 11.10.1.1/24
   (config-if)#commit
   (config-if)#exit
   (config)#interface lo
   (config)#ip address 1.1.1.1/32 secondary

   (config)#interface xe5
   (config-if)#description to_C1-SPINE2
   (config-if)#speed 10g
   (config-if)#load-interval 30
   (config-if)#ip address 12.10.1.1/24
   (config-if)#commit
   (config-if)#exit

   (config)interface xe22
   (config-if)#speed 10g
   (config-if)#channel-group 12 mode active
   (config-if)#commit
   (config-if)#exit
   ```

   VTEP2

   ```
   (config)#interface ce2
   ```

```
(config-if)#description to_C1-SPINE1
(config-if)#load-interval 30
(config-if)#ip address 21.10.1.1/24
(config-if)#commit
(config-if)#exit
(config)#interface lo
(config)#ip address 2.2.2.2/32 secondary

(config)#interface xe5
(config-if)#description to_C1-SPINE2
(config-if)#speed 10g
(config-if)#load-interval 30
(config-if)#ip address 22.10.1.1/24
(config-if)#commit
(config-if)#exit

(config)interface xe24
(config-if)#speed 10g
(config-if)#channel-group 12 mode active
(config-if)#commit
(config-if)#exit
```

VTEP3

```
(config)#interface ce49
(config-if)#description to_C1-SPINE2
(config-if)#load-interval 30
(config-if)#ip address 32.10.1.1/24
(config-if)#exit
(config)#interface lo
(config-if)#ip address 3.3.3.3/32 secondary
(config-if)#exit
(config)#interface xe4
(config-if)#description to_C1-SPINE1
(config-if)#load-interval 30
(config-if)#ip address 31.10.1.1/24
(config-if)#exit
(config)#interface xe11
(config-if)#channel-group 34 mode active
(config-if)#commit
(config-if)#exit
(config)#
```

VTEP4

```
(config)#interface ce49
(config-if)#description to_C1-SPINE2
(config-if)#load-interval 30
(config-if)#ip address 42.10.1.1/24
(config-if)#exit
(config)#interface lo
(config-if)#ip address 4.4.4.4/32 secondary
```

```
(config-if)#exit
(config)#interface xe5
(config-if)#description to_C1-SPINE1
(config-if)#load-interval 30
(config-if)#ip address 41.10.1.1/24
(config-if)#exit
(config)#interface xe15
(config-if)#channel-group 34 mode active
(config-if)#commit
(config-if)#exit
```

8. Configure BGP router mode and set router ID, peer groups, neighbors, and address families.

   on VTEP1

```
(config)#router bgp 100
(config-router)#bgp router-id 1.1.1.1

(config-router)#neighbor Overlay peer-group
(config-router)#neighbor Overlay remote-as 200
(config-router)#neighbor Overlay ebgp-multihop
(config-router)#neighbor Overlay update-source lo
(config-router)#neighbor Overlay advertisement-interval 0
(config-router)#neighbor Overlay fall-over bfd multihop
```

   Note: BGP EVPN overlay handles MAC/IP advertisement and tunnel mapping.

```
(config-router)#neighbor Underlay peer-group
(config-router)#neighbor Underlay remote-as 200
(config-router)#neighbor Underlay advertisement-interval 0
(config-router)#neighbor Underlay fall-over bfd
```

   Note: BGP underlay handles reachability between VTEPs

```
(config-router)#neighbor 11.10.1.2 peer-group Underlay
(config-router)#neighbor 11.11.11.11 peer-group Overlay
(config-router)#neighbor 11.11.11.11 description to_c1-spine1

(config-router)#neighbor 12.10.1.2 peer-group Underlay
(config-router)#neighbor 22.22.22.22 peer-group Overlay
(config-router)#neighbor 22.22.22.22 description to_c1-spine2
(config-router)#commit
(config-router)#exit

(config-router)#address-family ipv4 unicast
(config-router-af)#network 1.1.1.1/32
(config-router-af)#max-paths ebgp 2
(config-router-af)#neighbor Underlay activate
(config-router-af)#neighbor Underlay allowas-in 1
(config-router-af)#exit-address-family
```

   Note: This configuration establishes eBGP sessions 2 with the spines over physical links (`11.10.1.2`, `12.10.1.2`) and, advertises the VTEP IP (`1.1.1.1/32`) into the underlay. The `max-paths ebgp 2` allows ECMP load-balancing across both spines.

```
(config-router)# address-family l2vpn evpn
```

```
(config-router-af)#neighbor Overlay activate
(config-router-af)#neighbor Overlay allowas-in 1
(config-router-af)#exit-address-family
```

Note: This configuration establishes BGP EVPN sessions between the VTEPs and the spine route reflectors. It also enables the EVPN control plane for VxLAN MAC/IP route exchange for VxLAN overlays.

on VTEP2

```
(config)#router bgp 100
(config-router)#bgp router-id 2.2.2.2
(config-router)#neighbor Overlay peer-group
(config-router)#neighbor Overlay remote-as 200
(config-router)#neighbor Overlay ebgp-multihop
(config-router)#neighbor Overlay update-source lo
(config-router)#neighbor Overlay advertisement-interval 0
(config-router)#neighbor Overlay fall-over bfd multihop
(config-router)#neighbor Underlay peer-group
(config-router)#neighbor Underlay remote-as 200
(config-router)#neighbor Underlay advertisement-interval 0
(config-router)#neighbor Underlay fall-over bfd
(config-router)#neighbor 11.11.11.11 peer-group Overlay
(config-router)#neighbor 21.10.1.2 peer-group Underlay
(config-router)#neighbor 22.10.1.2 peer-group Underlay
(config-router)#neighbor 22.22.22.22 peer-group Overlay
(config-router)#commit
(config-router)#exit

(config-router)#address-family ipv4 unicast
(config-router-af)#network 2.2.2.2/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#neighbor Underlay activate
(config-router-af)#neighbor Underlay allowas-in 1
(config-router-af)#exit-address-family

(config-router)# address-family l2vpn evpn
(config-router-af)#neighbor Overlay activate
(config-router-af)#neighbor Overlay allowas-in 1
(config-router-af)#exit-address-family
```

on VTEP3

```
(config)#router bgp 100
(config-router)#bgp router-id 3.3.3.3
(config-router)#neighbor Overlay peer-group
(config-router)#neighbor Overlay remote-as 200
(config-router)#neighbor Overlay ebgp-multihop
(config-router)#neighbor Overlay update-source lo
(config-router)#neighbor Overlay advertisement-interval 0
(config-router)#neighbor Overlay fall-over bfd multihop
(config-router)#neighbor Underlay peer-group
(config-router)#neighbor Underlay remote-as 200
(config-router)#neighbor Underlay advertisement-interval 0
(config-router)#neighbor Underlay fall-over bfd
```

```
(config-router)#neighbor 11.11.11.11 peer-group Overlay
(config-router)#neighbor 22.22.22.22 peer-group Overlay
(config-router)#neighbor 31.10.1.2 peer-group Underlay
(config-router)#neighbor 32.10.1.2 peer-group Underlay
(config-router)#commit
(config-router)#exit

(config-router)#address-family ipv4 unicast
(config-router-af)#network 3.3.3.3/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#neighbor Underlay activate
(config-router-af)#neighbor Underlay allowas-in 1
(config-router-af)#exit-address-family

(config-router)# address-family l2vpn evpn
(config-router-af)#neighbor Overlay activate
(config-router-af)#neighbor Overlay allowas-in 1
(config-router-af)#exit-address-family
```

on VTEP4

```
(config)#router bgp 100
(config-router)#bgp router-id 4.4.4.4
(config-router)#neighbor Overlay peer-group
(config-router)#neighbor Overlay remote-as 200
(config-router)#neighbor Overlay ebgp-multihop
(config-router)#neighbor Overlay update-source lo
(config-router)#neighbor Overlay advertisement-interval 0
(config-router)#neighbor Overlay fall-over bfd multihop
(config-router)#neighbor Underlay peer-group
(config-router)#neighbor Underlay remote-as 200
(config-router)#neighbor Underlay advertisement-interval 0
(config-router)#neighbor Underlay fall-over bfd
(config-router)#neighbor 11.11.11.11 peer-group Overlay
(config-router)#neighbor 22.22.22.22 peer-group Overlay
(config-router)#neighbor 41.10.1.2 peer-group Underlay
(config-router)#neighbor 42.10.1.2 peer-group Underlay
(config-router)#commit
(config-router)#exit

(config-router)#address-family ipv4 unicast
(config-router-af)#network 4.4.4.4/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#neighbor Underlay activate
(config-router-af)#neighbor Underlay allowas-in 1
(config-router-af)#exit-address-family

(config-router)# address-family l2vpn evpn
(config-router-af)#neighbor Overlay activate
(config-router-af)#neighbor Overlay allowas-in 1
(config-router-af)#exit-address-family
```

The following VxLAN OAM configuration enables the VxLAN OAM monitoring for the VxLAN segment with VNID 2, continuously checks the health of VxLAN tunnels to three remote VTEP 2, VTEP 3 and VTEP 4 (2.2.2.2, 3.3.3.3, and 4.4.4.4) and verifies that all expected remote endpoints are reachable.

9. Configure VxLAN OAM and Ethernet CFM to support monitoring and diagnostics for VxLAN tunnels.

```
(config)#nvo vxlan-oam
```

10. Set up a CFM Maintenance Domain (MD) named `mdom2` with `level 7` and default MIP (Maintenance Intermediate Point) creation.

```
(config)#ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation default
```

where

- `domain-type character-string` specifies the domain name format as a string.
- `domain-name mdom2` specifies the maintenance domain named as `mdom2`.
- `level 7` sets the CFM level 7, typically used for customer-level monitoring.
- `mip-creation default` enables default creation of MIPs for intermediate devices.

11. Define a MA service named `ma2` within the `mdom2` domain, associated with EVPN ID 2.

```
(config-ether-cfm)#service ma-type string ma-name ma2
(config-ether-cfm-ma)#evpn 2
```

Note: The following options:

- `ma-type string` specifies the MA name format as a string.
- `ma-name ma2` names the MA ma2.
- `evpn 2` links the service to EVPN with VPN-ID 2, aligning with the VxLAN ID 2 from the previous configuration. Refer to step, Configure the access side subinterface for encapsulation, rewrite VLAN tag with POP, and EVPN mapping. This configuration connects VLAN 2 traffic from the access side (po12.2) to EVPN/VxLAN (VPN-ID 2) overlay.)

Note: This configuration monitors VxLAN tunnels that belong to VNID 2.

12. Configure downward-facing MEPs with

MPIDs 13, 12 and 14 for the remote VTEPs 3.3.3.3, 2.2.2.2 and 4.4.4.4 for enabling continuity check (CC) multicast on VTEP1

```
(config-ether-cfm-ma)#ethernet cfm mep down mpid 13 active true remote-vtep 3.3.3.3
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode

(config-ether-cfm-ma)#ethernet cfm mep down mpid 12 active true remote-vtep 2.2.2.2
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode

(config-ether-cfm-ma)#ethernet cfm mep down mpid 14 active true remote-vtep 4.4.4.4
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
```

Note: The following options:

- `mep down mpid 13` creates a downward-facing MEP with Maintenance Point ID 13.

- `active true` activates the MEP for monitoring.
- `remote-vtep 3.3.3.3` associates the MEP with the remote VTEP IP 3.3.3.3.
- `cc multicast state` enables continuity check messages in multicast mode to monitor connectivity.
- `exit-ether-ma-mep-mode` exits the MEP configuration sub-mode.

Note: This configuration creates a down MEP with ID 13/12/14 that monitors the VxLAN tunnel to remote VTEP 3.3.3.3. / 2.2.2.2. / 4.4.4.4 respectively.

MPIDs 23, 21 and 24 for the remote VTEPs 3.3.3.3, 1.1.1.1 and 4.4.4.4 for enabling continuity check (CC) multicast on VTEP2

```
(config-ether-cfm-ma)#ethernet cfm mep down mpid 23 active true remote-vtep
3.3.3.3
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
(config-ether-cfm-ma)#ethernet cfm mep down mpid 21 active true remote-vtep
1.1.1.1
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode

(config-ether-cfm-ma)#ethernet cfm mep down mpid 24 active true remote-vtep
4.4.4.4
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
```

Note:

MPIDs 32, 31 and 34 for the remote VTEPs 2.2.2.2, 1.1.1.1 and 4.4.4.4 for enabling continuity check (CC) multicast on VTEP3

```
(config-ether-cfm-ma)#ethernet cfm mep down mpid 32 active true remote-vtep
2.2.2.2
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
(config-ether-cfm-ma)#ethernet cfm mep down mpid 31 active true remote-vtep
1.1.1.1
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode

(config-ether-cfm-ma)#ethernet cfm mep down mpid 34 active true remote-vtep
4.4.4.4
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
```

MPIDs 42, 41 and 43 for the remote VTEPs 2.2.2.2, 1.1.1.1 and 3.3.3.3 for enabling continuity check (CC) multicast on VTEP4

```
(config-ether-cfm-ma)#ethernet cfm mep down mpid 42 active true remote-vtep
2.2.2.2
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
(config-ether-cfm-ma)#ethernet cfm mep down mpid 41 active true remote-vtep
1.1.1.1
(config-ether-cfm-ma-mep)#cc multicast state enable
(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
```

```
      (config-ether-cfm-ma)#ethernet cfm mep down mpid 43 active true remote-vtep
3.3.3.3
      (config-ether-cfm-ma-mep)#cc multicast state enable
      (config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
```

13. Enable crosschecking for remote MEPs. Cross-check verifies that these remote MEPs exist and are reachable, ensuring the full end-to-end path is operational. Defines expected remote MEPs with IDs

    31, 21, and 41 on VTEP1

    ```
    (config-ether-cfm-ma)#mep crosscheck mpid 31
    (config-ether-cfm-ma)#mep crosscheck mpid 21
    (config-ether-cfm-ma)#mep crosscheck mpid 41
    (config-ether-cfm-ma)#cc interval 1s
    (config-ether-cfm-ma)#exit-ether-ma-mode
    ```

    32, 12, and 42 on VTEP2

    ```
    (config-ether-cfm-ma)#mep crosscheck mpid 32
    (config-ether-cfm-ma)#mep crosscheck mpid 12
    (config-ether-cfm-ma)#mep crosscheck mpid 42
    (config-ether-cfm-ma)#cc interval 1s
    (config-ether-cfm-ma)#exit-ether-ma-mode
    ```

    23, 13, and 43 on VTEP3

    ```
    (config-ether-cfm-ma)#mep crosscheck mpid 23
    (config-ether-cfm-ma)#mep crosscheck mpid 13
    (config-ether-cfm-ma)#mep crosscheck mpid 43
    (config-ether-cfm-ma)#cc interval 1s
    (config-ether-cfm-ma)#exit-ether-ma-mode
    ```

    24, 14, and 34 on VTEP4

    ```
    (config-ether-cfm-ma)#mep crosscheck mpid 23
    (config-ether-cfm-ma)#mep crosscheck mpid 13
    (config-ether-cfm-ma)#mep crosscheck mpid 43
    (config-ether-cfm-ma)#cc interval 1s
    (config-ether-cfm-ma)#exit-ether-ma-mode
    ```

Sample running configuration on VTEP 1

```
C1-Leaf1#show running-config  (SP Node)
!
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
hardware-profile statistics cfm-ccm enable
!
hostname C1-Leaf1
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf vrf2
 evpn-vlan-service vlan-based
 rd 1.1.1.1:2
 route-target both 100:2
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan id 2 ingress-replication
```

```
 vxlan host-reachability-protocol evpn-bgp vrf2
!
interface po12
 switchport
 load-interval 30
 evpn-mac-holdtime 180
 evpn multi-homed system-mac 0000.0000.1212
!
interface po12.2 switchport
 encapsulation dot1q 2
 rewrite pop
 access-if-evpn
  map vpn-id 2
!
interface ce2
 description to_C1-SPINE1
 ip address 11.10.1.1/24
!
interface lo
 ip address 1.1.1.1/32 secondary
!
interface xe5
 description to_C1-SPINE2
 speed 10g
 load-interval 30
 ip address 12.10.1.1/24
!
interface xe22
 speed 10g
 channel-group 12 mode active
!
 exit
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor Overlay peer-group
 neighbor Overlay remote-as 200
 neighbor Overlay ebgp-multihop
 neighbor Overlay update-source lo
 neighbor Overlay advertisement-interval 0
 neighbor Overlay fall-over bfd multihop
 neighbor Underlay peer-group
 neighbor Underlay remote-as 200
 neighbor Underlay advertisement-interval 0
 neighbor Underlay fall-over bfd
 neighbor 11.10.1.2 peer-group Underlay
 neighbor 11.11.11.11 peer-group Overlay
 neighbor 11.11.11.11 description to_c1-spine1
 neighbor 12.10.1.2 peer-group Underlay
 neighbor 22.22.22.22 peer-group Overlay
 neighbor 22.22.22.22 description to_c1-spine2
 !
 address-family ipv4 unicast
 network 1.1.1.1/32
 max-paths ebgp 2
 neighbor Underlay activate
 neighbor Underlay allowas-in 1
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor Overlay activate
 neighbor Overlay allowas-in 1
 exit-address-family
 !
 exit
!
end

##############################

nvo vxlan-oam
!
ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation default
```

```
 service ma-type string ma-name ma2
   evpn 2
   ethernet cfm mep down mpid 13 active true remote-vtep 3.3.3.3
     cc multicast state enable
     exit-ether-ma-mep-mode
   ethernet cfm mep down mpid 12 active true remote-vtep 2.2.2.2
     cc multicast state enable
     exit-ether-ma-mep-mode
   ethernet cfm mep down mpid 14 active true remote-vtep 4.4.4.4
     cc multicast state enable
     exit-ether-ma-mep-mode
   mep crosscheck mpid 31
   mep crosscheck mpid 21
   mep crosscheck mpid 41
   cc interval 1s
   exit-ether-ma-mode
```

Sample running configuration on VTEP 2

```
C1-Leaf2#show running-config (SP Node)
!
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
hardware-profile statistics cfm-ccm enable
!
hostname C1-Leaf2
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf vrf2
 evpn-vlan-service vlan-based
 rd 2.2.2.2:2
 route-target both 100:2
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 2 ingress-replication
 vxlan host-reachability-protocol evpn-bgp vrf2
!
interface po12
 switchport
 load-interval 30
 evpn-mac-holdtime 180
 evpn multi-homed system-mac 0000.0000.1212
!
interface po12.2 switchport
 encapsulation dot1q 2
 rewrite pop
 access-if-evpn
  map vpn-id 2
!
interface ce2
 description to_C1-SPINE1
 load-interval 30
 ip address 21.10.1.1/24
!
interface lo
 ip address 2.2.2.2/32 secondary
!
interface xe5
 description to_C1-SPINE2
 speed 10g
 load-interval 30
 ip address 22.10.1.1/24
!
interface xe24
 speed 10g
 channel-group 12 mode active
```

```
!
 exit
!
router bgp 100
 bgp router-id 2.2.2.2
 neighbor Overlay peer-group
 neighbor Overlay remote-as 200
 neighbor Overlay ebgp-multihop
 neighbor Overlay update-source lo
 neighbor Overlay advertisement-interval 0
 neighbor Overlay fall-over bfd multihop
 neighbor Underlay peer-group
 neighbor Underlay remote-as 200
 neighbor Underlay advertisement-interval 0
 neighbor Underlay fall-over bfd
 neighbor 11.11.11.11 peer-group Overlay
 neighbor 21.10.1.2 peer-group Underlay
 neighbor 22.10.1.2 peer-group Underlay
 neighbor 22.22.22.22 peer-group Overlay
 !
 address-family ipv4 unicast
 network 2.2.2.2/32
 max-paths ebgp 4
 neighbor Underlay activate
 neighbor Underlay allowas-in 1
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor Overlay activate
 neighbor Overlay allowas-in 1
 exit-address-family
 !
 exit
!
end

!
C1-Leaf2#

##############################

nvo vxlan-oam
!
ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation default
 service ma-type string ma-name ma2
   evpn 2
   ethernet cfm mep down mpid 23 active true remote-vtep 3.3.3.3
     cc multicast state enable
     exit-ether-ma-mep-mode
   ethernet cfm mep down mpid 21 active true remote-vtep 1.1.1.1
     cc multicast state enable
     exit-ether-ma-mep-mode
   ethernet cfm mep down mpid 24 active true remote-vtep 4.4.4.4
     cc multicast state enable
     exit-ether-ma-mep-mode
   mep crosscheck mpid 32
   mep crosscheck mpid 12
   mep crosscheck mpid 42
   cc interval 1s
   exit-ether-ma-mode
```

Sample running configuration on VTEP 3

```
C1-Leaf3#show running-config (DC Node)
!
hardware-profile filter bfd-group enable
!
hostname C1-Leaf3
!
```

```
port-group 3 speed 10g
!
bridge 1 protocol rstp vlan-bridge
!
vlan database
 vlan 2 bridge 1 state enable
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf vrf2
 evpn-vlan-service vlan-based
 rd 3.3.3.3:2
 route-target both 100:2
!
nvo vxlan vtep-ip-global 3.3.3.3
!
nvo vxlan id 2 ingress-replication bridge-vlan 2
 vxlan host-reachability-protocol evpn-bgp vrf2
!
interface po34
 switchport
 bridge-group 1 spanning-tree disable
 switchport mode trunk
 switchport trunk allowed vlan add 2
 load-interval 30
 evpn-mac-holdtime 180
 access-if-vxlan
 evpn multi-homed system-mac 0000.0000.3434
!
interface ce49
 description to_C1-SPINE2
 load-interval 30
 ip address 32.10.1.1/24
!
interface lo
 ip address 3.3.3.3/32 secondary
!
interface xe4
 description to_C1-SPINE1
 load-interval 30
 ip address 31.10.1.1/24
!
interface xe11
 channel-group 34 mode active
!
 exit
!
router bgp 100
 bgp router-id 3.3.3.3
 neighbor Overlay peer-group
 neighbor Overlay remote-as 200
 neighbor Overlay ebgp-multihop
 neighbor Overlay update-source lo
 neighbor Overlay advertisement-interval 0
 neighbor Overlay fall-over bfd multihop
 neighbor Underlay peer-group
 neighbor Underlay remote-as 200
 neighbor Underlay advertisement-interval 0
 neighbor Underlay fall-over bfd
 neighbor 11.11.11.11 peer-group Overlay
 neighbor 22.22.22.22 peer-group Overlay
 neighbor 31.10.1.2 peer-group Underlay
 neighbor 32.10.1.2 peer-group Underlay
 !
 address-family ipv4 unicast
 network 3.3.3.3/32
 max-paths ebgp 4
 neighbor Underlay activate
 neighbor Underlay allowas-in 1
```

```
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor Overlay activate
 neighbor Overlay allowas-in 1
 exit-address-family
 !
 exit
!
end

!
C1-Leaf3#

##############################

nvo vxlan-oam
!
ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation default
 service ma-type string ma-name ma2
   evpn 2
   ethernet cfm mep down mpid 32 active true remote-vtep 2.2.2.2
     cc multicast state enable
     exit-ether-ma-mep-mode
   ethernet cfm mep down mpid 31 active true remote-vtep 1.1.1.1
     cc multicast state enable
     exit-ether-ma-mep-mode
   ethernet cfm mep down mpid 34 active true remote-vtep 4.4.4.4
     cc multicast state enable
     exit-ether-ma-mep-mode
   mep crosscheck mpid 23
   mep crosscheck mpid 13
   mep crosscheck mpid 43
   cc interval 1s
   exit-ether-ma-mode
```

Sample running configuration on VTEP 4

```
C1-Leaf4#show running-config (DC Node)
!
hardware-profile filter bfd-group enable
!
hostname C1-Leaf4
port-group 4 speed 10g
port-group 12 speed 10g
!
bridge 1 protocol rstp vlan-bridge
!
vlan database
 vlan-reservation 4037-4094
 vlan 2 bridge 1 state enable
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf vrf2
 evpn-vlan-service vlan-based
 rd 4.4.4.4:2
 route-target both 100:2
!
nvo vxlan vtep-ip-global 4.4.4.4
!
nvo vxlan id 2 ingress-replication bridge-vlan 2
 vxlan host-reachability-protocol evpn-bgp vrf2
!
interface po34
 switchport
 bridge-group 1 spanning-tree disable
 switchport mode trunk
 switchport trunk allowed vlan add 2
```

```
 load-interval 30
 evpn-mac-holdtime 180
 access-if-vxlan
 evpn multi-homed system-mac 0000.0000.3434
!
interface ce49
 description to_C1-SPINE2
 load-interval 30
 ip address 42.10.1.1/24
!
interface lo
 ip address 4.4.4.4/32 secondary
!
interface xe5
 description to_C1-SPINE1
 load-interval 30
 ip address 41.10.1.1/24
!
interface xe15
 channel-group 34 mode active
!
 exit
!
router bgp 100
 bgp router-id 4.4.4.4
 neighbor Overlay peer-group
 neighbor Overlay remote-as 200
 neighbor Overlay ebgp-multihop
 neighbor Overlay update-source lo
 neighbor Overlay advertisement-interval 0
 neighbor Overlay fall-over bfd multihop
 neighbor Underlay peer-group
 neighbor Underlay remote-as 200
 neighbor Underlay advertisement-interval 0
 neighbor Underlay fall-over bfd
 neighbor 11.11.11.11 peer-group Overlay
 neighbor 22.22.22.22 peer-group Overlay
 neighbor 41.10.1.2 peer-group Underlay
 neighbor 42.10.1.2 peer-group Underlay
 !
 address-family ipv4 unicast
 network 4.4.4.4/32
 max-paths ebgp 4
 neighbor Underlay activate
 neighbor Underlay allowas-in 1
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor Overlay activate
 neighbor Overlay allowas-in 1
 exit-address-family
 !
 exit
!
end

!
C1-Leaf4#

##############################

nvo vxlan-oam
!
ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation default
 service ma-type string ma-name ma2
   evpn 2
   ethernet cfm mep down mpid 42 active true remote-vtep 2.2.2.2
     cc multicast state enable
     exit-ether-ma-mep-mode
   ethernet cfm mep down mpid 41 active true remote-vtep 1.1.1.1
     cc multicast state enable
     exit-ether-ma-mep-mode
   ethernet cfm mep down mpid 43 active true remote-vtep 3.3.3.3
```

```
    cc multicast state enable
    exit-ether-ma-mep-mode
  mep crosscheck mpid 24
  mep crosscheck mpid 14
  mep crosscheck mpid 34
  cc interval 1s
  exit-ether-ma-mode
```

Configure the Spine 1 and Spine 2

The following configuration sets the underlay router and an EVPN Route Reflector (RR) in a VxLAN/EVPN network.

1. Login to Global Configuration Mode.

```
#configure terminal
(config)#hardware-profile statistics cfm-ccm enable
```

2. Configure the underlay network side

   interfaces `ce2,ce0, xe4` and `xe5` towards VTEP1, VTEP2, VTEP3, and VTEP4 with description and IP address on Spine 1.

```
(config)#hostname C1-Spine1

(config)#interface ce0
(config-if)#description to_C1-LEAF2
(config-if)#load-interval 30
(config-if)#ip address 21.10.1.2/24

(config)#interface ce2
(config-if)#description to_C1-LEAF1
(config-if)#load-interval 30
(config-if)#ip address 11.10.1.2/24

(config)#interface lo
(config-if)#ip address 11.11.11.11/32 secondary

(config)#interface xe4
(config-if)#description to_C1-LEAF3
(config-if)#load-interval 30
(config-if)#ip address 31.10.1.2/24

(config)#interface xe5
(config-if)#description to_C1-LEAF4
(config-if)#load-interval 30
(config-if)#ip address 41.10.1.2/24
```

   interfaces `ce1,ce31, xe4/1` and `xe4/2` towards VTEP1, VTEP2, VTEP3, and VTEP4 with description and IP address on Spine 2.

```
(config)#interface ce1
(config-if)#description to_C1-LEAF2
(config-if)#load-interval 30
(config-if)#ip address 32.10.1.2/24

(config)#interface ce31
(config-if)#description to_C1-LEAF4
```

```
(config-if)#load-interval 30
(config-if)#ip address 42.10.1.2/24

(config)#interface lo
(config-if)#ip address 22.22.22.22/32 secondary

(config)#interface xe4/1
(config-if)#description to_C1-LEAF1
(config-if)#load-interval 30
(config-if)#ip address 12.10.1.2/24

(config)#interface xe4/2
(config-if)#description to_C1-LEAF2
(config-if)#load-interval 30
(config-if)#ip address 22.10.1.2/24
```

3. Configure a BGP routing for both the underlay and overlay with router ID, peer groups, neighbors, and address families

   on Spine 1

```
(config)#router bgp 200
(config-router)#bgp router-id 11.11.11.11
(config-router)#no bgp inbound-route-filter
```

Note: These CLIs configure BGP control plane for both the underlay routing and EVPN overlay signaling.

```
(config-router)#neighbor Overlay peer-group
(config-router)#neighbor Overlay remote-as 100
(config-router)#neighbor Overlay ebgp-multihop
(config-router)#neighbor Overlay update-source lo
(config-router)#neighbor Overlay advertisement-interval 0
(config-router)#neighbor Overlay fall-over bfd multihop
```

Note: These CLIs configure an overlay BGP session for EVPN route exchange.

```
(config-router)#neighbor Underlay peer-group
(config-router)#neighbor Underlay remote-as 100
(config-router)#neighbor Underlay advertisement-interval 0
(config-router)#neighbor Underlay fall-over bfd
```

Note: These CLIs configure an eBGP session for reachability across the underlay network.

```
(config-router)#neighbor 1.1.1.1 peer-group Overlay
(config-router)#neighbor 2.2.2.2 peer-group Overlay
(config-router)#neighbor 3.3.3.3 peer-group Overlay
(config-router)#neighbor 4.4.4.4 peer-group Overlay
```

Note: These CLIs form an overlay BGP EVPN session from each VTEP to spine through loopback address.

```
(config-router)#neighbor 11.10.1.1 peer-group Underlay
(config-router)#neighbor 21.10.1.1 peer-group Underlay
(config-router)#neighbor 31.10.1.1 peer-group Underlay
(config-router)#neighbor 41.10.1.1 peer-group Underlay
```

Note: These CLIs form an eBGP session to each VTEP via the underlay physical interfaces.

```
(config-router)#address-family ipv4 unicast
```

```
(config-router-af)#network 11.11.11.11/32
(config-router-af)#max-paths ebgp 4
(config-router-af)#neighbor Underlay activate
(config-router-af)#exit-address-family
```

Note: These CLIs configuration establish redundant ECMP paths, ensuring that every VTEP can reach all other VTEPs.

```
(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor Overlay activate
(config-router-af)#exit-address-family
```

Note: Activates the L2VPN EVPN address family to advertise and reflect EVPN routes between VTEPs, enabling VxLAN tunnel establishment and MAC/IP reachability.

on Spine 2

```
(config)#router bgp 200
(config-router)#bgp router-id 22.22.22.22
(config-router)#no bgp inbound-route-filter
(config-router)#neighbor Overlay peer-group
(config-router)#neighbor Overlay remote-as 100
(config-router)#neighbor Overlay ebgp-multihop
(config-router)#neighbor Overlay update-source lo
(config-router)#neighbor Overlay advertisement-interval 0
(config-router)#neighbor Overlay fall-over bfd multihop
(config-router)#neighbor Underlay peer-group
(config-router)#neighbor Underlay remote-as 100
(config-router)#neighbor Underlay advertisement-interval 0
(config-router)#neighbor Underlay fall-over bfd
(config-router)#neighbor 1.1.1.1 peer-group Overlay
(config-router)#neighbor 2.2.2.2 peer-group Overlay
(config-router)#neighbor 3.3.3.3 peer-group Overlay
(config-router)#neighbor 4.4.4.4 peer-group Overlay
(config-router)#neighbor 12.10.1.1 peer-group Underlay
(config-router)#neighbor 22.10.1.1 peer-group Underlay
(config-router)#neighbor 32.10.1.1 peer-group Underlay
(config-router)#neighbor 42.10.1.1 peer-group Underlay

(config-router-af)address-family ipv4 unicast
(config-router-af)network 22.22.22.22/32
(config-router-af)max-paths ebgp 4
(config-router-af)neighbor Underlay activate
(config-router-af)exit-address-family

(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor Overlay activate
(config-router-af)#exit-address-family
```

The following VxLAN OAM configuration enables the spine to host MIPs that participate in OAM operations such as ping, path-trace, and continuity checks (CCMs).

4. Enable VxLAN OAM functionality in spine mode. Set up a CFM Maintenance Domain (MD) named mdom2 with level 7 and default MIP (Maintenance Intermediate Point) creation.

```
(config)#nvo vxlan-oam spine
```

```
(config)#ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation static
```

where

- `domain-type character-string` specifies the domain name format as a string.
- `domain-name mdom2` specifies the maintenance domain mdom2.
- `level 7` sets the CFM level 7, typically used for customer-level monitoring.
- `mip-creation static` enables static creation of MIPs for intermediate devices.

Note: These CLIs create MIPs within the `mdom2` domain (same as VTEP) that help to trace the VxLAN OAM path across the spine.

5. Define a MA service named ma2 within the mdom2 domain.

```
(config-ether-cfm)#service ma-type string ma-name ma2
(config-ether-cfm-ma)#exit-ether-ma-mode
```

where

- `ma-type string` specifies the MA name format as a string.
- `ma-name ma2` names the MA `ma2`

Note: These CLIs create a MA (`ma2`) within the domain `mdom2`. This ties the OAM configuration at the spine to the same MA (`ma2`) used by the VTEPs. When VTEP MEPs send OAM messages (e.g., ping or pathtrace), these messages traverse through the spine's MIP, which can respond or timestamp the message for troubleshooting. Thus, improving fault isolation and visibility in the VxLAN underlay or overlay path.

Sample running configuration on Spine 1

```
C1-Spine1#show running-config (SP Node)
!
hardware-profile statistics cfm-ccm enable
!
hostname C1-Spine1
!
interface ce0
 description to_C1-LEAF2
 load-interval 30
 ip address 21.10.1.2/24
!
interface ce2
 description to_C1-LEAF1
 load-interval 30
 ip address 11.10.1.2/24
!
interface lo
 ip address 11.11.11.11/32 secondary
!
interface xe4
 description to_C1-LEAF3
 load-interval 30
 ip address 31.10.1.2/24
!
interface xe5
 description to_C1-LEAF4
 load-interval 30
 ip address 41.10.1.2/24
!
 exit
!
router bgp 200
 bgp router-id 11.11.11.11
 no bgp inbound-route-filter
 neighbor Overlay peer-group
 neighbor Overlay remote-as 100
```

```
 neighbor Overlay ebgp-multihop
 neighbor Overlay update-source lo
 neighbor Overlay advertisement-interval 0
 neighbor Overlay fall-over bfd multihop
 neighbor Underlay peer-group
 neighbor Underlay remote-as 100
 neighbor Underlay advertisement-interval 0
 neighbor Underlay fall-over bfd
 neighbor 1.1.1.1 peer-group Overlay
 neighbor 2.2.2.2 peer-group Overlay
 neighbor 3.3.3.3 peer-group Overlay
 neighbor 4.4.4.4 peer-group Overlay
 neighbor 11.10.1.1 peer-group Underlay
 neighbor 21.10.1.1 peer-group Underlay
 neighbor 31.10.1.1 peer-group Underlay
 neighbor 41.10.1.1 peer-group Underlay
 !
 address-family ipv4 unicast
 network 11.11.11.11/32
 max-paths ebgp 4
 neighbor Underlay activate
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor Overlay activate
 exit-address-family
 !
 exit
!
!
end

!
C1-Spine1#

##############################

nvo vxlan-oam spine
!
ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation static
 service ma-type string ma-name ma2
   exit-ether-ma-mode
!
```

Sample running configuration on Spine 2

```
C1-Spine2#show running-config (DC Node)
!
hostname C1-Spine2
!
port ce4 breakout 4X10g
!
interface ce1
 description to_C1-LEAF2
 load-interval 30
 ip address 32.10.1.2/24
!
interface ce31
 description to_C1-LEAF4
 load-interval 30
 ip address 42.10.1.2/24
!
interface lo
 ip address 22.22.22.22/32 secondary
!
interface xe4/1
 description to_C1-LEAF1
 load-interval 30
 ip address 12.10.1.2/24
!
interface xe4/2
 description to_C1-LEAF2
 load-interval 30
```

```
 ip address 22.10.1.2/24
!
 exit
!
router bgp 200
 bgp router-id 22.22.22.22
 no bgp inbound-route-filter
 neighbor Overlay peer-group
 neighbor Overlay remote-as 100
 neighbor Overlay ebgp-multihop
 neighbor Overlay update-source lo
 neighbor Overlay advertisement-interval 0
 neighbor Overlay fall-over bfd multihop
 neighbor Underlay peer-group
 neighbor Underlay remote-as 100
 neighbor Underlay advertisement-interval 0
 neighbor Underlay fall-over bfd
 neighbor 1.1.1.1 peer-group Overlay
 neighbor 2.2.2.2 peer-group Overlay
 neighbor 3.3.3.3 peer-group Overlay
 neighbor 4.4.4.4 peer-group Overlay
 neighbor 12.10.1.1 peer-group Underlay
 neighbor 22.10.1.1 peer-group Underlay
 neighbor 32.10.1.1 peer-group Underlay
 neighbor 42.10.1.1 peer-group Underlay
 !
 address-family ipv4 unicast
 network 22.22.22.22/32
 max-paths ebgp 4
 neighbor Underlay activate
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor Overlay activate
 exit-address-family
 !
 exit
!
!
end

!
C1-Spine2#

##############################

nvo vxlan-oam spine
!
ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation static
 service ma-type string ma-name ma2
   exit-ether-ma-mode
!
```

Configure Switch 1 and Switch 2

The L2 access switch configuration sets up the VLAN trunking and uplinks toward VTEPs.

1. Login to Switch. Create a RSTP loop-free Layer-2 bridge supporting VLAN tagging.

   ```
   (config)#hostname C1-SW1
   (config)#bridge 1 protocol rstp vlan-bridge
   ```

2. Create a Port-Channel (LAG) redundancy interface between the switch and the VTEP.

   on VTEP 1 and VTEP2 from Switch1.

   ```
   (config)#interface po12
   (config-if)#switchport
   (config-if)#bridge-group 1
   (config-if)#switchport mode trunk
   ```

```
(config-if)#switchport trunk allowed vlan add 2
(config-if)#load-interval 30

(config)#interface xe22
(config-if)#channel-group 12 mode active

(config)#interface xe24
(config-if)#channel-group 12 mode active
```

on VTEP 3 and VTEP 4 from Switch 2.

```
(config)#interface po34
(config-if)#switchport
(config-if)#bridge-group 1
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add 2
(config-if)#load-interval 30

(config)#interface xe11
(config-if)#channel-group 34 mode active

(config)#interface xe14
(config-if)#channel-group 34 mode active
```

Sample running configuration on Switch 1

```
C1-SW1#show running-config
!

hostname C1-SW1
bridge 1 protocol rstp vlan-bridge

vlan database
 vlan 2 bridge 1 state enable
!
interface po12
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan add 2
 load-interval 30
!
interface xe22
 channel-group 12 mode active
!
interface xe24
 channel-group 12 mode active
!
!
end

!
C1-SW1#
```

Sample running configuration on Switch 2

```
C1-SW2#show running-config
!

hostname C1-SW2
```

```
bridge 1 protocol rstp vlan-bridge

vlan database
 vlan 2 bridge 1 state enable
!
interface po34
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan add 2
 load-interval 30
!
interface xe11
 channel-group 34 mode active
!
interface xe14
 channel-group 34 mode active
!
 exit
!
!
end
```

Validation

Execute following show commands to verify VxLAN OAM.

VTEP1

Verify the names and configuration details of the Maintenance Domain (MD) and Maintenance Association (MA) configured for VxLAN OAM.

```
C1-Leaf1#show ethernet cfm nvo vxlan-oam
------------------------------------------------------------------------------------------
---------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm  ma-evpn-id ma-name-type  mep-id mep-dir mep-stat vtep-ip
------------------------------------------------------------------------------------------
char-string  3      none      16777215   2-oct-int     8191   down    true     127.0.0.1
```

Display the traffic statistics for all CFM OAM messages, including CCM sent/received, Loopback, Link Trace, and Path Trace counts.

Note: CCM stats will not be counted when loss session is active for the MEP.

```
C1-Leaf1#show ethernet cfm statistics
Continuity Check Messages
   CCM Sent                 : 2421
   CCM Received             : 562
   CCM Drop                 : 0

Loop Back Messages
   LBM Sent                 : 0
   LBR Sent                 : 0
   LBM Drop                 : 0
   LBR Drop                 : 0
   LBR Received(Valid)      : 0
   LBR Received(Bad msdu)   : 0
   LBR Received(Out-of-Seq): 0
```

```
Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :0
   PTM Received (bad)        :0
   PTR Sent                  :0
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
   PTM Drop                  :0
```

Execute the following show commands to verify the connectivity status of MEP on the domain `mdom2` associated with MA name `ma2`.

Verify the Connectivity Status (e.g., Fully Connected) for all configured local MEPs within MD `mdom2`/`MA ma2`.

```
C1-Leaf1#show ethernet cfm ma status domain mdom2 mep all ma-name ma2

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
-------------------------------------------------------------
  13         Down MEP               -          Fully Connected
  12         Down MEP               -          Fully Connected
  14         Down MEP               -          Fully Connected
=============================================================
```

Check for active OAM defects (e.g., defRDICCM, defRemoteCCM) present on MEPs in the MD `mdom2`.

```
C1-Leaf1#show ethernet cfm errors domain mdom2

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
-----------------------------------------------------------------------
mdom2          ma2           7     NA    NA         13     .....
mdom2          ma2           7     NA    NA         12     .....
mdom2          ma2           7     NA    NA         14     .....


1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
```

Shows the total count of configured (MIPs, MEPs (Up/Down/Active), and active CCM sessions.

Note: A VxLAN OAM default MEP count is included in the MEPs count, in active state, but CCM feature is not supported by this default MEP.

```
C1-Leaf1#show ethernet cfm maintenance-points count
-----------------------------------------
Total No of MIPs                : 0
Total No of MEPs                : 4
```

```
Total No of UP MEPs              : 0
Total No of Down MEPs            : 4
Total No of Active MEPs          : 4
Total No of UP CCM sessions      : 3
Total No of Active Test sessions : 0
Total No of Active LM sessions   : 0
Total No of Active DM sessions   : 0
-----------------------------------------
```

Display the details of the local MEPs (ID, direction, CCM status/interval, Peer-Addr) on MD `mdom2`/MA `ma2`..

```
C1-Leaf1#show ethernet cfm maintenance-points local mep domain mdom2 ma-name ma2
-----------------------------------------------------
MPID Dir Lvl MA-vnid    CC-Stat CC-Intvl Def Peer-Addr
-----------------------------------------------------
13   Dn  7   2          Enable  1 sec    F   3.3.3.3
12   Dn  7   2          Enable  1 sec    F   2.2.2.2
14   Dn  7   2          Enable  1 sec    F   4.4.4.4
```

Display the state of the remote MEPs (ID, Rx CCM status, RDI, Peer-IP) detected or expected on MD `mdom2`/MA `ma2`.

```
C1-Leaf1#show ethernet cfm maintenance-points remote domain mdom2 ma-name ma2
 MA_NAME  RMEPID     LEVEL     Rx CCM    RDI    PEER-IP         TYPE
-------------------------------------------------------------------------------
 ma2       31         7         Yes       False  3.3.3.3         Configured
 ma2       21         7         Yes       False  2.2.2.2         Configured
 ma2       41         7         Yes       False  4.4.4.4         Configured
```

Display the detailed CCM status (State, Remote MEP IP and ID, VNID, last flap/recovery time) for a specific local MEP ID 12.

```
C1-Leaf1#show ethernet cfm ccm-status mep 12 domain mdom2 ma-name ma2
Remote MEP Identifier:     21  State: Up
Remote MEP: 2.2.2.2 VNID:     2
Last flapped: Wed Sep 10 10:35:11 2025
              (00:15:33 ago)
Last recovered: Wed Sep 10 10:41:45 2025
                (00:08:59 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:     1  flow id:  1
Remote defect indication: False
```

Display the detailed CCM status (State, Remote MEP IP and ID, VNID, last flap/recovery time) for a specific local MEP ID 13.

```
C1-Leaf1#show ethernet cfm ccm-status mep 13 domain mdom2 ma-name ma2
Remote MEP Identifier:     31  State: Up
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: Wed Sep 10 10:35:11 2025
              (00:15:33 ago)
Last recovered: Wed Sep 10 10:42:09 2025
                (00:08:35 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:     1  flow id:  1
Remote defect indication: False
```

Display the detailed CCM status (State, Remote MEP IP and ID, VNID, last flap/recovery time) for a specific local MEP ID 14.

```
C1-Leaf1#show ethernet cfm ccm-status mep 14 domain mdom2 ma-name ma2
Remote MEP Identifier:     41  State: Up
Remote MEP: 4.4.4.4 VNID:     2
Last flapped: Wed Sep 10 10:35:11 2025
              (00:15:35 ago)
Last recovered: Wed Sep 10 10:42:19 2025
                (00:08:27 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:     1  flow id:  1
Remote defect indication: False
```

VTEP2

Execute the following show commands to verify the name and other detailed information about MA and MD.

```
C1-Leaf2#show ethernet cfm nvo vxlan-oam

-------------------------------------------------------------------------------------------------------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm  ma-evpn-id ma-name-type  mep-id mep-dir mep-stat vtep-ip
-------------------------------------------------------------------------------------------------------
char-string  3      none      16777215   2-oct-int     8191   down    true     127.0.0.1
```

Execute the following show commands to verify the detail information about CFM statistics of CCM, LBM, LBR, LTM, LTR and PTM.

Note: CCM stats will not be counted when loss session is active for the MEP.

```
C1-Leaf2#show ethernet cfm statistics


NOTE: CCM stats will not be counted when loss session is active for the MEP

Continuity Check Messages
   CCM Sent                : 1716
   CCM Received            : 571
   CCM Drop                : 245

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 0
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0
```

```
Path Trace Messages
   PTM Sent                   :0
   PTM Received               :0
   PTM Received (bad)         :0
   PTR Sent                   :0
   PTR Received (Valid)       :0
   PTR received (Unexpected)  :0
   PTM Drop                   :0
!
```

Execute the following show commands to verify the connectivity status of MEP on the domain `mdom2` associated with MA name `ma2`.

```
C1-Leaf2#show ethernet cfm ma status domain mdom2 mep all ma-name ma2

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
------------------------------------------------------------
  23         Down MEP               -          Fully Connected
  21         Down MEP               -          Fully Connected
  24         Down MEP               -          Fully Connected
============================================================
```

Verify the defects present in MD named `mdom2`.

```
C1-Leaf2#show ethernet cfm errors domain mdom2

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
-----------------------------------------------------------------------
mdom2          ma2           7     NA    NA         23     .....
mdom2          ma2           7     NA    NA         21     .....
mdom2          ma2           7     NA    NA         24     .....


1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
```

Execute the following show commands to verify the information about the total CFM sessions count.

Note: A VxLAN OAM default MEP count is included in the MEPs count, in active state, but CCM feature is not supported by this default MEP.

```
C1-Leaf2#show ethernet cfm maintenance-points count
-------------------------------------------
Total No of MIPs                 : 0
Total No of MEPs                 : 4
Total No of UP MEPs              : 0
Total No of Down MEPs            : 4
Total No of Active MEPs          : 4
Total No of UP CCM sessions      : 3
Total No of Active Test sessions : 0
Total No of Active LM sessions   : 0
```

```
Total No of Active DM sessions   : 0
-----------------------------------------
```

Execute the following show commands to verify the information about the local MEP on the MD name `mdom2` associated with MA name `ma2`.

```
C1-Leaf2#show ethernet cfm maintenance-points local mep domain mdom2 ma-name ma2
-------------------------------------------------------
MPID Dir Lvl MA-vnid    CC-Stat CC-Intvl Def Peer-Addr
-------------------------------------------------------
23   Dn  7   2          Enable  1 sec    F   3.3.3.3
21   Dn  7   2          Enable  1 sec    F   1.1.1.1
24   Dn  7   2          Enable  1 sec    F   4.4.4.4
```

Execute the following show commands to verify the information about the remote MEPs on the MD named `mdom2` associated with MA name `ma2`.

```
C1-Leaf2#show ethernet cfm maintenance-points remote domain mdom2 ma-name ma2

 MA_NAME  RMEPID     LEVEL     Rx CCM    RDI    PEER-IP         TYPE
-------------------------------------------------------------------------------
 ma2       32         7         Yes       False  3.3.3.3         Configured
 ma2       12         7         Yes       False  1.1.1.1         Configured
 ma2       42         7         Yes       False  4.4.4.4         Configured
```

Execute the following show commands to verify the information about the CCM for the MEP ID 21 on the MD domain named `mdom2` associated with MA name `ma2`.

```
C1-Leaf2#show ethernet cfm ccm-status mep 21 domain mdom2 ma-name ma2
Remote MEP Identifier:     12  State: Up
Remote MEP: 1.1.1.1 VNID:     2
Last flapped: NA
Last recovered: NA
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    NA  flow id: NA
Remote defect indication: False
```

Execute the following show commands to verify the information about the CCM for the MEP ID 23 on the MD named `mdom2` associated with the MA name `ma2`.

```
C1-Leaf2#show ethernet cfm ccm-status mep 23 domain mdom2 ma-name ma2
Remote MEP Identifier:     32  State: Up
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: Wed Sep 10 10:43:01 2025
              (00:09:11 ago)
Last recovered: Wed Sep 10 10:43:21 2025
                (00:08:51 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:     1  flow id:  1
Remote defect indication: False
```

Execute the following show commands to verify the information about the CCM for the MEP ID 24 on MD named `mdom2` associated with the MA name `ma2`.

```
C1-Leaf2#show ethernet cfm ccm-status mep 24 domain mdom2 ma-name ma2
Remote MEP Identifier:     42  State: Up
Remote MEP: 4.4.4.4 VNID:     2
Last flapped: Wed Sep 10 10:43:01 2025
              (00:09:12 ago)
Last recovered: Wed Sep 10 10:43:31 2025
                (00:08:42 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:     1  flow id:  1
Remote defect indication: False
!
```

VTEP3

Execute the following show commands to verify the name and other detailed information about MA and MD.

```
############################## show outputs

C1-Leaf3#show ethernet cfm nvo vxlan-oam

------------------------------------------------------------------------------------
---------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm  ma-evpn-id ma-name-type  mep-id mep-dir mep-stat vtep-ip
------------------------------------------------------------------------------------
---------
char-string  3      none      16777215   2-oct-int     8191   down    true     127.0.0.1
```

Execute the following show commands to verify the detailed information about CFM statistics of CCM, LBM, LBR, LTM, LTR and PTM.

Note: CCM stats will not be counted when loss session is active for the MEP.

```
C1-Leaf3#show ethernet cfm statistics


Continuity Check Messages
   CCM Sent                : 1677
   CCM Received            : 1500
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 0
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0
```

```
Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :0
   PTM Received (bad)        :0
   PTR Sent                  :0
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
   PTM Drop                  :0
C1-Leaf3#

!
!
```

Execute the following show commands to verify the connectivity status of MEP on the domain `mdom2` associated with MA name `ma2`.

```
C1-Leaf3#show ethernet cfm ma status domain mdom2 mep all ma-name ma2

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
-------------------------------------------------------------
  32         Down MEP               -          Fully Connected
  31         Down MEP               -          Fully Connected
  34         Down MEP               -          Fully Connected
=============================================================
```

Verify the defects present in MD named `mdom2`.

```
C1-Leaf3#show ethernet cfm errors domain mdom2

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
-----------------------------------------------------------------------
mdom2          ma2           7     NA    NA         32     .....
mdom2          ma2           7     NA    NA         31     .....
mdom2          ma2           7     NA    NA         34     .....

1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
```

Execute the following show commands to verify the information about the total CFM sessions count.

Note: A VxLAN OAM default MEP count is included in the MEPs count, in active state, but CCM feature is not supported by this default MEP.

```
C1-Leaf3#show ethernet cfm maintenance-points count
-----------------------------------------
```

```
Total No of MIPs                : 0
Total No of MEPs                : 4
Total No of UP MEPs             : 0
Total No of Down MEPs           : 4
Total No of Active MEPs         : 4
Total No of UP CCM sessions     : 3
Total No of Active Test sessions : 0
Total No of Active LM sessions   : 0
Total No of Active DM sessions   : 0
-----------------------------------------
```

Execute the following show commands to verify the information about the local MEP on the MD name `mdom2` associated with MA name `ma2`.

```
C1-Leaf3#show ethernet cfm maintenance-points local mep domain mdom2 ma-name ma2
-----------------------------------------------------
MPID Dir Lvl MA-vnid    CC-Stat CC-Intvl Def Peer-Addr
-----------------------------------------------------
32   Dn  7   2          Enable  1 sec    F   2.2.2.2
31   Dn  7   2          Enable  1 sec    F   1.1.1.1
34   Dn  7   2          Enable  1 sec    F   4.4.4.4
```

Execute the following show commands to verify the information about the remote MEPs on the MD named `mdom2` associated with MA name `ma2`.

```
C1-Leaf3#show ethernet cfm maintenance-points remote domain mdom2 ma-name ma2
 MA_NAME  RMEPID     LEVEL     Rx CCM    RDI    PEER-IP         TYPE
-------------------------------------------------------------------------------
 ma2       23         7         Yes       False  2.2.2.2         Configured
 ma2       13         7         Yes       False  1.1.1.1         Configured
 ma2       43         7         Yes       False  4.4.4.4         Configured
```

Execute the following show commands to verify the information about the CCM for the MEP ID 31 on the MD named `mdom2` associated with MA name `ma2`.

```
C1-Leaf3#show ethernet cfm ccm-status mep 31 domain mdom2 ma-name ma2
Remote MEP Identifier:     13  State: Up
Remote MEP: 1.1.1.1 VNID:     2
Last flapped: NA
Last recovered: NA
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    NA  flow id: NA
Remote defect indication: False
```

Execute the following show commands to verify the information about the CCM for the MEP ID 32 on the MD named `mdom2` associated with the MA name `ma2`.

```
C1-Leaf3#show ethernet cfm ccm-status mep 32 domain mdom2 ma-name ma2
Remote MEP Identifier:     23  State: Up
Remote MEP: 2.2.2.2 VNID:     2
Last flapped: NA
Last recovered: NA
```

```
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    NA  flow id: NA
Remote defect indication: False
```

Execute the following show commands to verify the information about the CCM for the MEP ID 34 on MD named `mdom2` associated with the MA name `ma2`.

```
C1-Leaf3#show ethernet cfm ccm-status mep 34 domain mdom2 ma-name ma2
Remote MEP Identifier:     43  State: Up
Remote MEP: 4.4.4.4 VNID:     2
Last flapped: Wed Sep 10 10:43:10 2025
              (00:07:05 ago)
Last recovered: Wed Sep 10 10:43:16 2025
                (00:06:59 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:     1  flow id:  1
Remote defect indication: False
```

VTEP4

Execute the following show commands to verify the name and other detailed information about MA and MD.

```
C1-Leaf4#show ethernet cfm nvo vxlan-oam

------------------------------------------------------------------------------------
---------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm  ma-evpn-id ma-name-type  mep-id mep-dir mep-stat vtep-ip
------------------------------------------------------------------------------------
---------
char-string  3      none      16777215   2-oct-int     8191   down    true     127.0.0.1
```

Execute the following show commands to verify the detailed information about CFM statistics of CCM, LBM, LBR, LTM, LTR and PTM.

Note: CCM stats will not be counted when loss session is active for the MEP.

```
C1-Leaf4#show ethernet cfm statistics
Continuity Check Messages
   CCM Sent                : 1641
   CCM Received            : 0
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 0
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
```

```
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :0
   PTM Received (bad)        :0
   PTR Sent                  :0
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
   PTM Drop                  :0
```

Execute the following show commands to verify the information about the CCM for the MEP ID 41 on the MD domain named `mdom2` associated with MA name `ma2`.

```
C1-Leaf4#show ethernet cfm ccm-status mep 41 domain mdom2 ma-name ma2
Remote MEP Identifier:     14  State: Up
Remote MEP: 1.1.1.1 VNID:     2
Last flapped: NA
Last recovered: NA
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    NA  flow id: NA
Remote defect indication: False

Remote MEP Identifier:     24  State: Up
Remote MEP: 2.2.2.2 VNID:     2
Last flapped: NA
Last recovered: NA
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    NA  flow id: NA
Remote defect indication: False
```

Execute the following show commands to verify the information about the CCM for the MEP ID 43 on the MD named `mdom2` associated with the MA name `ma2`.

```
C1-Leaf4#show ethernet cfm ccm-status mep 43 domain mdom2 ma-name ma2
Remote MEP Identifier:     34  State: Up
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: NA
Last recovered: NA
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    NA  flow id: NA
Remote defect indication: False
```

Spine 1

Verify the names, level, and MIP permission of the Maintenance Domain (MD) and Maintenance Association (MA) configured for VxLAN OAM spine mode.

```
############################## show outputs
C1-Spine1#show ethernet cfm nvo vxlan-oam
```

```
! [execution timestamp : 2025 Sep 16 19:42:26]
-----------------------------------------------------------------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm ma-evpn-id ma-name-type mip-level
-----------------------------------------------------------------
char-string  3      static   16777215   2-oct-int    3
-----------------------------------------------------------------
md-name   : mdom2
ma-name   : ma2
name-type    level  mip-perm ma-evpn-id ma-name-type mip-level
-----------------------------------------------------------------
char-string  7      static   16777215   char-string  7
C1-Spine1#
```

Spine 2

Verify the names, level, and MIP permission of the Maintenance Domain (MD) and Maintenance Association (MA) configured for VxLAN OAM spine mode.

```
C1-Spine2#show ethernet cfm nvo vxlan-oam
-----------------------------------------------------------------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm ma-evpn-id ma-name-type mip-level
-----------------------------------------------------------------
char-string  3      static   16777215   2-oct-int    3
-----------------------------------------------------------------
md-name   : mdom2
ma-name   : ma2
name-type    level  mip-perm ma-evpn-id ma-name-type mip-level
-----------------------------------------------------------------
char-string  7      static   16777215   char-string  7
C1-Spine2#
```

Implementation Examples

The example configuration shared below ensures that VTEPs (Virtual Tunnel Endpoints) can proactively detect connectivity failures across the fabric. It is a proactive monitoring framework for four VTEPs (Leaf1 through Leaf4) to ensure 100% reachability of the EVPN instance in an overlay network. The goal is to detect underlay path failures or VTEP service interruptions within one second, prior to host-level application timeouts.

CFM configurations on All VTEPs

Following are the sample CFM configuration running on all the VTEPs.

Leaf1

```
C1-Leaf1#show running-config cfm
! [execution timestamp : 2025 Dec 23 12:38:49]
!
nvo vxlan-oam
```

```
      !
      ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation
default
       service ma-type string ma-name ma2
         evpn 2
         ethernet cfm mep down mpid 12 active true remote-vtep 2.2.2.2
           cc multicast state enable
           exit-ether-ma-mep-mode
         ethernet cfm mep down mpid 13 active true remote-vtep 3.3.3.3
           cc multicast state enable
           exit-ether-ma-mep-mode
         ethernet cfm mep down mpid 14 active true remote-vtep 4.4.4.4
           cc multicast state enable
           exit-ether-ma-mep-mode
         mep crosscheck mpid 21
         mep crosscheck mpid 31
         mep crosscheck mpid 41
         cc interval 1s
         exit-ether-ma-mode
      !
      !
```

Leaf2

```
      C1-Leaf2#show running-config cfm
      ! [execution timestamp : 2025 Dec 23 12:38:54]
      !
      nvo vxlan-oam
      !
      ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation
default
       service ma-type string ma-name ma2
         evpn 2
         ethernet cfm mep down mpid 23 active true remote-vtep 3.3.3.3
           cc multicast state enable
           exit-ether-ma-mep-mode
         ethernet cfm mep down mpid 21 active true remote-vtep 1.1.1.1
           cc multicast state enable
           exit-ether-ma-mep-mode
         ethernet cfm mep down mpid 24 active true remote-vtep 4.4.4.4
           cc multicast state enable
           exit-ether-ma-mep-mode
         mep crosscheck mpid 32
         mep crosscheck mpid 12
         mep crosscheck mpid 42
         cc interval 1s
         exit-ether-ma-mode
      !
      !
```

Leaf3

```
      C1-Leaf3#show running-config cfm
      ! [execution timestamp : 2025 Dec 23 12:38:35]
      !
      nvo vxlan-oam
      !
      ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation
default
       service ma-type string ma-name ma2
         evpn 2
         ethernet cfm mep down mpid 32 active true remote-vtep 2.2.2.2
           cc multicast state enable
           exit-ether-ma-mep-mode
         ethernet cfm mep down mpid 31 active true remote-vtep 1.1.1.1
           cc multicast state enable
           exit-ether-ma-mep-mode
         ethernet cfm mep down mpid 34 active true remote-vtep 4.4.4.4
           cc multicast state enable
           exit-ether-ma-mep-mode
         mep crosscheck mpid 23
         mep crosscheck mpid 13
         mep crosscheck mpid 43
         cc interval 1s
         exit-ether-ma-mode
      !
      !
```

Leaf4

```
      C1-Leaf4#show running-config cfm
      ! [execution timestamp : 2025 Dec 23 12:37:25]
      !
      nvo vxlan-oam
      !
      ethernet cfm domain-type character-string domain-name mdom2 level 7 mip-creation
default
       service ma-type string ma-name ma2
         evpn 2
         ethernet cfm mep down mpid 42 active true remote-vtep 2.2.2.2
           cc multicast state enable
           exit-ether-ma-mep-mode
         ethernet cfm mep down mpid 41 active true remote-vtep 1.1.1.1
           cc multicast state enable
           exit-ether-ma-mep-mode
         ethernet cfm mep down mpid 43 active true remote-vtep 3.3.3.3
           cc multicast state enable
           exit-ether-ma-mep-mode
         mep crosscheck mpid 24
         mep crosscheck mpid 14
         mep crosscheck mpid 34
         cc interval 1s
         exit-ether-ma-mode
```

```
!
!
```

Show output from all VTEPs before simulating CCM session down

The following show output ensures that all MEPs connectivity status are fully connected.

Leaf1

```
C1-Leaf1#terminal length 0
C1-Leaf1#
C1-Leaf1#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:57]

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
------------------------------------------------------------
  13         Down MEP              -          Fully Connected
  12         Down MEP              -          Fully Connected
  14         Down MEP              -          Fully Connected
============================================================
C1-Leaf1#
```

Leaf2

```
C1-Leaf2#terminal length 0
C1-Leaf2#
C1-Leaf2#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:57]

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
------------------------------------------------------------
  23         Down MEP              -          Fully Connected
  21         Down MEP              -          Fully Connected
  24         Down MEP              -          Fully Connected
============================================================
C1-Leaf2#
```

Leaf3

```
C1-Leaf3#terminal length 0

C1-Leaf3#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:40]

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
```

```
------------------------------------------------------------
  32         Down MEP               -          Fully Connected
  31         Down MEP               -          Fully Connected
  34         Down MEP               -          Fully Connected
============================================================
```

Leaf4

```
C1-Leaf4#terminal length 0
C1-Leaf4#
C1-Leaf4#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 12:39:32]

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
------------------------------------------------------------
  42         Down MEP               -          Fully Connected
  41         Down MEP               -          Fully Connected
  43         Down MEP               -          Fully Connected
============================================================
```

The following show output ensures that no defects found in any MEPs.

Leaf1

```
C1-Leaf1#show ethernet cfm errors domain mdom2
! [execution timestamp : 2025 Dec 23 12:40:57]

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
------------------------------------------------------------------------
mdom2          ma2           7     NA    NA         13     .....
mdom2          ma2           7     NA    NA         12     .....
mdom2          ma2           7     NA    NA         14     .....


1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
C1-Leaf1#
```

Leaf2

```
C1-Leaf2#show ethernet cfm errors domain mdom2
! [execution timestamp : 2025 Dec 23 12:40:57]

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
------------------------------------------------------------------------
mdom2          ma2           7     NA    NA         23     .....
mdom2          ma2           7     NA    NA         21     .....
mdom2          ma2           7     NA    NA         24     .....
```

```
1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
```

Leaf3

```
C1-Leaf3#show ethernet cfm errors domain mdom2
! [execution timestamp : 2025 Dec 23 12:40:40]

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
---------------------------------------------------------------------
mdom2          ma2           7     4098  NA         32     .....
mdom2          ma2           7     4098  NA         31     .....
mdom2          ma2           7     4098  NA         34     .....


1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
```

Leaf4

```
C1-Leaf4#show ethernet cfm errors domain mdom2
! [execution timestamp : 2025 Dec 23 12:39:32]

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
---------------------------------------------------------------------
mdom2          ma2           7     4098  NA         42     .....
mdom2          ma2           7     4098  NA         41     .....
mdom2          ma2           7     4098  NA         43     .....

1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
```

The following output displays the MEPs count and the total CCM sessions count on each MEP.

Leaf1

```
C1-Leaf1#show ethernet cfm maintenance-points count
! [execution timestamp : 2025 Dec 23 12:40:57]
-----------------------------------------
Total No of MIPs                 : 0
Total No of MEPs                 : 4
Total No of UP MEPs              : 0
Total No of Down MEPs            : 4
Total No of Active MEPs          : 4
Total No of UP CCM sessions      : 3
Total No of Active Test sessions : 0
Total No of Active LM sessions   : 0
Total No of Active DM sessions   : 0
-----------------------------------------
Note: A VxLAN OAM default MEP count is included in the MEPs count, in active
state, but CCM feature is not supported by this Default MEP
C1-Leaf1#
```

Leaf2

```
C1-Leaf2#show ethernet cfm maintenance-points count
! [execution timestamp : 2025 Dec 23 12:40:57]
-----------------------------------------
Total No of MIPs                 : 0
Total No of MEPs                 : 4
Total No of UP MEPs              : 0
Total No of Down MEPs            : 4
Total No of Active MEPs          : 4
Total No of UP CCM sessions      : 3
Total No of Active Test sessions : 0
Total No of Active LM sessions   : 0
Total No of Active DM sessions   : 0
-----------------------------------------
Note: A VxLAN OAM default MEP count is included in the MEPs count, in active
state, however, the CCM feature is not supported by default MEP.
C1-Leaf2#
```

Leaf3

```
C1-Leaf3#show ethernet cfm maintenance-points count
! [execution timestamp : 2025 Dec 23 12:40:40]
-----------------------------------------
Total No of MIPs                 : 0
Total No of MEPs                 : 4
Total No of UP MEPs              : 0
Total No of Down MEPs            : 4
Total No of Active MEPs          : 4
Total No of UP CCM sessions      : 3
Total No of Active Test sessions : 0
Total No of Active LM sessions   : 0
Total No of Active DM sessions   : 0
-----------------------------------------
Note: A VxLAN OAM default MEP count is included in the MEPs count, in active
state, however, CCM feature is not supported by default MEP.
C1-Leaf3##
```

Leaf4

```
C1-Leaf4#show ethernet cfm maintenance-points count
! [execution timestamp : 2025 Dec 23 12:39:32]
-----------------------------------------
Total No of MIPs                 : 0
Total No of MEPs                 : 4
Total No of UP MEPs              : 0
Total No of Down MEPs            : 4
Total No of Active MEPs          : 4
Total No of UP CCM sessions      : 3
Total No of Active Test sessions : 0
Total No of Active LM sessions   : 0
Total No of Active DM sessions   : 0
-----------------------------------------
```

```
      Note: A VxLAN OAM default MEP count is included in the MEPs count, in active
state, however, CCM feature is not supported by default MEP.
      C1-Leaf4#
```

The following output displays the MEP information of local and remote interface on each VTEPs. Also the CFM objects created on each VTEPs.

Note that the CCM feature is not supported for default MEP.

Leaf1

```
C1-Leaf1#show ethernet cfm maintenance-points local mep domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:57]
------------------------------------------------------
MPID Dir Lvl MA-vnid    CC-Stat CC-Intvl Def Peer-Addr
------------------------------------------------------
13   Dn  7   2          Enable  1 sec    F   3.3.3.3
12   Dn  7   2          Enable  1 sec    F   2.2.2.2
14   Dn  7   2          Enable  1 sec    F   4.4.4.4
C1-Leaf1#

C1-Leaf1#show ethernet cfm maintenance-points remote domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:57]
 MA_NAME  RMEPID     LEVEL     Rx CCM    RDI    PEER-IP         TYPE
-------------------------------------------------------------------------------
 ma2       31         7         Yes       False  3.3.3.3         Configured
 ma2       21         7         Yes       False  2.2.2.2         Configured
 ma2       41         7         Yes       False  4.4.4.4         Configured
C1-Leaf1#

C1-Leaf1#show ethernet cfm nvo vxlan-oam
! [execution timestamp : 2025 Dec 23 12:40:58]
-----------------------------------------------------------------------------------------------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm  ma-evpn-id ma-name-type  mep-id mep-dir mep-stat vtep-ip
-----------------------------------------------------------------------------------------------
char-string  3      none      16777215   2-oct-int     8191   down    true     127.0.0.1
```

Leaf2

```
C1-Leaf2#
C1-Leaf2#show ethernet cfm maintenance-points local mep domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:57]
------------------------------------------------------
MPID Dir Lvl MA-vnid    CC-Stat CC-Intvl Def Peer-Addr
------------------------------------------------------
23   Dn  7   2          Enable  1 sec    F   3.3.3.3
21   Dn  7   2          Enable  1 sec    F   1.1.1.1
24   Dn  7   2          Enable  1 sec    F   4.4.4.4

C1-Leaf2#show ethernet cfm maintenance-points remote domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:57]
```

```
 MA_NAME  RMEPID     LEVEL     Rx CCM    RDI    PEER-IP         TYPE
-------------------------------------------------------------------------------
 ma2       32         7         Yes       False  3.3.3.3         Configured
 ma2       12         7         Yes       False  1.1.1.1         Configured
 ma2       42         7         Yes       False  4.4.4.4         Configured

C1-Leaf2#show ethernet cfm nvo vxlan-oam
! [execution timestamp : 2025 Dec 23 12:40:58]
-----------------------------------------------------------------------------------------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm  ma-evpn-id ma-name-type  mep-id mep-dir mep-stat vtep-ip
-----------------------------------------------------------------------------------------
char-string  3      none      16777215   2-oct-int     8191   down    true     127.0.0.1
```

Leaf3

```
C1-Leaf3#show ethernet cfm maintenance-points local mep domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:40]
-----------------------------------------------------
MPID Dir Lvl MA-vnid    CC-Stat CC-Intvl Def Peer-Addr
-----------------------------------------------------
32   Dn  7   2          Enable  1 sec    F   2.2.2.2
31   Dn  7   2          Enable  1 sec    F   1.1.1.1
34   Dn  7   2          Enable  1 sec    F   4.4.4.4

C1-Leaf3#show ethernet cfm maintenance-points remote domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 12:40:40]
 MA_NAME RMEPID     LEVEL     Rx CCM    RDI    PEER-IP         TYPE
-------------------------------------------------------------------------------
 ma2       23         7         Yes       False       2.2.2.2         Configured
 ma2       13         7         Yes       False       1.1.1.1         Configured
 ma2       43         7         Yes       False       4.4.4.4         Configured

C1-Leaf3#show ethernet cfm nvo vxlan-oam
! [execution timestamp : 2025 Dec 23 12:40:41]
-----------------------------------------------------------------------------------------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm  ma-evpn-id ma-name-type  mep-id mep-dir mep-stat vtep-ip
-----------------------------------------------------------------------------------------
char-string  3      none      16777215   2-oct-int     8191   down    true     127.0.0.1
```

Leaf4

```
C1-Leaf4#show ethernet cfm maintenance-points local mep domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 12:39:32]
-----------------------------------------------------
MPID Dir Lvl MA-vnid    CC-Stat CC-Intvl Def Peer-Addr
-----------------------------------------------------
42   Dn  7   2          Enable  1 sec    F   2.2.2.2
41   Dn  7   2          Enable  1 sec    F   1.1.1.1
43   Dn  7   2          Enable  1 sec    F   3.3.3.3

C1-Leaf4#show ethernet cfm maintenance-points remote domain mdom2 ma-name ma2
```

```
! [execution timestamp : 2025 Dec 23 12:39:32]
 MA_NAME RMEPID     LEVEL     Rx CCM    RDI    PEER-IP         TYPE
-------------------------------------------------------------------------------
 ma2       24         7         Yes       False       2.2.2.2         Configured
 ma2       14         7         Yes       False       1.1.1.1         Configured
 ma2       34         7         Yes       False       3.3.3.3         Configured

C1-Leaf4#show ethernet cfm nvo vxlan-oam
! [execution timestamp : 2025 Dec 23 12:39:32]
-------------------------------------------------------------------------------------------------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm  ma-evpn-id ma-name-type  mep-id mep-dir mep-stat vtep-ip
-------------------------------------------------------------------------------------------------
char-string  3      none      16777215   2-oct-int     8191   down    true     127.0.0.1
```

The following output shows the CFM statistics: CCM sent and received, LBM sent and LBR received, LTM sent and LTR received.

Note that the CCM send and received count is not considered for inactive MEP.

Leaf1

```
C1-Leaf1#show ethernet cfm statistics
! [execution timestamp : 2025 Dec 23 14:08:16]

Note: CCM stats will not be counted when loss session is active for the MEP
Continuity Check Messages
   CCM Sent                : 25311
   CCM Received            : 24853
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 0
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :0
   PTM Received (bad)        :0
   PTR Sent                  :0
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
```

```
   PTM Drop                  :0
```

Leaf2

```
C1-Leaf2#show ethernet cfm statistics
! [execution timestamp : 2025 Dec 23 14:08:20]
Note: CCM stats will not be counted when loss session is active for the MEP

Continuity Check Messages
   CCM Sent                : 25317
   CCM Received            : 24030
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 0
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :0
   PTM Received (bad)        :0
   PTR Sent                  :0
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
   PTM Drop                  :0
```

Leaf3

```
C1-Leaf3#show ethernet cfm statistics
! [execution timestamp : 2025 Dec 23 14:08:05]
Note: CCM stats will not be counted when loss session is active for the MEP

Continuity Check Messages
   CCM Sent                : 25395
   CCM Received            : 24904
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 0
```

```
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :0
   PTM Received (bad)        :0
   PTR Sent                  :0
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
   PTM Drop                  :0

C1-Leaf3#show ethernet cfm statistics
! [execution timestamp : 2025 Dec 23 14:09:26]
Note: CCM stats will not be counted when loss session is active for the MEP
Continuity Check Messages
   CCM Sent                : 25635
   CCM Received            : 25145
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 5
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :1
   PTM Received (bad)        :0
   PTR Sent                  :1
   PTR Received (Valid)      :0
```

```
   PTR received (Unexpected) :0
   PTM Drop                  :0
```

Leaf4

```
C1-Leaf4#show ethernet cfm statistics
! [execution timestamp : 2025 Dec 23 14:06:59]
Note: CCM stats will not be counted when loss session is active for the MEP
Continuity Check Messages
   CCM Sent                : 25419
   CCM Received            : 24919
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 0
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :0
   PTM Received (bad)        :0
   PTR Sent                  :0
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
   PTM Drop                  :0
```

The following show output display the CCM status between the VTEPs for the CC interval equal 1.

Leaf1

```
 C1-Leaf1#show ethernet cfm ccm-status mep 12 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:10:56]
Remote MEP Identifier:     21  State: Up
Remote MEP: 2.2.2.2 VNID:     2
Last flapped: Tue Dec 23 12:04:15 2025
              (02:06:41 ago)
Last recovered: Tue Dec 23 12:11:08 2025
                (01:59:48 ago)
Before flap: Good sequence id:   992  flow id:  1
After flap: Good sequence id:  1408  flow id:  1
```

```
Remote defect indication: False
!

C1-Leaf1#show ethernet cfm ccm-status mep 13 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:10:56]
Remote MEP Identifier:     31  State: Up
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: Tue Dec 23 11:47:40 2025
              (02:23:16 ago)
Last recovered: Tue Dec 23 11:47:47 2025
                (02:23:09 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    27  flow id:  1
Remote defect indication: False
!

C1-Leaf1#show ethernet cfm ccm-status mep 14 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:10:57]
Remote MEP Identifier:     41  State: Up
Remote MEP: 4.4.4.4 VNID:     2
Last flapped: Tue Dec 23 11:47:40 2025
              (02:23:17 ago)
Last recovered: Tue Dec 23 11:47:47 2025
                (02:23:10 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    30  flow id:  1
Remote defect indication: False
!
```

Leaf2

```
C1-Leaf2#show ethernet cfm ccm-status mep 21 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:11:00]
Remote MEP Identifier:     12  State: Up
Remote MEP: 1.1.1.1 VNID:     2
Last flapped: Tue Dec 23 12:04:11 2025
              (02:06:49 ago)
Last recovered: Tue Dec 23 12:11:04 2025
                (01:59:56 ago)
Before flap: Good sequence id:   997  flow id:  1
After flap: Good sequence id:  1413  flow id:  1
Remote defect indication: False
!

C1-Leaf2#show ethernet cfm ccm-status mep 23 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:11:00]
Remote MEP Identifier:     32  State: Up
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: Tue Dec 23 12:04:11 2025
              (02:06:49 ago)
```

```
Last recovered: Tue Dec 23 12:11:04 2025
                (01:59:56 ago)
Before flap: Good sequence id:  1011  flow id:  1
After flap: Good sequence id:  1427  flow id:  1
Remote defect indication: False
!

C1-Leaf2#show ethernet cfm ccm-status mep 24 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:11:01]
Remote MEP Identifier:     42  State: Up
Remote MEP: 4.4.4.4 VNID:     2
Last flapped: Tue Dec 23 12:04:11 2025
              (02:06:50 ago)
Last recovered: Tue Dec 23 12:11:04 2025
                (01:59:57 ago)
Before flap: Good sequence id:  1014  flow id:  1
After flap: Good sequence id:  1430  flow id:  1
Remote defect indication: False
!
C1-Leaf2#
```

Leaf3

```
C1-Leaf3#show ethernet cfm ccm-status mep 31 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:10:46]
Remote MEP Identifier:     13  State: Up
Remote MEP: 1.1.1.1 VNID:     2
Last flapped: Tue Dec 23 11:47:02 2025
              (02:23:44 ago)
Last recovered: Tue Dec 23 11:47:22 2025
                (02:23:24 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    13  flow id:  1
Remote defect indication: False
!

C1-Leaf3#show ethernet cfm ccm-status mep 32 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:10:46]
Remote MEP Identifier:     23  State: Up
Remote MEP: 2.2.2.2 VNID:     2
Last flapped: Tue Dec 23 12:03:50 2025
              (02:06:56 ago)
Last recovered: Tue Dec 23 12:10:43 2025
                (02:00:03 ago)
Before flap: Good sequence id:   992  flow id:  1
After flap: Good sequence id:  1408  flow id:  1
Remote defect indication: False
!
C1-Leaf3#show ethernet cfm ccm-status mep 34 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:10:47]
Remote MEP Identifier:     43  State: Up
```

```
Remote MEP: 4.4.4.4 VNID:     2
Last flapped: Tue Dec 23 11:47:02 2025
              (02:23:45 ago)
Last recovered: Tue Dec 23 11:47:13 2025
                (02:23:34 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    21  flow id:  1
Remote defect indication: False
!
C1-Leaf3#
```

Leaf4

```
C1-Leaf4#show ethernet cfm ccm-status mep 41 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:09:42]
Remote MEP Identifier:     14  State: Up
Remote MEP: 1.1.1.1 VNID:     2
Last flapped: Tue Dec 23 11:45:46 2025
              (02:23:56 ago)
Last recovered: Tue Dec 23 11:46:10 2025
                (02:23:32 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    13  flow id:  1
Remote defect indication: False
!

C1-Leaf4#show ethernet cfm ccm-status mep 42 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:09:42]
Remote MEP Identifier:     24  State: Up
Remote MEP: 2.2.2.2 VNID:     2
Last flapped: Tue Dec 23 12:02:38 2025
              (02:07:04 ago)
Last recovered: Tue Dec 23 12:09:31 2025
                (02:00:11 ago)
Before flap: Good sequence id:   992  flow id:  1
After flap: Good sequence id:  1408  flow id:  1
Remote defect indication: False
!

C1-Leaf4#show ethernet cfm ccm-status mep 43 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:09:42]
Remote MEP Identifier:     34  State: Up
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: Tue Dec 23 11:45:46 2025
              (02:23:56 ago)
Last recovered: Tue Dec 23 11:46:01 2025
                (02:23:41 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    18  flow id:  1
Remote defect indication: False
!
```

```
C1-Leaf4#
```

The below output verifies the connectivity from Leaf1 (1.1.1.1) to Leaf3 (3.3.3.3).

```
C1-Leaf1#ping ethernet evpn 13 domain mdom2 ma ma2 destination 3.3.3.3 2 flow-
inport xe22 flow-dst 001a.9002.000c flow-vlan 4 flow-src 001a.0902.000c flow-ethtype
0x800
! [execution timestamp : 2025 Dec 23 14:09:23]
% Warning: Applicability of given flow parameter(s) on the packet depends on the
load-balance setting in the system
Response from 3.3.3.3-2: transid:1 rtt:1.012 ms
Response from 3.3.3.3-2: transid:2 rtt:0.804 ms
Response from 3.3.3.3-2: transid:3 rtt:0.771 ms
Response from 3.3.3.3-2: transid:4 rtt:0.740 ms
Response from 3.3.3.3-2: transid:5 rtt:0.740 ms
 success rate is  100 (5/5)
```

Identifies the specific underlay path in the transit hop from Leaf1 (1.1.1.1) to Spine1 (11.10.1.2 - IP configured on the ingress interface `ce2`) before reaching the target VTEP (Leaf3).

```
C1-Leaf1#pathtrace ethernet evpn 13 domain mdom2 ma ma2 destination 3.3.3.3 2
flow-inport xe22 flow-dst 001a.9002.000c flow-vlan 4 flow-src 001a.0902.000c flow-
ethtype 0x800
2025 Dec 23 14:09:25.894 : C1-Leaf1 : ONMD : NOTIF : [CFM_PATHTRACE_STATUS_4]:
Pathtrace operation is complete, response received from the target vtep 3.3.3.3
! [execution timestamp : 2025 Dec 23 14:09:25]
% Warning: Applicability of given flow parameter(s) on the packet depends on the
load-balance setting in the system

Sent 2 PTM packets, attempted 2 hops
=====================================================================
TTL  Reply               Ingress    Egress     State      rtt (ms)
---------------------------------------------------------------------
1    11.10.1.2           ce2        xe4        UP / UP    1.744
2    3.3.3.3             xe4        null       UP         0.911
=====================================================================
C1-Leaf1#
```

The following output shows the statics for all the local meps (meps 12,13,14)

```
C1-Leaf1#show ethernet cfm statistics
! [execution timestamp : 2025 Dec 23 14:09:36]
Note: CCM stats will not be counted when loss session is active for the MEP
Continuity Check Messages
   CCM Sent                : 25551
   CCM Received            : 25093
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 5
   LBR Sent                : 0
```

```
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 5
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :2
   PTM Received              :0
   PTM Received (bad)        :0
   PTR Sent                  :0
   PTR Received (Valid)      :2
   PTR received (Unexpected) :0
   PTM Drop                  :0
```

Show output from all VTEPs after simulating CCM session down

After disabling `cc multicast` state on leaf3, the MEP state becomes down.

Leaf1

```
C1-Leaf1#2025 Dec 23 14:15:02.799 : C1-Leaf1 : ONMD : NOTIF :
[CFM_NVO3_OAM_CCM_STATE_CHANGE_4]: CCM Session DOWN: mep 13 ma ma2 md mdom2 last good
seq 8854 flowid 1 occurred 2025-12-23T14:15:02Z

C1-Leaf1#show ethernet cfm ccm-status mep 13 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:15:31]
Remote MEP Identifier:     31  State: Down
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: Tue Dec 23 14:15:02 2025
              (00:00:29 ago)
Last recovered: NA
Before flap: Good sequence id:  8854  flow id:  1
After flap: Good sequence id:    NA  flow id: NA
Remote defect indication: False
!

C1-Leaf1#show ethernet cfm statistics mep 13 domain modm2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:16:02]
%% MD not found in the list
C1-Leaf1#show ethernet cfm statistics mep 13 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:16:36]

Note: CCM stats will not counted when loss session is active for the MEP
```

```
 CFM Statistics for MEP 13 of MD mdom2
 =====================================
Continuity Check Messages
   CCM Sent                : 8937
   CCM Received            : 8828

Loop Back Messages
   LBM Sent                : 5
   LBR Sent                : 0
   LBR Received(Valid)     : 5
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :2
   PTM Received              :0
   PTR Sent                  :0
   PTR Received (Valid)      :2
   PTR received (Unexpected) :0
```

Note: There is no CCM packets received on Leaf1 from Leaf3, hence CCM received count is not incremented for MEP 13

```
C1-Leaf1#show ethernet cfm statistics mep 13 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:16:40]

 CFM Statistics for MEP 13 of MD mdom2
 =====================================
Continuity Check Messages
   CCM Sent                : 8941
   CCM Received            : 8828

Loop Back Messages
   LBM Sent                : 5
   LBR Sent                : 0
   LBR Received(Valid)     : 5
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0
```

```
Path Trace Messages
   PTM Sent                   :2
   PTM Received               :0
   PTR Sent                   :0
   PTR Received (Valid)       :2
   PTR received (Unexpected) :0

C1-Leaf1#show ethernet cfm statistics mep 13 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:16:42]
Note: CCM stats will not counted when loss session is active for the MEP

 CFM Statistics for MEP 13 of MD mdom2
 =====================================
Continuity Check Messages
   CCM Sent                 : 8943
   CCM Received             : 8828

Loop Back Messages
   LBM Sent                 : 5
   LBR Sent                 : 0
   LBR Received(Valid)      : 5
   LBR Received(Bad msdu)   : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                 : 0
   LTR Sent                 : 0
   LTR Received(Valid)      : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                   :2
   PTM Received               :0
   PTR Sent                   :0
   PTR Received (Valid)       :2
   PTR received (Unexpected) :0
```

Leaf1 detects the loss of connectivity when `cc multicast` is disabled on Leaf3.

```
C1-Leaf1#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 14:20:19]

MA/MEG Name: ma2
MA/MEG Status: Partially Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
-------------------------------------------------------------
  13         Down MEP              -      Partially Connected
  12         Down MEP              -      Partially Connected
  14         Down MEP              -      Partially Connected
```

```
      ============================================================

      C1-Leaf1#show ethernet cfm errors domain mdom2
      ! [execution timestamp : 2025 Dec 23 14:20:20]

      Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
      ------------------------------------------------------------------------
      mdom2          ma2           7     NA    NA         13     ..3..
      mdom2          ma2           7     NA    NA         12     .....
      mdom2          ma2           7     NA    NA         14     .....


      1. defRDICCM    2. defMACstatus  3. defRemoteCCM
      4. defErrorCCM  5. defXconCCM
      C1-Leaf1#
```

Leaf2

```
      C1-Leaf2#
      C1-Leaf2#2025 Dec 23 14:14:59.878 : C1-Leaf2 : ONMD : NOTIF : [CFM_ALARM_EVENT_4]:
defect none: mep 21 ma ma2 md mdom2

      C1-Leaf2#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
      ! [execution timestamp : 2025 Dec 23 14:20:19]

      MA/MEG Name: ma2
      MA/MEG Status: Active

      MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
      ------------------------------------------------------------
        23         Down MEP              -          Fully Connected
        21         Down MEP              -          Fully Connected
        24         Down MEP              -          Fully Connected
      ============================================================
      C1-Leaf2##
      C1-Leaf2#show ethernet cfm errors domain mdom2
      ! [execution timestamp : 2025 Dec 23 14:20:19]

      Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
      ------------------------------------------------------------------------
      mdom2          ma2           7     NA    NA         23     .....
      mdom2          ma2           7     NA    NA         21     1....
      mdom2          ma2           7     NA    NA         24     .....


      1. defRDICCM    2. defMACstatus  3. defRemoteCCM
      4. defErrorCCM  5. defXconCCM
      C1-Leaf2#

      C1-Leaf2#show ethernet cfm errors domain mdom2
      ! [execution timestamp : 2025 Dec 23 14:24:51]
```

```
Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
-------------------------------------------------------------------------
mdom2          ma2           7     NA    NA         23     .....
mdom2          ma2           7     NA    NA         21     .....
mdom2          ma2           7     NA    NA         24     .....


1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
C1-Leaf2#
```

Leaf3

```
C1-Leaf3#
Enter configuration commands, one per line.  End with CNTL/Z.
C1-Leaf3(config)#ethernet cfm domain-type character-string domain-name mdom2
level 7 mip-creation default
C1-Leaf3(config-ether-cfm)#service ma-type string ma-name ma2
C1-Leaf3(config-ether-cfm-ma)#ethernet cfm mep down mpid 31 active true remote-
vtep 1.1.1.1
C1-Leaf3(config-ether-cfm-ma-mep)#no cc
C1-Leaf3(config-ether-cfm-ma-mep)#cc multicast state disable
C1-Leaf3(config-ether-cfm-ma-mep)#commit
C1-Leaf3(config-ether-cfm-ma-mep)#end


C1-Leaf3#2025 Dec 23 14:14:39.039 : C1-Leaf3 : ONMD : NOTIF : [CFM_ALARM_EVENT_4]:
defect none: mep 31 ma ma2 md mdom2
C1-Leaf3#show ethernet cfm ccm-status mep 31 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:14:54]
Remote MEP Identifier:     13  State: Up
Remote MEP: 1.1.1.1 VNID:     2
Last flapped: Tue Dec 23 11:47:02 2025
              (02:27:52 ago)
Last recovered: Tue Dec 23 11:47:22 2025
                (02:27:32 ago)
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    13  flow id:  1
Remote defect indication: True
!

C1-Leaf3#show ethernet cfm statistics mep 31 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:16:44]

 CFM Statistics for MEP 31 of MD mdom2
 =====================================
Continuity Check Messages
   CCM Sent               : 0
   CCM Received           : 0

Loop Back Messages
```

```
   LBM Sent                : 0
   LBR Sent                : 5
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :1
   PTR Sent                  :1
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
C1-Leaf3#show ethernet cfm statistics mep 31 domain mdom2 ma-name ma2
! [execution timestamp : 2025 Dec 23 14:16:46]

 CFM Statistics for MEP 31 of MD mdom2
 =====================================
Continuity Check Messages
   CCM Sent                : 0
   CCM Received            : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 5
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
   LTR Received(Valid)     : 0
   LTR Received(unexpected): 0

Path Trace Messages
   PTM Sent                  :0
   PTM Received              :1
   PTR Sent                  :1
   PTR Received (Valid)      :0
   PTR received (Unexpected) :0
C1-Leaf3#
C1-Leaf3#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 14:20:00]
```

```
MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
------------------------------------------------------------
  32         Down MEP              -          Fully Connected
  31         Down MEP              -          Fully Connected
  34         Down MEP              -          Fully Connected
============================================================

C1-Leaf3#show ethernet cfm errors domain mdom2
! [execution timestamp : 2025 Dec 23 14:20:00]

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
------------------------------------------------------------------------
mdom2          ma2           7     4098  NA         32     .....
mdom2          ma2           7     4098  NA         31     1....
mdom2          ma2           7     4098  NA         34     .....


1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
C1-Leaf3#
```

Leaf4

```
C1-Leaf4#2025 Dec 23 14:13:26.963 : C1-Leaf4 : ONMD : NOTIF : [CFM_ALARM_EVENT_4]:
defect none: mep 41 ma ma2 md mdom2

C1-Leaf4#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 14:18:50]

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
------------------------------------------------------------
  42         Down MEP              -          Fully Connected
  41         Down MEP              -          Fully Connected
  43         Down MEP              -          Fully Connected
============================================================

C1-Leaf4#show ethernet cfm errors domain mdom2
! [execution timestamp : 2025 Dec 23 14:18:51]

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
------------------------------------------------------------------------
mdom2          ma2           7     4098  NA         42     .....
mdom2          ma2           7     4098  NA         41     1....
mdom2          ma2           7     4098  NA         43     .....
```

```
      1. defRDICCM    2. defMACstatus  3. defRemoteCCM
      4. defErrorCCM  5. defXconCCM
      C1-Leaf4#
```

Show output from all VTEPs after enabling CCM session up

After enabling the `cc multicast` state on Leaf 3 (mep 31), the connectivity is recovered.

```
      C1-Leaf3(config)#ethernet cfm domain-type character-string domain-name mdom2
level 7 mip-creation default
      C1-Leaf3(config-ether-cfm)#service ma-type string ma-name ma2
      C1-Leaf3(config-ether-cfm-ma)#ethernet cfm mep down mpid 31 active true remote-
vtep 1.1.1.1
      C1-Leaf3(config-ether-cfm-ma-mep)#cc multicast state enable
      C1-Leaf3(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
      C1-Leaf3(config-ether-cfm-ma)#commit
      C1-Leaf3(config-ether-cfm-ma)#end

      C1-Leaf3#start
      C1-Leaf3#show ethernet cfm ccm-status mep 31 domain mdom2 ma-name ma2
      ! [execution timestamp : 2025 Dec 23 14:23:59]
      Remote MEP Identifier:     13  State: Up
      Remote MEP: 1.1.1.1 VNID:     2
      Last flapped: Tue Dec 23 11:47:02 2025
                    (02:36:57 ago)
      Last recovered: Tue Dec 23 11:47:22 2025
                      (02:36:37 ago)
      Before flap: Good sequence id:    NA  flow id: NA
      After flap: Good sequence id:    13  flow id:  1
      Remote defect indication: False
      !
      C1-Leaf3#
```

Leaf1

```
      C1-Leaf1#clear ethernet cfm statistics
      C1-Leaf1#show ethernet cfm ccm-status mep 13 domain mdom2 ma-name ma2
      ! [execution timestamp : 2025 Dec 23 14:23:00]
      Remote MEP Identifier:     31  State: Down
      Remote MEP: 3.3.3.3 VNID:     2
      Last flapped: Tue Dec 23 14:15:02 2025
                    (00:07:58 ago)
      Last recovered: NA
      Before flap: Good sequence id:  8854  flow id:  1
      After flap: Good sequence id:    NA  flow id: NA
      Remote defect indication: False
      !

      C1-Leaf1#2025 Dec 23 14:23:56.596 : C1-Leaf1 : ONMD : NOTIF :
[CFM_NVO3_OAM_CCM_STATE_CHANGE_4]: CCM Session UP: mep 13 ma ma2 md mdom2 recovered good
seq 1 flowid 1 occurred 2025-12-23T14:23:56Z

      C1-Leaf1#show ethernet cfm ccm-status mep 13 domain mdom2 ma-name ma2
```

```
! [execution timestamp : 2025 Dec 23 14:24:28]
Remote MEP Identifier:     31  State: Up
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: Tue Dec 23 14:15:02 2025
              (00:09:26 ago)
Last recovered: Tue Dec 23 14:23:56 2025
                (00:00:32 ago)
Before flap: Good sequence id:  8854  flow id:  1
After flap: Good sequence id:     1  flow id:  1
Remote defect indication: False
!

C1-Leaf1#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 14:24:43]

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
-------------------------------------------------------------
  13         Down MEP              -          Fully Connected
  12         Down MEP              -          Fully Connected
  14         Down MEP              -          Fully Connected
=============================================================

C1-Leaf1#show ethernet cfm errors domain mdom2
! [execution timestamp : 2025 Dec 23 14:24:44]

Domain Name    MA Name       Level VLAN  InnerVLAN  MEPID  Defects
-----------------------------------------------------------------------
mdom2          ma2           7     NA    NA         13     .....
mdom2          ma2           7     NA    NA         12     .....
mdom2          ma2           7     NA    NA         14     .....

1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
C1-Leaf1#
```

Leaf3

```
C1-Leaf3#show ethernet cfm ma status domain mdom2 mep all ma-name ma2
! [execution timestamp : 2025 Dec 23 14:24:26]

MA/MEG Name: ma2
MA/MEG Status: Active

MEPID       DIRECTION       INTERFACE      CONNECTIVITY STATUS
-------------------------------------------------------------
  32         Down MEP              -          Fully Connected
  31         Down MEP              -          Fully Connected
```

```
  34          Down MEP                -          Fully Connected
============================================================
C1-Leaf3##
C1-Leaf3#show ethernet cfm errors domain mdom2
! [execution timestamp : 2025 Dec 23 14:24:26]

Domain Name     MA Name        Level VLAN  InnerVLAN  MEPID  Defects
-------------------------------------------------------------------------
mdom2           ma2            7     4098  NA         32     .....
mdom2           ma2            7     4098  NA         31     .....
mdom2           ma2            7     4098  NA         34     .....


1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
C1-Leaf3#
```

CLI Commands

Following are the new CLIs introduced:

- ping ethernet evpn
- pathtrace ethernet evpn
- show ethernet cfm ccm-status
- nvo vxlan-oam
- show ethernet cfm nvo vxlan-oam

ping ethernet evpn

Use this VxLAN OAM ping command to test the reachability of VTEPs involved in the VxLAN tunnel. This CLI provides flow parameters to simulate the host packet arriving from the access port.

Command Syntax

```
ping ethernet evpn ((MEPID domain DOMAIN_NAME ma MA_NAME)|) destination A.B.C.D
  EVPNID (interface-ip-address|)( flow-inport IFNAME flow-dst XXXX.XXXX.XXXX flow-
  vlan (VLANID (flow-innervlan VLANID|)|untagged)({flow-src XXXX.XXXX.XXXX | flow-
  ethtype ETHTYPE}|)| IFNAME|)
```

Parameters

`MEPID` Enter the MEP identifier <1-8191>. (Optional parameter in implicit mode.)

`domain DOMAIN_NAME`

Enter the name of the domain. Name length can be maximum of 43 characters. (Optional parameter in implicit mode.)

`ma MA_NAME` Enter maintenance association name. Name length can be maximum of 6 characters. (Optional parameter in implicit mode.)

`destination A.B.C.D`

Specify the destination VTEP IPv4 address.

`EVPNID` Enter the EVPN identifier in range <1-16777215>.

`interface-ip-address`

To fetch the ingress interface IP addresses of spine and leaf.

`flow-inport IFNAME`

Specify the access port name to simulate host packet arriving from the access port.

`flow-dst XXXX.XXXX.XXXX`

Specify the destination MAC address to simulate host packet arriving from the access port.

`flow-vlan VLANID`

Specify the outer VLAN ID mapped to the `flow-inport` interface to simulate host packet arriving from access port.

`flow-innervlan VLANID`

Specify the Inner VLAN ID to simulate host packet arriving from the access port.

`untagged` Specify untagged host packet arriving from access port.

`flow-src XXXX.XXXX.XXXX`

Specify the source MAC address to simulate host packet arriving from the access port.

`flow-ethtype ETHTYPE`

Specify the Ethertype value (0x600 - 0xFFFF) to simulate host packet arriving from the access port.

`IFNAME` Specify the PTM packet outgoing interface name.

Command Mode

Privilege execute mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

To verify ping to a VTEP passing through MA within MD:

```
#ping ethernet evpn 21 domain T_MD ma ma1 destination 1.1.1.1 10 flow-inport xe1 flow-dst 0010.7700.0001 flow-vlan 102 flow-ethtype 0x8100
% Warning: Applicability of given flow parameter(s) on the packet, depends on the load-balance rtag7 setting in the system
Response from 1.1.1.1-10: transid:1 rtt:3.231 ms
Response from 1.1.1.1-10: transid:2 rtt:3.332 ms
Response from 1.1.1.1-10: transid:3 rtt:3.285 ms
Response from 1.1.1.1-10: transid:4 rtt:3.162 ms
Response from 1.1.1.1-10: transid:5 rtt:3.320 ms
 success rate is  100 (5/5)
```

Note: To check the path of a particular flow `flow-inport`, `flow-dst` and `flow-vlan` are mandatory parameters irrespective of the load-balancing configuration.

Note: In Q2A devices, it is mandatory to use both `flow-vlan` and `flow-ethtype` parameters to get the exact VxLAN packet path. The following message is displayed to the user when not used together:

```
      #ping ethernet evpn destination 4.4.4.4 2 flow-inport xe11 flow-dst
0000.0241.0002 flow-vlan 2
      %% Option 'flow-vlan' must be accompanied by 'flow-ethtype'
```

pathtrace ethernet evpn

Use this command to find the exact path of a flow between leaf nodes. This CLI provides flow parameters to simulate the host packet arriving from the access port.

Command Syntax

```
pathtrace ethernet evpn ((MEPID domain DOMAIN_NAME ma MA_NAME)|) destination
  A.B.C.D EVPNID (interface-ip-address|) ( flow-inport IFNAME flow-dst
  XXXX.XXXX.XXXX   flow-vlan (VLANID (flow-innervlan VLANID|)|untagged)({flow-src
  XXXX.XXXX.XXXX | flow-ethtype ETHTYPE}|)| IFNAME| )
```

Parameters

`MEPID` Enter the MEP identifier <1-8191>. (Optional parameter in implicit mode.)

`domain DOMAIN_NAME`

Enter the name of the domain. Name length can be maximum of 43 characters. (Optional parameter in implicit mode.)

`ma MA_NAME` Enter maintenance association name. Name length can be maximum of 6 characters. (Optional parameter in implicit mode.)

`destination A.B.C.D`

Specify the destination VTEP IPv4 address.

`EVPNID` Enter the EVPN identifier in range <1-16777215>

`interface-ip-address`

To fetch the ingress interface IP addresses of spine and leaf.

`flow-inport IFNAME`

Specify the access port name to simulate host packet arriving from the access port.

`flow-dst XXXX.XXXX.XXXX`

Specify the destination MAC address to simulate host packet arriving from the access port.

`flow-vlan VLANID`

Specify the outer VLAN ID mapped to the `flow-inport` interface to simulate host packet arriving from access port.

`flow-innervlan VLANID`

Specify the Inner VLAN ID to simulate host packet arriving from the access port.

`untagged` Specify untagged host packet arriving from access port.

`flow-src XXXX.XXXX.XXXX`

Specify the source MAC address to simulate host packet arriving from the access port.

`flow-ethtype ETHTYPE`

Specify the Ethertype value (0x600 - 0xFFFF) to simulate host packet arriving from the access port.

`IFNAME`Specify the PTM packet outgoing interface name.

Command Mode

Privilege execute mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

To verify pathtrace with outgoing interface

```
#pathtrace ethernet evpn destination 5.5.5.5 30 vlan1.65
```

To verify pathtrace with CFM objects

```
#pathtrace ethernet evpn 10 domain testdomain ma testma destination 10.10.10.10 100
```

To verify pathtrace with flow parameters:

```
#pathtrace ethernet evpn 21 domain T_MD ma ma1 destination 1.1.1.1 10 flow-inport xe1 flow-dst 0010.7700.0001 flow-vlan 102 flow-ethtype 0x8100 ce2
% Warning: Applicability of given flow parameter(s) on the packet, depends on the load-balance rtag7 setting in the system

      Sent 2 PTM packets, attempted 2 hops
      ======================================================================
      TTL  Reply                Ingress    Egress     State      rtt (ms)
      ----------------------------------------------------------------------
      1    21.10.1.2            ce4        ce2        UP / UP    2.787
      2    1.1.1.1              ce50       null       UP         2.875
```

Note: To check the path of a particular flow `flow-inport`, `flow-dst` and `flow-vlan` are mandatory parameters irrespective of the load-balancing configuration.

Note: In Q2A devices, it is mandatory to use both `flow-vlan` and `flow-ethtype` parameters to get the exact VxLAN packet path. The following message is displayed to the user when not used together:

```
      #pathtrace ethernet evpn destination 4.4.4.4 2 flow-inport xe11 flow-dst 0000.0241.0002 flow-vlan 2
      %% Option 'flow-vlan' must be accompanied by 'flow-ethtype'
```

show ethernet cfm ccm-status

Use this command to display the CCM status between the VTEPs.

Command Syntax

```
show ethernet cfm ccm-status mep MEPID domain DOMAIN_NAME ma-name <string>
```

Parameters

`MEPID`	MEP identifier <1-8191>. (Optional parameter in implicit mode.)
`domain DOMAIN_NAME`	Enter the name of the domain. Name must be of 5 characters if type is character-string otherwise no_name if domain-type is no-name1. (Optional parameter in implicit mode.)
`ma MA_NAME`	Enter maintenance association name. If ma-type is character string then maximum length of ma-name is 6 else if it's integer then maximum is 2-octets. (Optional parameter in implicit mode)

Command Mode

Privilege execute mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

```
#configure terminal
#show ethernet cfm ccm-status mep 43 domain mdom2 ma-name ma2
Remote MEP Identifier:     34  State: Up
Remote MEP: 3.3.3.3 VNID:     2
Last flapped: NA
Last recovered: NA
Before flap: Good sequence id:    NA  flow id: NA
After flap: Good sequence id:    NA  flow id: NA
Remote defect indication: False
```

nvo vxlan-oam

Use this command to enable the VxLAN OAM feature. Executing this command installs the necessary hardware filters.

Command Syntax

```
nvo vxlan-oam (|spine)
```

Parameters

None

Command Mode

Configuration Mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

Executing the CLI on the leaf node creates implicit mode CFM objects.

```
#configure terminal
#nvo vxlan-oam
```

Executing the CLI on the spine node creates Implicit mode CFM objects.

```
#configure terminal
#nvo vxlan-oam spine
```

Verifying ping in implicit mode.

```
ping ethernet epvn destination 10.10.10.1 100 xe3
```

Executing the following CLIs on the spine node creates explicit mode CFM objects.

```
  (config)#nvo vxlan-oam spine
  (config)#ethernet cfm domain-type character-string domain-name testmd level 0
mip-creation static
  (config-ether-cfm)#service ma-type string ma-name testma
```

Verifying pathtrace in explicit mode

```
      pathtrace ethernet evpn 10 domain testdomain ma testma destination 10.10.10.1
100 flow-inport xe58 flow-dst 0000.0001.0002 flow-vlan 105 flow-src 0000.aaaa.bbbb
```

show ethernet cfm nvo vxlan-oam

Use this command to display the CFM objects created, including the CFM objects.

Command Syntax

```
show ethernet cfm nvo vxlan-oam
```

Parameters

None

Default

NA

Command Mode

Configuration Mode

Applicability

This command was introduced in OcNOS version 7.0.0.

Example

```
#configure terminal
#show ethernet cfm nvo vxlan-oam

------------------------------------------------------------
md-name   : NVO3BaseMode
ma-name   : 65532
name-type    level  mip-perm ma-evpn-id ma-name-type mip-level
------------------------------------------------------------
char-string  3      static   16777215   2-oct-int    3
------------------------------------------------------------
md-name   : mdom2
ma-name   : ma2
name-type    level  mip-perm ma-evpn-id ma-name-type mip-level
------------------------------------------------------------
char-string  7      static   16777215   char-string  7
```

Modified CLIs

The following existing CLIs have been modified:

- ethernet cfm domain-type
- ethernet cfm mep
- show ethernet cfm statistics

ethernet cfm domain-type

Use this command to create a CFM Maintenance Domain (MD) and enter into Ethernet CFM mode. The CLI is modified to add a static option to the MD.

Command Syntax

```
ethernet cfm domain-type (no-name|character-string) domain-name DOMAIN_NAME level
<0-7>  mip-creation (none|default|explicit|static)
no ethernet cfm domain-name DOMAIN_NAME
```

Parameters

`static`	Creates static MIP. It is used in Explicit Mode where the Spine participates actively by responding to pathtrace requests.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 3.0. The parameter **`static`** is introduced in OcNOS version 7.0.0.

Examples

To create a static MD.

```
OcNOS(config)#
OcNOS(config)#ethernet cfm domain-type character-string domain-name 12347
level 7 mip-creation static
OcNOS(config-ether-cfm)#
```

ethernet cfm mep

The CLI is modified to add a remote-vtep <VTEP-IP> parameter.

Command Syntax

```
ethernet cfm mep (down|up) mpid <1-8191> active (false|true) (IFNAME ((vlan <1-
  4094> remote-vtep <VTEP-IP> (inner-vlan <1-4094>)|)| evpn <1-16777215>)
no ethernet cfm mpid <1-8191>
```

Parameters

`remote-vtep <VTEP-IP>` Specify the remote VTEP IP address.

Command Mode

Ethernet CFM MA mode

Applicability

This command was introduced in OcNOS version 3.0. The parameter `remote-vtep <VTEP-IP>` is introduced in OcNOS version 7.0.0.

Example

To configure a Down MEP for the connectivity status of a remote VTEP.

```
      #config-ether-cfm-ma)#ethernet cfm mep down mpid 13 active true remote-vtep
3.3.3.3
      #(config-ether-cfm-ma-mep)#cc multicast state enable
      #(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
```

show ethernet cfm statistics

Use this command to display CFM statistics: CCM sent and received, LBM sent and LBR received, LTM sent and LTR received. This CLI is modified to include pathtrace messages PTM and PTR statistics.

Command Syntax

```
show ethernet cfm statistics mep MEPID domain DOMAIN_NAME ma-name MA_NAME
```

Applicability

This command was introduced before OcNOS version 3.0. The output is modified to add path trace in OcNOS version 7.0.0

Example

To display CCM, Ping and Pathtrace statistics using CFM command.

```
C1-Leaf1#show ethernet cfm statistics

NOTE: CCM stats will not be counted when loss session is active for the MEP

Continuity Check Messages
   CCM Sent                : 2421
   CCM Received            : 562
   CCM Drop                : 0

Loop Back Messages
   LBM Sent                : 0
   LBR Sent                : 0
   LBM Drop                : 0
   LBR Drop                : 0
   LBR Received(Valid)     : 0
   LBR Received(Bad msdu)  : 0
   LBR Received(Out-of-Seq): 0

Link Trace Messages
   LTM Sent                : 0
   LTR Sent                : 0
```

```
    LTR Received(Valid)      : 0
    LTR Received(unexpected): 0

Path Trace Messages
    PTM Sent                    :0
    PTM Received                :0
    PTM Received (bad)          :0
    PTR Sent                    :0
    PTR Received (Valid)        :0
    PTR received (Unexpected)   :0
    PTM Drop                    :0
```

Troubleshooting

When troubleshooting the VxLAN OAM feature, focus on verifying the core EVPN and CFM components, as these form the foundation for OAM messages. Here are some suggested troubleshooting tips, grouped by logical area:

General OAM and Mode Checks

Confirm that VxLAN OAM is enabled on both the VTEPs and Spines.

- VTEP (Leaf): Ensure `nvo vxlan-oam` is configured.
- Spine: Ensure `nvo vxlan-oam spine` is configured.

Recognize the functionality differences based on mode.

- Implicit Modes: Only ping and pathtrace are supported; CCM is not supported.
- Explicit Mode: All functions (ping, pathtrace, and CCM are supported. If CCM fails, ensure you are in explicit mode and CCM is not disabled per MEP.
- Verify the necessary VxLAN hardware capability (refer to Prerequisites) is enabled on the VTEPs.

Connectivity Faults (Ping and CCM)

Before troubleshooting the VxLAN overlay, ensure the underlay network can route IP traffic between all VTEP-IP addresses (loopbacks).

- Check BGP neighbor status on VTEPs and Spines to confirm IPv4 unicast sessions (Underlay peer-group) are Up.

Use OAM Ping Command: Use the specific VxLAN OAM ping command to test reachability over the L2 VxLAN tunnel.

- Example: `ping ethernet evpn destination 2.2.2.2 2 po12`
- If ping fails, run a `pathtrace` to isolate the point of failure (see below).

Use the detailed CCM status command to find recent flaps, the remote MEP ID, and RDI (Remote Defect Indication) status.

- Command: `show ethernet cfm ccm-status mep <MEPID> domain mdom2 ma-name ma2.`
- A remote defect indication of True means the peer MEP is reporting a defect.
- Verify that all expected remote MEPs are configured for cross-check on the local VTEP.

Path Isolation (Pathtrace Faults)

Use the OAM `pathtrace` command to determine the exact path the VxLAN packet takes, identifying which Spine (MIP) or VTEP is the point of failure.

- Command Structure: `pathtrace ethernet evpn destination A.B.C.D EVPNID....`

If `pathtrace` stops at a Spine, confirm that the Spine has the matching MD and MA configured with `mip-creation static`.

If troubleshooting a specific flow, the `flow-inport, flow-dst`, and `flow-vlan` parameters are mandatory to select the correct ECMP path.

- For Qumran 2A devices, both `flow-vlan` and `flow-ethtype` are mandatory to get the exact path

Configuration Discrepancy Checks

A local MEP must be configured as `active true` for CCM, `ping` and `pathtrace` to work.

- The `remote-vtep` IP in the MEP configuration must match the actual VTEP-IP of the remote device.

Ensure the MD name, MA name, level, and the linked EVPN ID are configured identically on all VTEPs and spines.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
CCM	Continuity Check Message. One-way, periodic messages exchanged between configured MEPs to detect loss of connectivity, similar to keepalive or heartbeat messages.
CFM	Connectivity Fault Management. The technology (often IEEE 802.1ag or similar) providing OAM functions like Continuity Check Messages (CCMs).
CLOS Network Topology	A multi-stage switching network design, typically used in data centers, which includes Leaf and Spine layers.
Data Center	A facility where networking (like Leaf and Spine CLOS topology) is deployed, with L2 and L3 VxLAN overlay packets forwarded over an IP underlay.
Decapsulation	The process where a target leaf switch removes the VxLAN tunnel header to forward the original (inner) packets to the destination server.
ECMP	Equal-Cost Multi-Path. A routing strategy that provides multiple redundant paths for traffic forwarding between leaf and spine switches.
Encapsulation	The process where a leaf switch wraps original packets into VxLAN tunnels for transit across the underlay network.
EVPN	Ethernet VPN. A control-plane protocol (typically BGP-based) used in the VxLAN overlay for MAC/IP advertisement and learning.
Explicit Mode	The operational mode for VTEPs/Leaves when there is explicit MEP configuration (active/inactive) on the VxLAN tunnel, supporting ping, pathtrace, and CCM.
Fault Isolation	An OAM service function that utilizes trace messages (pathtrace) to identify and isolate the exact path of a flow between leaf nodes.
Fault Verification	An OAM service function that uses loopback messages (ping) to verify network reachability between leaf nodes.

Implicit Mode	The operational mode when no explicit MEP configuration exists, where Maintenance Domain (MD), Maintenance Association (MA), and MEP objects are implicitly created. Only ping and pathtrace are supported.
Leaf	A layer of switches in the CLOS topology (also acting as VTEPs) connected to server nodes and spine switches.
Loopback Function (Ping)	An OAM function used for fault verification; the messages (interchangeable with ping) are used to detect errors and path failures.
MA	Maintenance Association. A service defined within an MD, typically associated with a specific EVPN ID/VNID.
MD	Maintenance Domain. A management boundary within the VxLAN OAM network topology.
MEP	Maintenance End Point. A logical endpoint of an OAM Maintenance Entity located at a VTEP. It generates, transmits, and terminates OAM frames.
MEP Placement (Downward)	MEP direction that monitors the physical (underlay) network, the direct path.
MEP Placement (Upward)	MEP direction that monitors the VxLAN tunnel (overlay), across the VxLAN fabric.
MIP	Maintenance Intermediate Point. A passive point inside the VxLAN overlay path, typically a spine node. It can respond to OAM requests like pathtrace to trace the overlay path hop by hop.
Overlay	The logical network is built over the physical underlay, which uses VxLAN to tunnel L2 and L3 packets. VxLAN OAM functions like MEPs and MIPs exist within the overlay.
Pathtrace Function (Traceroute)	An OAM function used for fault isolation. It finds the list of switches traversed by a frame to reach the destination.
Spine	A layer of switches in the CLOS topology connected to the leaf nodes; they often act as MIPs in the VxLAN OAM architecture.
Underlay	The physical IP network infrastructure (Leaf and Spine) over which the VxLAN overlay packets are forwarded.
VNI	VxLAN Network Identifier. A 24-bit identifier that uniquely specifies a logical VxLAN segment. OAM is tenant-specific and exists per VNI.
VTEP	VxLAN Tunnel Endpoint. A network node, typically a leaf switch, that handles the encapsulation and decapsulation of VxLAN packets. Connectivity is verified between VTEPs using OAM.
VxLAN OAM	VxLAN Operations, Administration, and Maintenance. The set of services and functions necessary for verifying connectivity, detecting faults (loss, delay, misconfiguration), and ensuring service level agreements (SLAs) for VxLAN overlay networks.

VxLAN Command Reference

CHAPTER 1 VXLAN Commands

This chapter describes the VXLAN commands:

- nvo vxlan vtep-info
- nvo vxlan vtep-ip-global
- show nvo vxlan
- show nvo vxlan access-if-config
- show nvo vxlan arp-cache
- show nvo vxlan counters access-port
- show nvo vxlan counters network-port
- show nvo vxlan mac-table
- show nvo vxlan nd-cache
- show nvo vxlan static host state
- show nvo vxlan tunnel
- show nvo vxlan vni-tunnel
- show running-config nvo vxlan
- show evpn esi
- show evpn multi-homing all
- show evpn multihoming-status
- show nvo vxlan route-count
- show nvo vxlan vni-name
- show hsl evpn multihoming esi
- shutdown
- vxlan host-reachability-protocol evpn-bgp
- vlan-xlate-1 large

arp-cache disable

Use this command to disable the ARP cache for MAC/IP.

When the ARP cache is disabled on a VxLAN access port, OcNOS does not reply to any ARP arriving on this port from the cache. OcNOS withdraws all MAC/IPs configured/learned on this access port and removes the MAC/IP entry for this access port from the local ARP cache.

OcNOS also makes sure that on withdrawing the MAC/IP route, the MAC does not become unknown. If all routes for this MAC are being withdrawn because of this command, then OcNOS advertises a MAC-only route. This is done so that the MAC does not become unknown and only the cache functionality becomes disabled.

Use the `no` form of this command to enable ARP cache for MAC/IP.

Note: On enabling the cache, an IP will be in conflict, then the cache enable will fail. The conflict has to be manually removed and then the cache enabled.

Command Syntax

```
arp-cache disable
```

```
no arp-cache disable
```

Parameters

None

Default

By default, the arp-cache option is enabled.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#arp-cache disable
(config-nvo-acc-if)#exit
```

arp-nd flood-suppress

Use this command to *completely* restrict the flood of ARP/ND packets towards remote VTEPs or other access ports.

This command applies only when the ARP cache and ND cache are enabled. When the ARP cache is disabled, ARP flooding is not suppressed even if this command is given. When the ND cache is disabled, ND flooding is not disabled, even if this command is given.

Use the `no` form of this command to not restrict the flood of ARP/ND packets.

Command Syntax

```
arp-nd flood-suppress
no arp-nd flood-suppress
```

Parameters

None

Default

By default, the arp-nd flood-suppress option is disabled.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#arp-nd flood-suppress
(config-nvo-acc-if)#exit
```

arp-nd refresh timer

Use this command to configure aging out the arp-cache and nd-cache entries for given time multiplied by 3 in secs

Use the no form of this command to remove the configuration.

Command Syntax

```
nvo vxlan arp-nd refresh-timer <3-190> mac (XX-XX-XX-XX-XX-
  XX|XX:XX:XX:XX:XX:XX|XXXX.XXXX.XXXX)
no nvo vxlan arp-nd refresh-timer
```

Parameters

`<3-190>` refresh time in seconds

`XX-XX-XX-XX-XX-XX`

v-mac is mandatory for MH

Command Mode

Config mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#config mode
(config)#nvo vxlan arp-nd refresh-timer 100
(config)#no nvo vxlan arp-nd refresh-timer
```

Example to configure in MH node

```
(config)#nvo vxlan arp-nd refresh-timer 100 mac 0000.1111.2222
```

clear mac address table dynamic vxlan

Use this command to clear dynamically learned MACs.

Command Syntax

```
clear mac address table dynamic vxlan
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear mac address table dynamic vxlan
```

clear nvo vxlan counters

Use this command to clear the counters of access ports or network ports.

Command Syntax

```
clear nvo vxlan counters((access-port (port IFNAME | port-vlan IFNAME (VLAN_ID |
  outer-vlan) | all)) | (network-port (dst A.B.C.D | all)))
```

Parameters

`port`	Port
`IFNAME`	Interface name
`port-vlan`	VLAN port
`IFNAME`	Interface name
`VLAN_ID`	VLAN identifier
`Outer-vlan`	Outer Vlan
`A.B.C.D`	Tunnel destination IPv4 address
`all`	All access or network ports

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Example for clearing a VLAN port counter:

```
#clear nvo vxlan counters access-port port-vlan xe1 2
```

Example for clearing all access port counters:

```
#clear nvo vxlan counters access-port all
```

Example for clearing network port counters:

```
#clear nvo vxlan counters network-port dst 1.1.1.1
```

Example for clearing all network port counters:

```
#clear nvo vxlan counters network-port all
```

clear nvo vxlan tunnels

Use this command to clear the nvo vxlan tunnels to re-establish the tunnel after mapping/un-mapping the QoS profile to vxlan tunnel.

Command Syntax

```
clear nvo vxlan tunnels (|dst-ip A.B.C.D)
```

Parameters

`dst-ip`	VXLAN tunnel destination
`A.B.C.D`	destination IPv4 address

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
#clear nvo vxlan tunnels dst-ip 2.2.2.2
```

Note: `clear nvo vxlan tunnels` This command will clear all the VXLAN tunnels destination.

`clear nvo vxlan tunnels dst-ip A.B.C.D` - This command to clear individual tunnel destination(i.e A.B.C.D).

clear nvo vxlan mac-stale-entries

Use this command to clear MAC entries that are in discard state in the forwarding database.

Command Syntax

```
clear nvo vxlan mac-stale-entries (vnid <1-16777215> |)
```

Parameters

`<1-16777215>`	VXLAN network identifier

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear nvo vxlan mac-stale-entries vnid 100
```

description

Use this command to set a description for a port.

Use the `no` form of this command to remove the description for a port.

Command Syntax

```
description LINE
no description
```

Parameters

`LINE`	Maximum 32 characters describing this port.

Default

No default value is specified for description LINE commands.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#description member-port xe1 with vlan 2
(config-nvo-acc-if)#exit
```

disable-l3-termination

Use this command to disable l3 termination for a MAC address of a particular VPN ID.

Command Syntax

```
disable-l3-termination XXXX.XXXX.XXXX
```

Parameters

XXXX.XXXX.XXXX	Host MAC address

Command Mode

NVO Mode

Applicability

This command is introduced in OcNOS version 6.5.4 and OcNOS version 6.6.0.

Example

```
#configure terminal
(config)#nvo vxlan id 123241
(config-nvo)#disable-l3-termination 0000.1111.2222
```

dynamic-learning disable

Use this command to disable dynamic learning of MACs at the access port. This command also disables dynamic learning of MAC/IP from ARP/ND messages received on this access port.

Use the `no` form of this command to enable dynamic learning of MACs at the access port.

Command Syntax

```
dynamic-learning disable
no dynamic-learning disable
```

Parameters

None

Default

By default, the dynamic-learning option is enabled.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#dynamic-learning disable
(config-nvo-acc-if)#exit
```

encapsulation

Use this command to assign a Tag Protocol Identifier (TPID) to an access port.

Use the `no` form of this command to set the default TPID (0x8100: IEEE 802.1Q VLAN-tagged frame) to an access port.

Command Syntax

```
encapsulation TPID
no encapsulation
```

Parameters

`TPID`	Tag Protocol Identifier: Ox88A8: IEEE 802.1ad Provider Bridging Ox9100: IEEE 802.1Q VLAN-tagged frame with double tagging

Default

The encapsulation TPID default is 0X8100.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#encapsulation 0x9100
(config-nvo-acc-if)#exit
```

evpn esi holdtime

Use this command to allow some time for the tunnels to come at the time of vxlan initialization before making the esi up. This avoids traffic to be black-holed when a new PE is added and connected to an already running CE for multi-homing.

Use the no form of this command to make the esi up immediately when configuring the access-if cli.

Command Syntax

```
evpn esi holdtime <10-300>
no evpn esi holdtime <10-300>
```

Parameters

<10-300>	Hold time in seconds

Default

Default value is 0.

Command Mode

Configuration Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)# evpn esi holdtime 100
(config)# exit
```

evpn-mac-holdtime

Use this command to delay route withdraw during downlink failure. Configure it under parrent access interface.

Command Syntax

```
evpn-mac-holdtime <1-600>
```

Parameters

<1-600>	MAC Hold time range from 1 to 600 in seconds.

Default

Default value is 0.

Command Mode

Interface Mode

Applicability

This command was introduced before OcNOS version 6.5.3.

Example

```
#configure terminal
(config)# interface xe1
(config-if)# evpn-mac-holdtime 100
(config-if)# exit
```

evpn vxlan multi-homing enable

Use this command to enable evpn vxlan multi-homing

Use the no form of this command to disable evpn vxlan multi-homing.

Note: Node will have to be restarted for this to be applicable. If there are some nodes in topology which have multi-homed CEs, then nodes which do not have multi-homed CEs should also enable multihoming so, that they can load share traffic to the multi-homed CEs.

Command Syntax

```
evpn vxlan multihoming enable
no evpn vxlan multihoming enable
```

Parameters

None

Default

By Default Multi-homing will be d]isabled

Command Mode

Configuration Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
#(config)# evpn vxlan multihoming enable
#(config)# exit
```

evpn multi-homed

Use this command to configure interfaces as multi-homed and configure esi-value in case of physical and static lag and system-mac in case of Dynamic lag.

Use the no parameter of this command to unconfigure multi-homed on the interface.

Command Syntax

```
evpn multi-homed (esi XX:XX:XX:XX:XX:XX:XX:XX:XX | system-mac (XX-XX-XX-XX-XX-
  XX|XX:XX:XX:XX:XX:XX|XXXX.XXXX.XXXX)
no evpn multi-homed (esi | system-mac)
```

Parameters

`XX:XX:XX:XX:XX:XX:XX:XX:XX:XX`	ESI value in HH:HH:HH:HH:HH:HH:HH:HH:HH:HH - 9 octet format
`XX-XX-XX-XX-XX-XX`	Host MAC address (Option 1)
`XX:XX:XX:XX:XX:XX`	Host MAC address (Option 2)
`XXXX.XXXX.XXXX`	Host MAC address (Option 3)

Default

Default value is 0.

Command Mode

Interface Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)# interface xe1
(config)# evpn multi-homed esi 00:11:22:33:44:55:66:77:88
(config)# exit

configure terminal
(config)# interface po1
(config)# evpn multi-homed system-mac 0000.0000.1111
(config)# exit
```

evpn-vlan-service

Use this command to configure EVPN-Service type.

Use the `no` form of this command to delete the evpn vlan service.

Note: If access port mappings to vnid exists already and VLAN service is configured later and mapped to tenant then we should not allow the mapping, user should either remove and reconfigure the access port mappings.

Command Syntax

```
evpn-vlan-service (vlan-aware-bundle | vlan-based)
no evpn-vlan-service
```

Parameters

`vlan-based`	VLAN-based mapping of a mac-vrf to single VNI (1 to 1)
`vlan-aware-bundle`	VLAN-aware-bundle mapping of a mac-vrf to multiple VNI (1 to many)

Default

By default, `vlan-aware-bundle` service is applied when we no `evpn-vlan-service` configuration is specified explicitly.

Command Mode

MAC vrf mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#mac vrf vrf1
(config-vrf)# evpn-vlan-service vlan-based
(config-vrf)# no evpn-vlan-service vlan-based
```

garp-gna enable

Use this command to enable GARP/GNA packets per vxlan on the outgoing access interfaces where it is configured.

Use the no form this command to disable GARP/GNA packets on the outgoing access interfaces.

By default it will be disable

Command Syntax

```
garp-gna enable
```

```
garp-gna enable
```

Parameters

```
None
```

Command Mode

config-nvo-acc-if mode

Applicability

This command is introduced in OcNOS version 1.3.6.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 10
config-nvo-acc-if)#garp-gna enable
```

load-balance rtag7 vxlan inner-l2

Use this command to enable RTAG7 load balancing for VxLAN for inner-L2 packets on the access side for load sharing of traffic on tunnel paths. This command also enables random UDP source port generation for the VxLAN UDP tunnels based on the L2 frame being received on the VxLAN access port.

Use the `no` parameter of this command to disable load balance for RTAG7 for VxLAN for inner-L2 frames.

Command Syntax

```
load-balance rtag7 vxlan inner-l2 (dest-mac | src-mac)
no load-balance rtag7 vxlan inner-l2 (dest-mac | src-mac)
```

Parameter

`dest-mac`	Destination MAC Address.
`src-mac`	Source MAC Address.

Default

By default, the load-balance rtag7 vxlan inner-l2 default is src-mac.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#load-balance rtag7 vxlan inner-l2 dest-mac

(config)#no load-balance rtag7 vxlan inner-l2 dest-mac
```

load-balance rtag7 vxlan inner-l3

Use this command to enable RTAG7 load balancing for VxLAN for inner-L3 IPv4 and IPv6 packets on the access side for load sharing of traffic on tunnel paths. This command also enables random UDP source port generation for the VxLAN UDP tunnels based on the L3 frame being received on the VxLAN access port.

Use the `no` parameter of this command to disable load balance for RTAG7 for VxLAN for inner-L3 IPv4 and IPv6 frames.

Command Syntax

```
load-balance rtag7 vxlan inner-l3 (dest-ip | destl4-port | protocol-id | src-ip |
  srcl4-port)
no load-balance rtag7 vxlan inner-l3 (dest-ip | destl4-port | protocol-id | src-ip
  | srcl4-port)
```

Parameter

`dest-ip`	Destination IP.
`destl4-port`	Destination l4 port.
`protocol-id`	Protocol (IPv4).
`src-ip`	Source IP.
`srcl4-port`	Source l4 port.

Default

By default, the load-balance rtag7 vxlan inner-l3 default is src-ip.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#load-balance rtag7 vxlan inner-l3 dest-ip

(config)#no load-balance rtag7 vxlan inner-l3 dest-ip
```

mac

Use this command to associate a static MAC address and a static IPv4 or IPv6 address on an access interface.

Use the `no` form of this command to disassociate a static MAC address and an IPv4 or IPv6 address for an access interface.

Note: When a static host is configured on an access port which is in the down state, its state is Inactive.

Note: The same static mac configuration is not allowed on a different access port as then there will be a chance of conflict. However, if a dynamic packet is sent at another access port which is up and running with the same MAC, it learns as usual. As soon as the port on which the static MAC is configured comes up, static learning is given precedence and the dynamically learned MAC is moved to the port where it is configured statically.

Command Syntax

```
mac XXXX.XXXX.XXXX
mac XXXX.XXXX.XXXX (ip A.B.C.D | ipv6 X:X::X:X)
no mac XXXX.XXXX.XXXX (ip A.B.C.D | ipv6 X:X::X:X)
```

Parameters

`XXXX.XXXX.XXXX`	Static MAC address. The following formats are supported:
	XX-XX-XX-XX-XX-XX Source MAC address (Option 1)
	XX:XX:XX:XX:XX:XX Source MAC address (Option 2)
	XXXX.XXXX.XXXX Source MAC address (Option 3)
`A.B.C.D`	Static IPv4 address.
`X:X::X:X`	Static IPv6 address.

Default

No default value is specified for mac command.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#mac 0000.0000.aaaa ip 10.10.10.1
(config-nvo-acc-if)#mac 0000.0000.aaaa ipv6 1201::1
(config-nvo-acc-if)#exit
```

mac-ageing

Use this command to enable or disable MAC ageing based on either “src or dest” hit or “src only” hit

Command Syntax

```
mac-ageing (dst-src|src)
```

Parameter

`dst-src`	Enable destination-or-source MAC hit based ageing
`src`	Enable source MAC hit based ageing

Default

Source

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 6.6.0.

Example

```
#configure terminal
(config)#mac-ageing src
```

mac vrf

Use this command to create a MAC VRF to use in EVPN routes.

See also vxlan host-reachability-protocol evpn-bgp.

Use the `no` parameter of this command to delete the MAC VRF.

Command Syntax

```
mac vrf WORD
no mac vrf WORD
```

Parameter

`WORD`	MAC routing or forwarding instance name.

Default

No default value is specified for mac vrf WORD command.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#mac vrf vrf1

(config)#no mac vrf vrf1
```

mac-holdtime

Use this command to set the MAC hold time for a MAC/IP or MAC.

The feature holds the MAC in hardware until BGP has withdrawn from the neighbors. This helps to reduce flooding to other access ports.

This setting applies when the access port is shut down, the physical port on which the access port is down, or the access port is removed from the VNID using the `no` form of the map vnid command.

When the MAC hold time is configured as -1, then the MAC is not removed from the hardware and is also not withdrawn from EVPN BGP.

Use the `no` form of this command to remove the MAC hold time for the MAC/IP or MAC.

Note: When a MAC is moved to discard state, traffic to and from this MAC is discarded. This is applicable only on statically configured MAC/MAC-IPs.

Command Syntax

```
mac-holdtime <-1-300>
no mac-holdtime
```

Parameters

`<-1-300>`	MAC hold time in seconds. Specify -1 to “never expire”.

Default

The default holdtime for mac is 3 seconds.

Command Mode

NVO mode and NVO_ACC_IF_MODE mode

Note: When configured in both modes, then the NVO_ACC_IF_MODE value takes preference for that access port.

Applicability

This command was introduced before OcNOS version 1.3.

This command in NVO_ACC_IF_MODE mode is introduced in OcNOS version 1.3.4.

Example

```
#configure terminal
(config)#nvo vxlan id 3 ingress-replication inner-vid-disabled
(config-nvo)#mac-holdtime -1
(config-nvo)#exit
```

map vnid

Use this command to map a tenant to an access-port.

Use the `no` form of this command to remove the tenant from an access-port

Command Syntax

```
map vnid <1-16777215>
no map vnid <1-16777215>
```

Parameters

`<1-16777215>`	VxLAN network identifier.

Default

No default value is specified for map vnid command.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#map vnid 100
(config-nvo-acc-if)#exit
```

nd-cache disable

Use this command to disable ND cache for MAC/IPv6.

When the ND cache is disabled on a VxLAN access port, OcNOS does not reply to any ND arriving on this port from the cache. OcNOS withdraws all MAC/IPs configured/learned on this access port and removes the MAC/IP entry for this access port from the local ND cache.

OcNOS also makes sure that on withdrawing the MAC/IP route, the MAC does not become unknown. If all routes for this MAC are being withdrawn because of this command, then OcNOS advertises a MAC-only route. This is done so that the MAC does not become unknown and only the cache functionality becomes disabled.

See also arp-cache disable.

Use the `no` form of this command to enable ND cache for MAC/IPv6.

Note: On enabling the cache, an IP will be in conflict, then the cache enable will fail. The conflict has to be manually removed and then the cache enabled.

Command Syntax

```
nd-cache disable
no nd-cache disable
```

Parameters

None

Default

By default, the nd-cache option is enabled.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#nd-cache disable
(config-nvo-acc-if)#exit
```

no nvo vxlan

Use this command to delete a given tenant or all tenants/VPNs.

Command Syntax

```
no nvo vxlan (id <1-16777215>|all-vnids)
```

Parameters

`<1-16777215>`	VXLAN Network Identifier.
`all-vnids`	Delete all VxLAN network identifiers.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#no nvo vxlan id 100
```

nvo vxlan

Use this command to enable or disable VxLAN.

Command Syntax

```
nvo vxlan (enable | disable)
```

Parameters

None

Default

By default, the nvo vxlan option is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan enable

(config)#nvo vxlan disable
```

nvo vxlan id

Use this command to add a tenant and the type of VPN. This command changes the mode to NVO mode.

Use `no` form of this command to unconfigure the VXLAN ID.

Before using this command, enable VXLAN by using the nvo vxlan command.

Command Syntax

```
nvo vxlan id <1-16777215> (multicast |default) | (etree-leaf|) ((ingress-
  replication (inner-vid-disabled | bridge-vlan VLAN-ID | etree-leaf)
no nvo vxlan id <1-16777215>
```

Parameters

`default`	Not a specific tenant which falls default type. It is not supported.
`vxaln id <1-16777215>`	Specifies the VXLAN Network Identifier (VNID) to identify the virtual VXLAN overlay network segment.
`multicast`	(Optional) Specifies a multicast mode (point to multipoint) used for transmitting VXLAN encapsulated multicast packets. It is not supported.
`ingress-replication`	(Optional) Enables head-end replication for forwarding BUM traffic. It is not supported.
`bridge-vlan VLAN-ID`	(Optional) Specifies the VLAN ID.
`inner-vid-disabled`	(Optional) Disables the transmission of the VLAN ID (VID) with traffic leaving the network port.
`etree-leaf`	(Optional) Configures the device as a leaf node within the E-Tree topology, implying its role in forwarding BUM traffic within the E-Tree service model.
`xconnect target-vxlan-id <1-16777215>`	Specifies the cross-connect Identifier.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3. Introduced the `etree-leaf` parameter in the OcNOS version 6.5.1

Example

```
#configure terminal
(config)#nvo vxlan id 300 ingress-replication
(config-nvo)#exit

(config)#nvo vxlan id 200 ingress-replication inner-vid-disabled
(config-nvo)#exit
```

To configure a VXLAN instance with VXLAN ID for ingress replication, with inner VLAN ID disabled, and as an E-Tree leaf node:

```
(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
(config-nvo)#exit
```

To configure a VXLAN instance with VLAN-VNID mapping for ingress replication on the leaf node:

```
(config)#nvo vxlan id 101 ingress-replication bridge-vlan 101
(config-nvo)#exit
```

nvo vxlan access-if

Use this command to map a complete interface or a VLAN or VLAN range on an interface to identify the tenant traffic and to enter NVO access interface mode.

The command "nvo vxlan access-if port <if_name> default" accepts all tagged, double tagged and untagged traffic received on the mapped physical port.

Use the `no` form of this command to unmap an interface or a VLAN.

Note: When a VxLAN access interface configured as a port VLAN as vlan-range or port as default, then arp-cache and nd-cache should be disabled and only VNID with inner-vid-enable is mapped.

Note: Inner-vid-enable is one-to-one mapping. VNID is mapped to only on one access-port.

Note: Vlan Range not allowed to be configured for Vxlan stacked access-port.

Command Syntax

```
nvo vxlan access-if (port IFNAME (| default) | port-vlan IFNAME VLAN_RANGE (|
  inner-vlan <2-4094>))
no nvo vxlan access-if (port IFNAME | port-vlan IFNAME VLAN_RANGE(| inner-vlan <2-
  4094>))
```

Parameters

`port`	A physical port.
`IFNAME`	Interface name (Physical/Static lag/Dynamic lag)
`default`	Default access interface
`port-vlan`	The physical port on which VLANs are configured
`IFNAME`	Interface name (Physical/Static lag/Dynamic lag)
`VLAN_RANGE`	Configure VlanId or Vlan-Range for outer vlan
`Inner-vlan`	Inner-vlan id
`<2-4094>`	VLAN id

Default

By default, the nvo vxlan access-if option is port VLAN ID and VLAN_RANGE option is introduced in OcNOS version 5.0..

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2-10
(config-nvo-acc-if)#exit

#configure terminal
```

```
(config)#nvo vxlan access-if port xe1
(config-nvo-acc-if)#exit

#configure terminal
(config)#nvo vxlan access-if port xe1 default
(config-nvo-acc-if)#exit

#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2 inner-vlan 10
(config-nvo-acc-if)#exit
```

nvo vxlan mac-ageing-time

Use this command to set the dynamically learned MAC aging time.

Use the `no` form of this command to set the age out the MACs in hardware to its default (300 seconds).

Command Syntax

```
nvo vxlan mac-ageing-time <10-572>
no nvo vxlan mac-ageing-time
```

Parameters

`<10-572>`	Ageing time in seconds.

Default

The default age out time is 300 seconds.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan mac-ageing-time 10
```

nvo vxlan max-cache-disable

Use this command to configure the maximum number of ARP and ND cache disables on access ports configured with the port+VLAN options. This command does not limit the ARP and ND cache disables on access ports created with only the port option.

Use the `no` form of this command to set the maximum number of ARP and ND cache disables to its default (0).

Note: If any MAC is in conflict when the max cache disable is being unconfigured, then the corresponding caches will not be enabled. This can be enabled after manually, removing the conflict. Caches where there is no conflict, will be enabled.

Command Syntax

```
nvo vxlan max-cache-disable <1-200>
no nvo vxlan max-cache-disable
```

Parameters

`<-1-200>`	Number of ARP/ND cache disable allowed

Default

The default maximum number of ARP and ND cache disables is 0.

Command Mode

Configuration mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan max-cache-disable 10
(config-nvo)#exit
```

nvo vxlan mh-mac-relocate-scan

Use this command to monitor the BCM MAC table in the control plane for MAC addresses nearing age-out and learned as Dynamic Remote. The identified MAC addresses are updated in the BCM to prevent their expiration, ensuring seamless traffic flow. This avoids traffic loss in scenarios where a MAC relocates to a Multi-Homed (MH) pair, but reverse traffic for the MAC continues to arrive at the original node.

Use either `mac-ageing dst-src` or `nvo vxlan mh-mac-relocate-scan enable` for multihoming. recommended.

Command Syntax

```
nvo vxlan mh-mac-relocate-scan (disable|enable)
```

Parameters

None

Default

Disable.

Command Mode

Config Mode

Applicability

This command was introduced before OcNOS version 6.6.0.

Example

```
(config)#nvo vxlan mh-mac-relocate-scan enable
```

show nvo vxlan vni-tunnel

Use this command to display VPN endpoints and its connected tunnel end points.

Command Syntax

```
show nvo vxlan vni-tunnel <| vni>
```

Parameters

vni	VXLAN Network Identifier

Default

None

Command Mode

Configuration Mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

```
Leaf1#show nvo vxlan vni-tunnel 100

VXLAN Network tunnel Entries per VNI
VNI Tunnel-endpoints
=====================================================================
100 34.34.34.34, 32.32.32.32
101 33.33.33.33

Total number of entries are 2
Leaf1# number of entries are 2
Leaf1#
```

nvo vxlan vtep-info

Use this command to configure remote VTEP router information.

Command Syntax

```
nvo vxlan vtep-info A.B.C.D (redund-type, description)
```

Parameters

A.B.C.D	Specify the VTEP's IP Address
description	Specify a description for VTEP.
redund-type	Specify the VTEP redundancy type as:
m-home	Multi homing
mc-lag	MC-LAG
other	Other
s-home	Single home
sh-mclag	Single homing and MC-LAG
sh-mh	Single homing and Multi homing

Command Mode

Configure NVO Vtep Info Mode

Example

```
OCNOS(config)#nvo vxlan vtep-info 2.2.2.2
OCNOS(config_nvo_vtep_info)#redund-type mc-lag
OCNOS(config_nvo_vtep_info)#description OCNOS_mc-lag-2


OCNOS(config)#nvo vxlan vtep-info 2.2.2.2
OCNOS(config_nvo_vtep_info)#no redund-type
OCNOS(config_nvo_vtep_info)#no description
OCNOS(config)#no nvo vxlan vtep-info 2.2.2.2
```

nvo vxlan vtep-ip-global

Use this command to set the source IP address of the VxLAN tunnels.

Use the `no` form of this command to remove the source IP address of the VxLAN tunnels.

Command Syntax

```
nvo vxlan vtep-ip-global A.B.C.D
no nvo vxlan vtep-ip-global A.B.C.D
```

Parameters

`A.B.C.D`	Source VTEP IP address of the global configuration

Default

No default value is specified for nvo vxlan vtep-ip-global command.

Command Mode

NVO mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config-nvo)#nvo vxlan vtep-ip-global 10.10.11.1
```

show nvo vxlan

Use this command to display VXLAN information.

Command Syntax

```
show nvo vxlan (vnid <1-16777215>|)
```

Parameters

<1-16777215>	VXLAN network identifier.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#sh nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     Vni-name     Type Interface ESI                             Vlan DF-
Status Src-addr         Dst-addr
_____________________________________________________________________________
_____________________________________
10       ----         NW   ----      ----                            ---- ----
      1.1.1.1          3.3.3.3
10       ----         NW   ----      ----                            ---- ----
      1.1.1.1          2.2.2.2
10       ----         AC   ce21/1    00:00:11:22:33:44:55:66:77:88 2
DF        ----            ----
20       ----         NW   ----      ----                            ---- ----
      1.1.1.1          3.3.3.3
20       ----         NW   ----      ----                            ---- ----
      1.1.1.1          2.2.2.2
20       ----         AC   ce21/1    00:00:11:22:33:44:55:66:77:88 3    NON-
DF    ----            ----

Total number of entries are 6
```

Table 1-3 explains the fields in the output.

Table 1-3: VxLAN fields

Field	Description
VNID	VXLAN network identifier.
Type	NW - Network Port: VxLAN tunnel AC - Access Port: Host connection

Table 1-3: VxLAN fields

Field	Description
Interface	Name of the Interface.
Vlan	VLAN identifier
Src-addr	Source address in the interface.
Dst-addr	Destination address in the interface.
Total number of entries	The total number of entries listed.

show nvo vxlan access-if-config

Use this command to display the current running configuration of the access interface.

Command Syntax

```
show nvo vxlan access-if-config (LINE|)
```

Parameters

LINE	Access port description.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3,

Added new cli "garp-gna enable" under access-if cli in OcNOS version 1.3.6.

Example

```
#show nvo vxlan access-if-config
nvo vxlan access-if port-vlan xe1 2
map vnid 100
garp-gna enable
access-if-description member-port with xe1 as vlan 2
dynamic-learning disable
arp-nd flood-suppress
arp-cache disable
nd-cache disable
shutdown
mac 0000.0000.1111
mac 0000.0000.aaaa  ip 12.12.12.1
mac 0000.0000.bbbb  ipv6 1201::1
map qos-profile ingress 100
cos 2 egress
     !
      nvo vxlan access-if port-vlan po1 6 inner-vlan 5
      encapsulation 0x9100
      no shutdown
      map vnid 100
     !
```

show nvo vxlan arp-cache

Use this command to display the ARP cache information.

Command Syntax

```
show nvo vxlan arp-cache (vnid <1-16777215>|summary|)
```

Parameters

`<1-16777215>`	VXLAN network identifier.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3, modified to include new fields – Age-Out and Retries-Left in OcNOS version 1.3.5.

Remote Static and Dynamic keyword got changed to Static and Dynamic Remote in MAC table in OcNOS version 1.3.6.

Example

```
#sh nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
ND Timeout :  300 sec   Random-Jitter-Max : 640

VNID     Ip-Addr          Mac-Addr       Type         Age-Out    Retries-Left
_____________________________________________________________________________
10       11.11.11.2       0000.0000.bbbb Dynamic Local  224      2
10       11.11.11.1       0000.0000.aaaa Dynamic Remote----
Total number of entries are 2
```

Table 1-8 explains the output fields.

Table 1-4: ARP cache fields

Field	Description
VNID	VXLAN network identifier
Ip-Addr	IP address of the vxlan
Mac-Addr	Device MAC address
Type	How a host learns a MAC/IP pair: Dynamic: Learned by data plane source learning Evpn: Learned by EVPN Type 2 (MAC/IP) routes Static Remote: Statically configured for remote; used only for static VxLAN, not with EVPN Static Local: Configured on local VTEP

Table 1-4: ARP cache fields

Field	Description
Age=Out	ARP entry expire time
Total number of entries	The total number of entries listed.

show nvo vxlan counters access-port

Use this command to display the receive and transmit counters of an access port including ARP, ND and GARP counters in the same command.

Note: Due to a limitation in the hardware, the transmit packet counters includes the BUM traffic received on that port.

Command Syntax

```
show nvo vxlan counters access-port (port IFNAME | port-vlan IFNAME VLAN_ID)
```

Parameters

`port`	Port Mapping
`IFNAME`	Access port name
`port-vlan`	Port-vlan Mapping
`IFNAME`	Access port name
`VLAN_ID_RANGE`	Vlan Id or Vlan Range <2-4094>
`INNER_VLAN_ID`	Inner-Vlan Id
`all`	All ports and VLANs

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3 and added Control packet (ARP, ND and GARP) counters in OcNOS version 1.3.5. The VLAN_RANGE option is introduced in OcNOS version 5.0.

Example

```
#show nvo vxlan counters access-port port-vlan xe1 2
     Data packets:
      *If ARP/ND cache is enabled, TX doesn't count ARP/ND replies
     from  ARP/ND cache and ARP/ND forwarded after uplifting to
     the control plane.
      RX:  packets  : 2774939
           bytes    : 210553516
      TX:  packets  : 4322274
           bytes    : 326026474

     Control Packets:
      *ARP/ND uplifted and sent/replied from control plane:
     Rx Vxlan Arp discard count           : 0
     Rx Vxlan Nd discard count            : 2
     Tx Vxlan Arp discard count           : 0
     Tx Vxlan Nd discard count            : 0
     Rx Vxlan Arp Request count           : 2
     Tx Vxlan Arp Request count           : 0
     Rx Vxlan Arp Reply count             : 5
     Tx Vxlan Arp Reply count             : 1
     Rx Vxlan Neighbor Solicitation count : 6
     Tx Vxlan Neighbor Solicitation count : 0
```

```
 Rx Vxlan Neighbor Advertisement count: 0
 Tx Vxlan Neighbor Advertisement count: 4
 Rx Vxlan Gratuitous ARP count        : 0
 Tx Vxlan Gratuitous ARP count        : 0
 Rx Vxlan Gratuitous Neighbor Advertisement count: 3
 Tx Vxlan Gratuitous Neighbor Advertisement count: 0
#
```

Table 1-5 explains the fields in the output.

Table 1-5: access port counters

Field	Description
RX: packets	Number of packets received.
RX: bytes	Number of bytes received.
TX: packets	Number of packets transmitted.
TX: bytes	Number of bytes transmitted.
Rx Vxlan Nd discard count	Number of discarded ND that is received from neighbor.
Tx Vxlan Arp discard count	Number of discarded Arp that is transmitted to peer.
Rx Vxlan Nd discard count	Number of discarded ND that is transmitted to peer.
Rx Vxlan Arp Request count	Number of request ARP that is received from neighbor.
Tx Vxlan Arp Request count	Number of request ARP that is transmitted to peer.
Rx Vxlan Arp Reply count	Number of replied ARP that is received from neighbor.
Tx Vxlan Arp Reply count	Number of replied ARP which is transmitted to peer.
Rx Vxlan Neighbor Solicitation count	Number of request ND that is received from neighbor.
Tx Vxlan Neighbor Solicitation count	Number of replied ND that is transmitted to peer.
Rx Vxlan Neighbor Advertisement count	Number of Neighbor Advertisement that is received from neighbor.

Table 1-5: access port counters (Continued)

Field	Description
Tx Vxlan Neighbor Advertisement count	Number of Neighbor Advertisement that is transmitted to peer
Rx Vxlan Gratuitous ARP count	Number of Gratuitous ARP that is received from neighbor.
Tx Vxlan Gratuitous ARP count	Number of Gratuitous ARP which is transmitted to peer.
Rx Vxlan Gratuitous Neighbor Advertisement count	Number of Gratuitous Neighbor Advertisement that is received from neighbor.
Tx Vxlan Neighbor Gratuitous Advertisement count	Number of Gratuitous Neighbor Advertisement which is transmitted to peer.

show nvo vxlan counters network-port

Use this command to display the receive and transmit counters of a network port including ARP, ND and GARP counters in the same command.

Command Syntax

```
show nvo vxlan counters network-port dst A.B.C.D
```

Parameters

A.B.C.D	Tunnel IPv4 address

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3 and added control packets (ARP, ND, and GARP) counters in OcNOS version 1.3.5.

Example

```
VTEP1#sh nvo vxlan counters network-port dst 100.2.2.2

Data packets:
   *If ARP/ND cache is enabled, TX doesn't count ARP/ND replies
 from  ARP/ND cache and ARP/ND forwarded after uplifting to
 the control plane.
   RX:  packets  : 814327
        bytes    : 92833544
   TX:  packets  : 1508023
        bytes    : 171914622

Control Packets:
 *ARP/ND uplifted and sent/replied from control plane:
  TX VXLAN Arp discard count           : 0
  TX VXLAN ND discard count            : 0
  Tx Vxlan Arp Request count           : 0
  Tx Vxlan Arp Reply count             : 0
  Tx Vxlan Neighbor Solicitation count : 0
  Tx Vxlan Neighbor Advertisement count: 0
  Rx Vxlan Gratuitous ARP count: 0
  Tx Vxlan Gratuitous ARP count: 0
  Rx Vxlan Gratuitous Neighbor Advertisement count: 0
  Tx Vxlan Gratuitous Neighbor Advertisement count: 0
```

Table 1-6 explains each network entry fields.

Table 1-6: show nvo vxlan counters network-port output fields

Field	Description
RX: packets	Number of hello packets received from neighbor.
RX: bytes	Number of hello packets received from neighbor in bytes received.
TX: packets	Number of hello packets transmitted to neighbor.
TX: bytes	Number of hello packets transmitted to neighbor in bytes transmitted.
Rx Vxlan Nd discard count	Number of discarded ND that is received from neighbor.
Tx Vxlan Arp discard count	Number of discarded Arp that is transmitted to peer.
Rx Vxlan Nd discard count	Number of discarded ND that is transmitted to peer.
Rx Vxlan Arp Request count	Number of request ARP that is received from neighbor.
Tx Vxlan Arp Request count	Number of request ARP that is transmitted to peer.
Rx Vxlan Arp Reply count	Number of replied ARP that is received from neighbor.
Tx Vxlan Arp Reply count	Number of replied ARP which is transmitted to peer.
Rx Vxlan Neighbor Solicitation count	Number of request ND that is received from neighbor.
Tx Vxlan Neighbor Solicitation count	Number of replied ND that is transmitted to peer.
Rx Vxlan Neighbor Advertisement count	Number of Neighbor Advertisement that is received from neighbor.
Tx Vxlan Neighbor Advertisement count	Number of Neighbor Advertisement that is transmitted to peer.
Rx Vxlan Gratuitous ARP count	Number of Gratuitous ARP that is received from neighbor.
Tx Vxlan Gratuitous ARP count	Number of Gratuitous ARP which is transmitted to peer.
Rx Vxlan Gratuitous Neighbor Advertisement count	Number of Gratuitous Neighbor Advertisement that is received from neighbor.
Tx Vxlan Neighbor Gratuitous Advertisement count	Number of Gratuitous Neighbor Advertisement which is transmitted to peer.

show nvo vxlan mac-table

Use this command to display the host MAC address table. Use the `hardware` option to see the age out time for dynamically learned macs.

Command Syntax

```
show nvo vxlan mac-table (vnid <1-16777215>|) (summary | hardware |)
```

Parameters

`<1-16777215>`	VXLAN network identifier.
`summary`	Count the MAC addresses.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Remote Static and Dynamic keyword got changed to Static and Dynamic Remote in MAC table in OcNOS version 1.3.6.

Example

```
#show nvo vxlan mac-table
================================================================================
================================================================
                      VXLAN MAC Entries
================================================================================
================================================================
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/
ESI                    Type                 Status              AccessPortDesc
--------------------------------------------------------------------------------
----------------------------------------------------------------

10        ce21/1     2      ----         0000.0000.1111
1.1.1.1                      Static Local         -------
partner-port
10        ----       ----   ----         0000.0000.2222
3.3.3.3                      Static Remote        -------             ------
-
20        ce21/1     3      ----         0000.0000.cccc
1.1.1.1                      Static Local         Discard             ------
-
20        ----       ----   ----         0000.0000.dddd
3.3.3.3                      Static Remote        -------             ------
-

Total number of entries are : 4

#show nvo vxlan mac-table hardware
================================================================================
================================================================
                        VXLAN MAC Entries
================================================================================
================================================================
```

```
VNID       Interface VlanId Inner-VlanId Mac-Addr       VTEP-Ip/
ESI                   Type                 Status            Time-out
AccessPortDesc
_____________________________________________________________________________
_____________________________________________________________________________

10         ce21/1    2      ---          0000.0000.1111
1.1.1.1                      Static Local         -----             ---
partner-port
10         ---       ---    ---          0000.0000.2222
3.3.3.3                      Static Remote               -----             -
--
10         ---       ---    ---          0000.0000.aa11
3.3.3.3                      Static Remote               -----             -
--
10         ce21/1    2      ---          0000.0000.bb11
1.1.1.1                      Dynamic Local        -----             300
partner-port
10         ce21/1    2      ---          0000.0000.bb12
1.1.1.1                      Dynamic Local        -----             277
partner-port
20         ce21/1    ---    ---          0000.0000.cccc
1.1.1.1                      Static Local         Discard           ---
20         ---       ---    ---          0000.0000.dddd
3.3.3.3                      Static Remote               -----             -
--

Total number of entries are 7
```

Table 1-7 explains the fields in the output.

Table 1-7: MAC table fields

Field	Description
VNID	VXLAN network identifier
Interface	Interface name
VlanId	VLAN identifier
Mac-Addr	MAC address
VTEP-Ip	VTEP identifier/Ethernet Segment Identifier
Type	How a host learns a MAC/IP pair: Remote: Statically configured for remote; used only for static VxLAN, not with EVPN Static Local: Configured on local VTEP Dynamic Local: Learned by data plane source learning
Status	Max Move conflict: When a MAC has moved too many times (5 or more times in 180 seconds). This is according to the procedures defined in RFC 7432, Section 15.1. Discard: If a MAC hold time is configured, then if the VxLAN access port goes down (admin or operational), the MAC is moved to the discard state for the period of the hold time. The MAC is also moved to the discard state if the VNID is unmapped from the port. In dynamically learned cases, the MAC is also moved to discard when learning is disabled.
Time-out	Age timeout for dynamically learned MACs.

Table 1-7: MAC table fields (Continued)

Field	Description
AccessPortDesc	Access port description.
Total number of entries	The total number of entries listed.

show nvo vxlan nd-cache

Use this command to display the Neighbor Discovery cache.

Command Syntax

```
show nvo vxlan nd-cache (vnid <1-16777215>|summary|)
```

Parameters

`<1-16777215>`	VXLAN network identifier.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3, modified to include new fields – Age-Out and Retries-Left in OcNOS version 1.3.5.

Remote Static and Dynamic keyword got changed to Static and Dynamic Remote in ND table in OcNOS version 1.3.6.

Example

```
#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
==========================
ND Timeout :  300 sec   Random-Jitter-Max : 640

VNID     Ip-Addr          Mac-Addr       Type          Age-Out    Retries-Left
_____________________________________________________________________________
10       2001::1       0000.0000.1111   Dynamic Local   274          2
10       2001::2       0000.0000.2222   Dynamic Remote  ----
Total number of entries are 2
```

Table 1-8 explains the output fields.

Table 1-8: ND cache fields

Field	Description
VNID	VXLAN network identifier
Ip-Addr	IP address
Mac-Addr	MAC address
Type	How a host learns a MAC/IP pair: Static Remote: Statically configured for remote; used only for static VxLAN, not with EVPN Static Local: Configured on local VTEP Dynamic: Learned by data plane source learning Evpn: Learned by EVPN Type 2 (MAC/IP) routes

Table 1-8: ND cache fields

Field	Description
Age-Out	ND entry expire Time
Total number of entries	The total number of entries listed.

show nvo vxlan static host state

Use this command to display the state of the host which is configured statically.

Command Syntax

```
show nvo vxlan static host state
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#sh nvo vxlan static host state
VNID Ifname    Vlan      Ip-Addr        Mac-Addr         Status
_____________________________________________________________________
100      xe1          2      12.12.12.1       0000.0000.aaaa      Inactive
100      xe2          2      1201::1          0000.0000.bbbb      Conflict
100      xe3          2      120.120.120.1    0000.0000.bb11      Active
```

Table 1-9 explains the output fields.

Table 1-9: Static host fields

Field	Description
VNID	VXLAN network identifier
Ifname	Interface name
Vlan	VLAN name
Ip-Addr	IP address

Table 1-9: Static host fields

Field	Description
Mac-Addr	MAC address
Status	Status of the MAC/IP on the host: Conflict: When a MAC/IP was configured, the conflict was not known as the VNID was not mapped to the access port. After the VNID is mapped, if the same MAC/IP is present statically on some other port on the same VNID, then it is in conflict state. Learnt Conflict: When a MAC/IP was configured, the conflict was not known. However, it is now in conflict because the same MAC/IP is configured on an access port on VTEP1 and on an access port on VTEP2. Because the BGP session/tunnel was not up, the MAC/IP was not known to the other VTEP and the configuration was allowed. When the BGP session/tunnel comes up and it finds such a conflicted route, it marks the state as Learnt Conflict. Inactive: Configured but not operating, such as when the port is not mapped to any VNID. The port is down and the ARP/ND cache is disabled. Active: Operating host MAC/IP.

show nvo vxlan tunnel

Use this command to view the (remote VTEP router) source, destination, and status of the VxLAN tunnel entries.

Command Syntax

```
show nvo vxlan tunnel
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output of the `show nvo vxlan tunnel` command.

```
#show nvo vxlan tunnel
VXLAN Network tunnel Entries
 Source Destination      Status Up/Down Update
=====================================================
1.1.1.1 2.2.2.2          Installed 00:00:20 00:00:20
Total number of entries are 1
#
```

Table 1-10 explains the output fields.

Table 1-10: VxLAN tunnel fields

Field	Description
Source	Tunnel source IP address.
Destination	Tunnel destination IP address.
Status	Installed: Tunnel Installed in the hardware and operating. Resolved: Tunnel destination IP is reachable, but VxLAN tunnel not installed in hardware. Therefore, not operating. Unresolved: Tunnel destination IP not reachable because L3 route is down.
Up/Down	When the tunnel came up or went down
Update	When the tunnel was last updated
Total number of entries	The total number of entries listed.

show running-config nvo vxlan

Use this command to display the current running configuration of VxLANs.

Command Syntax

```
show running-config nvo vxlan
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in a version before OcNOS version 1.3 and added few CLIs for IRB and ARP-ND refresh timer for OcNOS version 1.3.5.

Example

```
#show running-config nvo vxlan
!
     nvo vxlan multihoming enable
!
     nvo vxlan enable
     !
     nvo vxlan vtep-ip-global 1.1.1.1
     !
     nvo vxlan max-cache-disable 10
     !
     nvo vxlan tunnel qos-map-mode cos-dscp ingress QUE_DSCP
     !
     nvo vxlan tunnel qos-map-mode cos-dscp egress DSCP_QUE
     !
     nvo vxlan id 100 ingress-replication inner-vid-disabled
      vxlan host-reachability-protocol evpn-bgp vrf1
      mac-holdtime 20
     !
     nvo vxlan access-if port-vlan xe1 2
      map vnid 100
access-if-description member-port with xe1 as vlan 2
dynamic-learning disable
arp-nd flood-suppress
arp-cache disable
nd-cache disable
shutdown
mac 0000.0000.1111
mac 0000.0000.aaaa  ip 12.12.12.1
mac 0000.0000.bbbb  ipv6 1201::1
map qos-profile cos-to-queue COS_QUE
map qos-profile queue-color-to-cos QUE_COS
     !
     nvo vxlan access-if port-vlan po1 6 inner-vlan 5
      encapsulation 0x9100
```

```
 no shutdown
 map vnid 100
!
```

show evpn esi

Use this command to display host ESI's (multihomed segments) and its connected VTEPs end point's IP address list.

Command Syntax

```
show evpn esi
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

```
Leaf1#show evpn esi
ESI entries and their endpoints
ESI                             PE-list
============================================================
00:0e:00:56:02:16:00:00:00:00  172.20.254.9,  172.20.254.34
00:0e:00:56:02:16:16:00:00:00  172.20.254.10, 172.20.254.38
Total number of entries are 2
Leaf1#
```

show evpn multi-homing all

Use this command to display the multi-homed VTEP details.

Command Syntax

```
show evpn multi-homing (all |)
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show evpn multi-homing all
ESI                             Access-IF    PE-IP-ADDRESS
============================================================
00:00:11:22:33:44:55:66:77:88  ce21/1       1.1.1.1
00:00:11:22:33:44:55:66:77:88  ----         2.2.2.2
Total number of entries are 2
```

Table 1-11 explains the output fields.

Table 1-11: show evpn multi-homing all output details

Field	Description
ESI	An Ethernet segment has an unique nonzero identifier, called the Ethernet segment identifier (ESI). The ESI is encoded as a 10-octet integer that identifies this segment. When manually configuring an ESI value, the most significant octet, known as the type byte, must be 00. When a single-homed CE device is attached to an Ethernet segment, the entire ESI value is zero.
Access-IF	Map the access port ce21/1 for evpn.
PE-IP-ADDRESS	Address of the provider edge router in the interface.

show evpn multihoming-status

Use this command to display the status of multihoming on a VTEP.

Command Syntax

```
show evpn multihoming-status
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show evpn multihoming-status
Multihoming is ACTIVE in Hardware
```

show nvo vxlan route-count

Use this command to display the vxlan active route (MAC-IP.MAC-IPv6 and MAC-only) count information.

Command Syntax

```
show nvo vxlan route-count (|vnid <1-16777215>)
```

Parameters

`<1-16777215>`	Range supported for VNID.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3 and modified to include new fields – GW-IPv4, GW_IPv6, Prefix IPv4 and Prefix IPv6 in OcNOS version 1.3.5.

Example

```
#show nvo vxlan route-count
VXLAN Active route count information
====================================
Max route count   : 32768
Active route count: 7

Note: Prefix count refers only local routes
      It does not include evpn prefix routes

-------------------------------------------------------------------------------
----------
 VNID     Total    MACONLY  MACIPv4  MACIPv6  GW         GW         Prefix
Prefix
                                               MACIPv4    MACIPv6    IPv4
IPv6
-------------------------------------------------------------------------------
----------
 2000     3        1        1        1        0          0          0
0
 1000     0        0        0        0        0          0          0
0
 10002    4        0        0        0        1          1          1
1

Total number of entries are 3
```

Table 1-12 explains the output fields.

Table 1-12: show nvo vxlan route-count output details

Field	Description
Max route count	Maximum number of route count in vxlan.
Active route count	Number of active route count in the interface.

Table 1-12: show nvo vxlan route-count output details

Field	Description
VNID	VNID is used to identify Layer 2 segments and to maintain Layer 2 isolation between the segments.
Total	Total number of entries for the interface.
MACONLY	The MAC-only route for the local interface appears in the VXLAN instance route table.
MACIPv4	IPv4 media access control (MAC) address for a default virtual gateway.
MACIPv6	IPv6 media access control (MAC) address for a default virtual gateway.

show nvo vxlan vni-name

Use this command to display the vxlan results bashed on vni-name.

Command Syntax

```
show nvo vxlan vni-name (WORD)
```

Parameters

WORD	VNI name of max size 10 character and should not be only numeric.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show nvo vxlan vni-name SITEA-PRO
VXLAN Information
=================
Codes: NW - Network Port
AC - Access Port
(u) - Untagged
VNID Vni-name Type Interface ESI Vlan DF-Status Src-addr Dst-addr
_____________________________________________________________________________________
_____________________________
1 ---- NW ---- ---------- ---- ------ 10.0.1.1 10.0.6.8
1 ---- NW ---- ---------- ---- ------ 10.0.1.1 10.0.6.9
1 ---- NW ---- ---------- ---- ------ 10.0.1.1 10.0.3.1
1 ---- NW ---- ---------- ---- ------ 10.0.1.1 10.0.1.2
1 ---- NW ---- ---------- ---- ------ 10.0.1.1 10.0.5.1
1 ---- NW ---- ---------- ---- ------ 10.0.1.1 10.0.2.2
1 ---- NW ---- ---------- ---- ------ 10.0.1.1 10.0.2.1
1 SITEA-PRO AC xe7 — Single Homed port — 2 ------ ---- ----
1 SITEA-PRO AC xe1 — Single Homed port — 1010 ------ ---- ----
1 SITEA-PRO AC xe1 — Single Homed port — 100 ------ ---- ----
1 SITEA-PRO AC xe1 — Single Homed port — 2020 ------ ---- ----
1 SITEA-PRO AC po1 — Single Homed port — 100 ------ ---- ----
1 SITEA-PRO AC po1 — Single Homed port — 2 ------ ---- ----
1 SITEA-PRO AC po1 — Single Homed port — 200 ------ ---- ----
1 SITEA-PRO AC xe8 — Single Homed port — ---- ------ ---- ----
1 SITEA-PRO AC po2 — Single Homed port — ---- ------ ---- ----
Total number of entries are 16
```

show hsl evpn multihoming esi

Use this command to display the hsl status of evpn multihome esi.

Command Syntax

```
show hsl evpn multihoming esi (mac-list | vni)
```

Parameters

`mac-list`	Mac address list.
`vni`	VNID.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 6.6.0.

Example

```
##show hsl evpn multihoming esi
 L-means ESI Local, R-means ESI Remote
-------------------------------------------------------------------------------
-------------------------------------------------
ESI                               ESI-Status If       Ifstatus   VpLag
Tunnels            T-port      Tun-egrobj-id
-------------------------------------------------------------------------------
-------------------------------------------------
00:00:00:aa:aa:bb:bb:00:00:00(L) Up         po10     Up         --
--                 0x00000000   0x00000000
00:00:00:aa:aa:bb:bb:00:00:00(R) --         --       --         --
101.1.1.1          0x80000004   0x000186AA

#show hsl evpn multihoming esi
 L-means ESI Local, R-means ESI Remote
-------------------------------------------------------------------------------
-------------------------------------------------
ESI                               ESI-Status If       Ifstatus   VpLag
Tunnels            T-port      Tun-egrobj-id
-------------------------------------------------------------------------------
-------------------------------------------------
00:00:00:aa:aa:bb:bb:00:00:00(R) --         --       --         2113
101.1.1.1          0x80000004   0x000186AC
                                                                2113
102.1.1.1          0x80000002   0x000186A9
00:00:00:cc:cc:dd:dd:00:00:00(L) Down       po10     Down       --
--                 --           --

#show hsl evpn multihoming esi vni
 L-means ESI Local, R-means ESI Remote
-------------------------------------------------------------------------------
-------------------------------------------------
VNID    ESI                           ESI-Status If       Ifstatus   VpLag
Tunnels            T-port      Tun-egrobj-id
-------------------------------------------------------------------------------
-------------------------------------------------
```

```
102     00:00:00:aa:aa:bb:bb:00:00:00(L) Up        po10    Up        --
--                0x00000000   0x00000000
102     00:00:00:aa:aa:bb:bb:00:00:00(R) --        --      --        --
102.1.1.1         0x80000006   0x000186AE
103     00:00:00:aa:aa:bb:bb:00:00:00(L) Up        po10    Up        --
--                0x00000000   0x00000000
103     00:00:00:aa:aa:bb:bb:00:00:00(R) --        --      --        --
102.1.1.1         0x80000006   0x000186AE
104     00:00:00:aa:aa:bb:bb:00:00:00(L) Up        po10    Up        --
--                0x00000000   0x00000000
104     00:00:00:aa:aa:bb:bb:00:00:00(R) --        --      --        --
102.1.1.1         0x80000006   0x000186AE

#show hsl evpn multihoming esi vni
 L-means ESI Local, R-means ESI Remote
-------------------------------------------------------------------------------
-------------------------------------------------
VNID    ESI                           ESI-Status If       Ifstatus   VpLag
Tunnels           T-port      Tun-egobj-id
-------------------------------------------------------------------------------
-------------------------------------------------
102     00:00:00:aa:aa:bb:bb:00:00:00(R) --        --      --         2113
101.1.1.1         0x80000005   0x000186AC
                                                                     2113
102.1.1.1         0x80000002   0x000186A9
103     00:00:00:aa:aa:bb:bb:00:00:00(R) --        --      --         2113
101.1.1.1         0x80000005   0x000186AC
                                                                     2113
102.1.1.1         0x80000002   0x000186A9
104     00:00:00:aa:aa:bb:bb:00:00:00(R) --        --      --         2113
101.1.1.1         0x80000005   0x000186AC
                                                                     2113
102.1.1.1         0x80000002   0x000186A9
150     00:00:00:cc:cc:dd:dd:00:00:00(L) Down      po10    Down       --
--                --           --

#show hsl evpn multihoming esi mac-list
-------------------------------------------------------------------------------
---------------
MAC                  VNI        VPNID    ESI                             vplag-
id
-------------------------------------------------------------------------------
----------------
aa:aa:bb:bb:00:00:00 102        28674    00:00:00:aa:aa:bb:bb:00:00:00  2113
aa:aa:bb:bb:00:00:01 102        28674    00:00:00:aa:aa:bb:bb:00:00:00  2113
aa:aa:bb:bb:00:00:02 102        28674    00:00:00:aa:aa:bb:bb:00:00:00  2113
```

shutdown

Use this command to administratively shut down an NVO access interface.

Use the `no` form of this command to start an NVO access interface.

Command Syntax

```
shutdown
no shutdown
```

Parameters

None

Default

The NVO access interface is running by default.

Command Mode

NVO access interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#shutdown
(config-nvo-acc-if)#exit
```

vxlan host-reachability-protocol evpn-bgp

Use this command to set the host reachable protocol to Ethernet-VPN over BGP. This defines BGP as the mechanism for host reachability advertisement.

Use use the no form of this command to remove Ethernet-VPN as the host reachable protocol.

Command Syntax

```
vxlan host-reachability-protocol evpn-bgp NAME
no vxlan host-reachability-protocol evpn-bgp
```

Parameters

NAME	Name of the VRF to carry VNID routes.

Default

No default value is specified for vxlan host-reachability-protocol command.

Command Mode

NVO mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#nvo vxlan id 3
(config-nvo)#vxlan host-reachability-protocol evpn-bgp Blue
```

vlan-xlate-1 large

Use this command to increase the size of the VLAN_XLATE_1_DOUBLE table to 20k. Internally, the BCM maintains VLAN_XLATE_1_DOUBLE table for access interfaces, L2vnid,and L3vnid. The default size of the table is 16k. After configuring the CLI, the table size is increased to 20k.

Use the `no` form of this command to change the VLAN_XLATE_1_DOUBLE table size to default.

Note: Reboot the system after configuring the CLI.

Command Syntax

```
vlan-xlate-1 large
no vlan-xlate-1 large
```

Parameters

None

Default

None

Command Mode

Configuration Mode

Applicability

This command was introduced before OcNOS version 6.5.3.

Examples

```
(config)#vlan-xlate-1 large
OcNOS(config)#commit
%% System Reboot required, please save the config and reboot the board.
OcNOS(config)#end
OcNOS#
```

CHAPTER 2 VXLAN - IRB Commands

This chapter describes the VXLAN - IRB commands: These commands are applicable for TR3 and Maverick platform only.

evpn irb

Use this command to configure default gateway behavior on a VTEP for particular VNID.

Use the no form this command to disable default gateway behavior on a VTEP for the particular VNID.

Note: Map an IRB interface to an L2 VNID. This IRB interface can have multiple IP address as configured in IRB IP address CLI and can serve all subnets attached to the L2 VNID.

Command Syntax

```
evpn  <NAME>
no evpn  <NAME>
```

Parameters

NAME	IRB interface name

Command Mode

NVO Mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#config mode
(config)# nvo vxlan id 2000 ingress-replication inner-vid-disabled
(config-nvo)# evpn irb1
(config-nvo)# no evpn irb1
```

evpn irb-forwarding anycast-gateway-mac

Use this command to configure common anycast mac-address for all the IRB interfaces

Use the `no` form of this command to remove the global MAC address on all the the IRB interfaces.

Command Syntax

```
evpn irb-forwarding anycast-gateway-mac XXXX.XXXX.XXXX
no evpn irb-forwarding anycast-gateway-mac
```

Parameters

`XX-XX-XX-XX-XX-XX`	Source MAC address (Option 1)
`XX:XX:XX:XX:XX:XX`	Source MAC address (Option 2)
`XXXX.XXXX.XXXX`	Source MAC address (Option 3)

Command Mode

Configuration Mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#configure terminal
(config)#evpn irb-forwarding anycast-gateway-mac 0000.0000.1313
Or
(config)#evpn irb-forwarding anycast-gateway-mac 00:00:00:00:13:13
Or
(config)#evpn irb-forwarding anycast-gateway-mac 00-00-00-00-13-13
(config)# no evpn irb-forwarding anycast-gateway-mac
```

evpn irb-if-forwarding anycast-gateway-mac

Use this command to enable an IRB interface to use the global anycast IRB mac-address.

Use the `no` form of this command to un-configure anycast MAC at IRB interface.

Command Syntax

```
evpn irb-if-forwarding anycast-gateway-mac
no evpn irb-if-forwarding anycast-gateway-mac
```

Parameters

None

Command Mode

IRB_IF_Mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#configure teminal
(config)# interface irb 1
(config-irb-if)# ip vrf forwarding vrfip
(config-irb-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-irb-if)#no evpn irb-if-forwarding anycast-gateway-mac
```

interface irb

Use this command to configure logical IRB interface.

Use the `no` form of this command to un-configure logical IRB interface.

Command Syntax

```
interface irb <1-80000>
no interface irb <1-80000>
```

Parameters

`<1-80000>`	IRB interface number

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#configure teminal
(config)#interface irb1
(config)#no interface irb1
```

l3vni

Use this command to configure L3 Virtual Network Identifier for an ip vrf

Use the no form of this command to remove L3 Virtual Network Identifier

This identifies a tenant, with this one tenant can have L3VNI as its identifier and he can have multiple L2 networks identified with L2VNI's.

Note: L3 VNID cannot be same as L2 VNID.

Command Syntax

```
l3vni <L3 VNID>
no l3vni <L3 VNID>
```

Parameters

<1-16777215>	L3 VNID. Cannot be same as L2 VNID

Command Mode

Configure VRF mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#configure teminal
(config)#ip vrf vrfip
(config-vrf)#l3vni 10002
(config-vrf)#no l3vni 10002
```

nvo vxlan irb

Use this command to enable IRB functionality.

Use the `no` form of this command to disable IRB functionality.

Note: Remove the existing L2 VNID configuration to enable IRB.

Command Syntax

```
nvo vxlan irb
no nvo vxlan irb
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#configure terminal
(config)#nvo vxlan irb
(config)#no nvo vxlan irb
```

show interface irb

Use this command to display the current running configuration of IRB interface.

Command Syntax

```
show interface irb <1-80000>
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
rtr1#show interface irb1
Interface irb1
  Hardware is IRB  Current HW addr: 0000.0000.ff10
  Physical:(Not Applicable)  Logical:0000.0000.ff10
  Port Mode is Router
  Interface index: 700001
  Metric 0 mtu 1500
  Debounce timer: disable
  ARP ageing timeout 1500
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Associated with vrf1
  Label switching is disabled
  Administrative Group(s): None
  DHCP client is disabled.
  Last Flapped: Never
  Statistics last cleared: 2019 Mar 14 17:57:06 (00:21:31 ago)
  inet 80.80.80.1/24 broadcast 80.80.80.255
  inet6 8001::1/64
  inet6 fe80::200:ff:fe00:ff10/64
  RX
    unicast packets 0 multicast packets 0 broadcast packets 0
    input packets 0 bytes 0
    jumbo packets 0
    undersize 0 oversize 0 CRC 0 fragments 0 jabbers 0
    input error 0
    input with dribble 0 input discard 0
    Rx pause 0
  TX
    unicast packets 0 multicast packets 0 broadcast packets 0
    output packets 0 bytes 0
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
    Tx pause 0
rtr1#
```

show evpn l3vni-map

Use this command to display the L3 VNI, L2 VNI and IRB interface mapping.

Command Syntax

```
show evpn l3vni-map
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
rtr3#show evpn l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 45001        10           irb1
 45001        20           irb2

rtr3#
```

show evpn irb-status

Use this command to display the status of the IRB on a VTEP.

Command Syntax

```
show evpn irb-status
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 5.1.

Example

```
#show evpn irb-status
IRB is ACTIVE in Hardware
```

show running-config interface irb

Use this command to display the current running configuration of IRB interface.

Command Syntax

```
show running-config interface irb<1-4094>
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#show running-config interface irb1
!
      interface irb1
      ip vrf forwarding vrfip
      ip address 144.144.144.1/24
      ipv6 address 1201::1/48
      evpn irb-if-forwarding anycast-gateway-mac
      mac-address 0000.0000.1234
      qos map-profile dscp-to-queue DSCP_QUE
      qos map-profile queue-color-to-dscp QUE_DSCP
      shutdown
!
```

CHAPTER 3 VXLAN Quality of Service Commands

This chapter describes the VXLAN commands for QoS (Quality of Service):

clear nvo vxlan tunnels

Use this command to clear the nvo vxlan tunnels to re-establish the tunnel after mapping/un-mapping the QoS profile to vxlan tunnel.

Command Syntax

```
clear nvo vxlan tunnels (|dst-ip A.B.C.D)
```

Parameters

`dst-ip`	VXLAN tunnel destination
`A.B.C.D`	destination IPv4 address

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
#clear nvo vxlan tunnels dst-ip 2.2.2.2
```

Note: `clear nvo vxlan tunnels` This command will clear all the VXLAN tunnels destination.

`clear nvo vxlan tunnels dst-ip A.B.C.D` - This command to clear individual tunnel destination(i.e A.B.C.D).

cos queue

Use this command to configure user defined mapping for cos and queue.

Use the `no` form of this command to remove the mapping.

Command Syntax

```
cos <0-7> queue <0-7>
no cos <0-7>
```

Parameters

`<0-7>`	COS and Queue ranger

Default

Default cos and queue value is one-one default mapping if it is not configured.

Command Mode

QoS config mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#qos profile cos-to-queue ac_port_ingress
(config-ingress-cos-map)#cos 1 queue 7
(config-ingress-cos-map)#no cos  1
```

dscp queue

Use this command to configure user defined mapping for DSCP to queue. This will be mapped with nvo VXLAN tunnel of remote VTEP.

Use the `no` form of this command to delete the mapping.

Command Syntax

```
dscp <0-63> queue <0-7> (color (green|yellow|red)|) (dscp <0-63>|)
no dscp <0-63>
```

Parameters

`<0-63>`	DSCP
`<0-7>`	Queue number
`color`	Color to map
`green`	Set mapping for green packets
`red`	Set mapping for red packets
`yellow`	Set mapping for yellow packets
`<0-63>`	out DSCP value

Default

Default queue and cos value is one-one default mapping if it is not configured.

Command Mode

Ingress-dscp-map mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#qos profile dscp-to-queue nw_profile
(config-ingress-dscp-map)#dscp 50 queue 1
(config-ingress-dscp-map)#no dscp 50
```

map qos-profile

Use this command to map (attach) the qos profile to an IRB interface.

Use the `no` form of this command to remove a profile.

Use the following qos profile type for mapping from/to on the IRB interface.

- dscp-to-queue for ingress traffic and
- queue-color-to-dscp profile for egress traffic

Removing the map qos-profile, applies the default profile to the IRB interface.

Note: Default profile is applied to all the IRB interface only when qos is enabled.

Command Syntax

```
map qos-profile (dscp-to-queue | queue-color-to-dscp) <NAME>
no map qos-profile (dscp-to-queue | queue-color-to-dscp) <NAME>
```

Parameters

`NAME`	Profile name

Default

By default, the default dscp-to-queue and queue-color-to-dscp profile is attached to all IRB interface.

Changing the value in the default profile, will impact both the IRB interfaces and L3 interfaces.

Command Mode

IRB_IF_Mode

Default

Default mapping between queue and DSCP value is one-one.

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
#configure terminal
(config)#interface irb 1
(config-irb-if)# map qos-profile queue-color-to-dscp QUE_DSCP
(config-irb-if)# no map qos-profile queue-color-to-dscp QUE_DSCP

(config)#interface irb 1
(config-irb-if)# map qos-profile dscp-to-queue DSCP_QUE
(config-irb-if)# no map qos-profile dscp-to-queue DSCP_QUE
```

map qos-profile cos-to-queue

Use this command to map the cos-to-queue profile to vxlan access port on the local VTEP.

Use the `no` form of the command to remove the mapping.

Command Syntax

```
map qos-profile cos-to-queue NAME
no map qos-profile cos-to-queue NAME
```

Parameters

`NAME`	Profile name

Default

None

Command Mode

NVO access interface mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
(config)#nvo vxlan access-if port-vlan xe1 10
(config-nvo-acc-if)#map qos-profile cos-to-queue ac_port_ingress
(config-nvo-acc-if)#no map qos-profile cos-to-queue ac_port_ingress
```

map qos-profile queue-color-to-cos

Use this command to map the queue-color-to-cos profile to vxlan access port on the remote VTEP.

Use the `no` form of the command to remove the mapping.

Command Syntax

```
map qos-profile queue-color-to-cos NAME
no map qos-profile queue-color-to-cos NAME
```

Parameters

`NAME`	Profile name

Default

None

Command Mode

NVO access interface mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
(config)#nvo vxlan access-if port-vlan xe2 10
(config-nvo-acc-if)#map qos-profile queue-color-to-cos ac_profile
(config-nvo-acc-if)#no map qos-profile queue-color-to-cos ac_profile
```

nvo vxlan disable-arp-storm-control-for-cpu

Use this command to uplift the ARP/ND packet to CPU if packet is marked for DROP by storm control.

Use the `no` form of this command not to uplift the ARP/ND packet to CPU if marked for drop by storm-control. This is the default behavior.

Command Syntax

```
nvo vxlan disable-arp-storm-control-for-cpu
no nvo vxlan disable-arp-storm-control-for-cpu
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Example

```
(config)#nvo vxlan disable-arp-storm-control-for-cpu
```

nvo vxlan tunnel qos-map-mode cos-dscp

Use this command to map QoS profile for network side to nvo vxlan tunnel. For outgoing/incoming traffic, you need to provide the direction with the keyword `egress`/`ingress`.

Use the `no` form of this command to delete the mapping.

You must give the clear nvo vxlan tunnels command to do the network port setting for QoS profile mapped.

Command Syntax

```
nvo vxlan tunnel qos-map-mode cos-dscp (ingress|egress) NAME
no nvo vxlan tunnel qos-map-mode cos-dscp (ingress|egress)
```

Parameters

`NAME`	Profile name
`ingress`	Ingress direction
`egress`	Egress direction

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#nvo vxlan tunnel qos-map-mode cos-dscp egress nw_profile
(config)#no nvo vxlan tunnel qos-map-mode cos-dscp egress
(config)#nvo vxlan tunnel qos-map-mode cos-dscp ingress nw_profile
(config)#no nvo vxlan tunnel qos-map-mode cos-dscp ingress
```

qos profile cos-to-queue

Use this command to configure cos-to-queue profile. This profile has to be mapped to VXLAN access port on the local VTEP.

Use the `no` form of this command to delete the qos profile.

Command Syntax

```
qos profile cos-to-queue (NAME|default)
no qos profile cos-to-queue NAME
```

Parameters

`NAME`	QoS profile name for cos-to-queue
`default`	Default name

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#qos profile cos-to-queue ac_port_ingress
(config)#no qos profile cos-to-queue ac_port_ingress
```

qos profile dscp-to-queue

Use this command to configure QoS profile for DSCP to Queue mapping. This profile will be mapped to nvo vxlan tunnel of remote VTEP. The created profile will support remarking of the data packets.

Use the `no` form of this command to delete the QoS profile

Command Syntax

```
qos profile dscp-to-queue (NAME|default}
no qos profile dscp-to-queue NAME
```

Parameters

`NAME`	Profile name
`default`	Default name

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#qos profile dscp-to-queue nw_profile
(config-ingress-dscp-map)#exit
(config)#no qos profile dscp-to-queue nw_profile
(config)#
```

qos profile queue-color-to-cos

Use this command to configure queue-color-to-cos profile. This profile has to be mapped to VXLAN access port on remote VTEP.

Use the `no` form of this command to delete the qos profile.

Command Syntax

```
qos profile queue-color-to-cos (NAME|default)
no qos profile queue-color-to-cos NAME
```

Parameters

`NAME`	Profile name
`default`	Default name

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#qos profile queue-color-to-cos ac_profile
(config)#no qos profile queue-color-to-cos ac_profile
```

qos profile queue-color-to-dscp

Use this command to create a QoS profile queue-color-to-dscp. This profile will be mapped to nvo vxlan tunnel of local VTEP. The created profile supports remarking of the data packets.

Use the `no` form of this command to delete the profile.

Command Syntax

```
qos profile queue-color-to-dscp (NAME|default)
no qos profile queue-color-to-dscp NAME
```

Parameters

`NAME`	Profile name
`default`	Default name

Default

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#qos profile queue-color-to-dscp nw_profile
(config-egress-dscp-map)#exit
(config)#no qos profile queue-color-to-dscp nw_profile
```

queue cos

Use this command to configure user defined mapping for queue and cos.

Use the no form of this command to remove the mapping.

Command Syntax

```
queue <0-7> (color(green|yellow|red|all)|) cos <0-7>
no queue <0-7> (color(green|yellow|red|all)|)
```

Parameters

`<0-7>`	Queue and cos range
`color`	Color to map
`all`	Set mapping for all packets
`green`	Set mapping for green packets
`red`	Set mapping for red packets
`yellow`	Set mapping for yellow packets

Default

Default queue and cos value is one-one default mapping if it is not configured.

Command Mode

QoS config mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#qos profile queue-color-to-cos ac_profile
(config-egress-cos-map)#queue 5 cos 2
(config-egress-cos-map)#no queue 5
```

queue dscp

Use this command to configure user-defined mapping for queue to DSCP. This will be mapped with nvo VXLAN tunnel of local VTEP.

Use the `no` form of this command to remove the queue-to-DSCP mapping.

Command Syntax

```
queue <0-7> (color(green|yellow|red|all)|) dscp <0-63>
no queue <0-7> (color(green|yellow|red|all)|)
```

Parameters

`<0-7>`	Queue number
`<0-63>`	DSCP
`color`	Color to map
`all`	Set mapping for all packets
`green`	Set mapping for green packets
`red`	Set mapping for red packets
`yellow`	Set mapping for yellow packets

Default

Default queue and cos value is one-one default mapping if it is not configured.

Command Mode

QoS config mode

Applicability

This command was introduced in OcNOS version 1.3.

Example

```
(config)#qos profile queue-color-to-dscp nw_profile
(config-egress-dscp-map)# queue 1 dscp 63
(config-egress-dscp-map)#no queue 1
```