Virtual Extensible LAN 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.

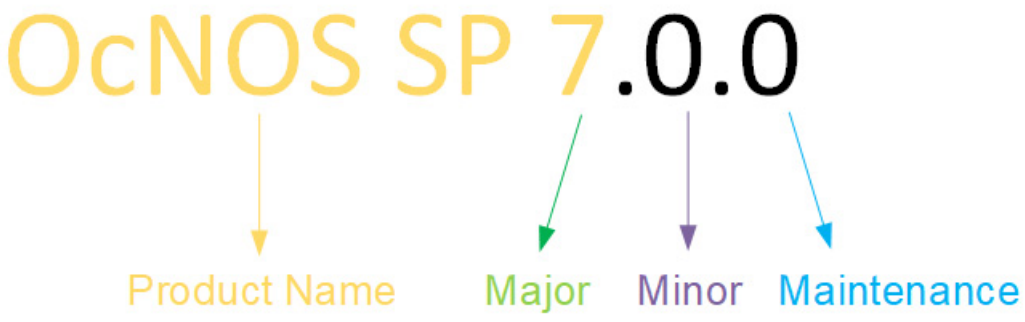

Product Name: IP Infusion Product Family

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

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.

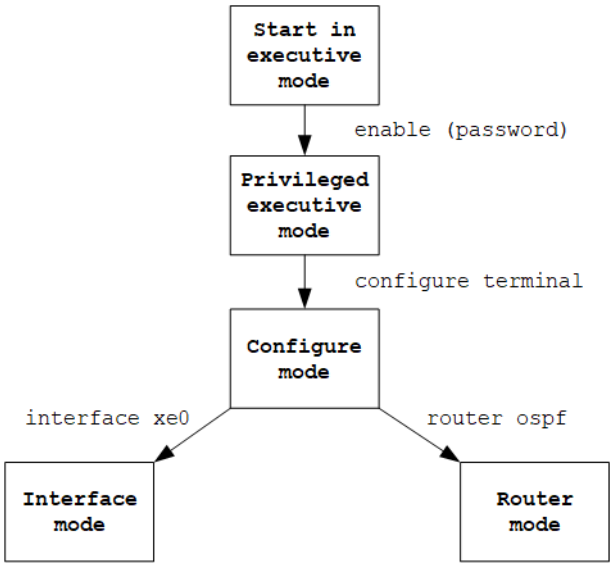

Figure 1: Common command modes

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 Overview

This chapter provides an overview of Virtual eXtensible Local Area Network (VXLAN) and its implementation with OcNOS. VXLAN 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 server using IP tunnels. The endpoints of the virtualized tunnel formed using VXLAN are called VTEPs (VXLAN Tunnel EndPoints). This technology allows the network to support several tenants with minimum changes in the network. The VTEPs carry tenant data in L3 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/VM's.

OcNOS supports VXLAN IPv4 tunnels, but both IPv4 and IPv6 hosts are supported.

Note: To configure VXLAN mapping to access ports, use one of the following two methods:

- Enable NVO access-if mode on a physical interface.: This method is supported in the Qumran1 series platforms.
- Activate access-if-evpn mode on an L2 sub-interface.: This method is supported in both Qumran1 and Qumran2 series platforms.
- Due to the BCM limitation on 9600-56DX boards, if TWAMP and VXLAN are used simultaneously, the hardware profile for NVO VXLAN should be configured last.

Terminology

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

IGMP	Internet Group Management Protocol
PIM	Protocol Independent Multicast
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 Layer 2 overlay network over which VMs communicate
VXLAN Gateway	An entity that forwards traffic between VXLANs

VXLAN Architecture

VXLAN runs over the existing networking infrastructure. It provides a means to "stretch" a Layer 2 network. In short, VXLAN is a Layer 2 overlay scheme on a Layer 3 network.

Each overlay is termed as a VXLAN segment. Only VMs within the same VXLAN segment can communicate with each other. 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.

VNI identifies the scope of the inner MAC frame originated by the individual VM. Hence, we can have overlapping MAC addresses across segments but never have traffic “cross over” since the traffic is isolated using the VNI. The VNI is in an outer header that encapsulates the inner MAC frame originated by the VM.

Any packets (including ARP-ND) that are uplifted to the VXLAN CPU queue from any port are rate limited to 500 packets/second. This is done to protect the system and CPU during an ARP storm.

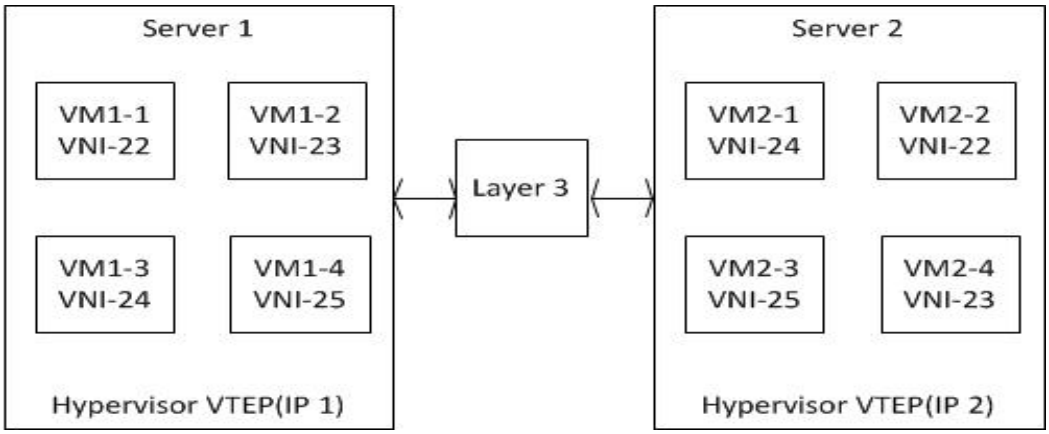

Figure 1-1: VXLAN Deployment - VTEPs across a Layer 3 Network

VxLAN - Data Forwarding Configuration

CHAPTER 1 VxLAN Unicast Configuration

This chapter contains basic Static-VxLAN unicast configuration examples.

Port Mapping

In this example, each VTEP (VTEP1 and VTEP2) is a multilayer switch where xe2 is an access port and xe1 is network port.

After we start sending a packet (say untagged ipv4 packet) from VM1, it hits the VTEP1 which does an encapsulation based on the VNID configured and sends it on xe1. Once the packet reaches VTEP2, it is decapsulated and is sent out on access port and reaches the destination VM, VM2.

Topology

The procedures in this section use the following topology:

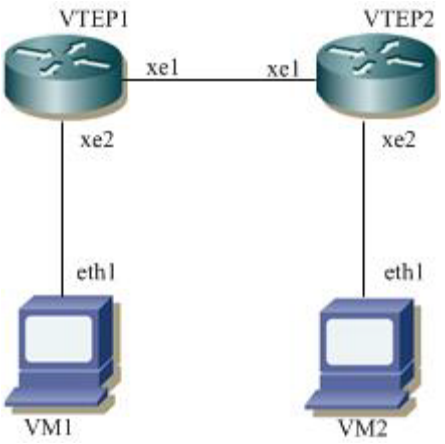

Figure 1-2: VxLAN unicast

VTEP1

`#configure terminal`	Enter the configure mode
`(config)#interface xe2`	Enter interface mode
`(config-if)#switchport`	Configure the interface as switchport
`(config-if)#no shutdown`	Bring the interface into operation with the no shutdown
`(config-if)#exit`	Exit interface mode
`(config)#interface lo`	Enter the configure mode
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure IP address on the interface xe1
`(config-if)#no shutdown`	Bring the interface into operation with the no shutdown
`(config-if)#exit`	Exit interface mode
`(config)#interface xe1`	Enter the configure mode
`(config-if)#ip address 5.5.5.2/24`	Configure IP address on the interface xe1

(config-if)#no shutdown	Bring the interface into operation with the no shutdown
(config-if)#exit	Exit interface mode
(config)#ip route 2.2.2.2/32 5.5.5.4	Configure static route
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN
(config)#nvo vxlan enable	Enable VxLAN globally on this vtep
(config)#interface Tunnel 2	Specify the interface (Tunnel2)to be configured and Enter interface mode
(config-if)#tunnel mode vxlan	Assign this tunnel mode as VxLAN
(config-if)#tunnel source 1.1.1.1	Configure the source IP of this tunnel as an IP address which is configured on the interface xe1 of VTEP1
(config-if)#tunnel destination 2.2.2.2	Configure the destination IP of this tunnel as an IP address which is configured on the interface xe1 of VTEP2.
(config-if)#exit	Exit interface mode
(config-if)#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 inner-vid-disabled	Configure a VNID on this VTEP and enter the nvo mode.
(config-nvo)#vxlan map-network tunnel Tunnel2	Map the tunnel 2 with VNID 2
(config-nvo)#exit	Exit the nvo mode
(config)# interface xe2.10 switchport	Enters sub-interface configuration mode for interface xe2, VLAN 10.
(config-if)# encapsulation dot1q 10	Configure encapsulation dot1q with vlan-id 10
(config-if)# access-if-evpn	Enter access-if-evpn mode
(config-access-if)# map vpn-id 100	Map vnid 100 to access interface xe2.1
(config-access-if)# exit	Exit access-if mode
(config)#commit	Perform commit operation for the changes to take effect.

VTEP2

#configure terminal	Enter the configure mode
(config)#interface xe2	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#no shutdown	Bring the interface into operation with the no shutdown
(config-if)#exit	Exit interface mode
(config)#interface lo	Enter the configure mode
(config-if)#ip address 2.2.2.2/32 secondary	Configure IP address on the interface xe1
(config-if)#no shutdown	Bring the interface into operation with the no shutdown
(config-if)#exit	Exit interface mode
(config)#interface xe1	Enter the configure mode
(config-if)#ip address 5.5.5.4/24	Configure IP address on the interface xe1
(config-if)#no shutdown	Bring the interface into operation with the no shutdown

(config-if)#exit	Exit interface mode
(config)#ip route 1.1.1.1/32 5.5.5.2	Configure static route
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN
(config)#nvo vxlan enable	Enable VxLAN globally on this vtep
(config)#interface tunnel 2	Specify the interface (Tunnel2)to be configured and Enter interface mode.
(config-if)#tunnel mode vxlan	Assign this tunnel mode as VxLAN
(config-if)#tunnel source 2.2.2.2	Configure the source IP of this tunnel as an IP address which is configured on the interface xe1 of VTEP1.
(config-if)#tunnel destination 1.1.1.1	Configure the destination IP of this tunnel as an IP address which is configured on the interface xe1 of VTEP2.
(config-if)#exit	Exit interface mode.
(config-if)#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 inner-vid-disabled	Configure a VNID on this VTEP and enter the nvo mode.
(config-nvo)#vxlan map-network tunnel Tunnel2	Map the tunnel 4 with VNID 2
(config-nvo)#exit	Exit the nvo mode.
(config)# interface xe2.10 switchport	Enters sub-interface configuration mode for interface xe2, VLAN 10.
(config-if)# encapsulation dot1q 10	Configure encapsulation dot1q with vlan-id 10
(config-if)# access-if-evpn	Enter access-if-evpn mode
(config-access-if)# map vpn-id 100	Map vnid 100 to access interface xe2.10
(config-access-if)# exit	Exit access-if mode
(config)#commit	Perform commit operation for the changes to take effect.

Validation

VTEP1

```
VTEP1#sh nvo vxlan tunnel
VxLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
1.1.1.1          2.2.2.2          Installed        00:04:59       00:04:59

Total number of entries are 1
VTEP1#sh nvo vxlan mac-table
================================================================================
================================================================
                                                        VxLAN MAC Entries
================================================================================
================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI
       Type            Status     MAC move AccessPortDesc

________________________________________________________________________________
```

```
_____________________________________________________________

100        xe2      10        ----      903c.b393.e001 1.1.1.1
       Dynamic Local   -------    0        -------

Total number of entries are : 1

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
VxLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type     Interface ESI          VLAN         Vlan
-Range   DF-Status    Src-Addr     Dst-Addr
_____________________________________________________________________________
_____________________________________________
100      ----         L2       NW   ----      ----                          ----
 ----      1.1.1.1          2.2.2.2
100      ----         --       AC   xe2      --- Single Homed Port ---     10
 ----      ----      ----             ----

Total number of entries are 2
VTEP1#

VTEP1#sh nvo vxlan access-if brief

                 Inner   Vlan                    Admin   Link
Interface  Vlan  vlan  Range   Ifindex  Vnid      status  status
-----------------------------------------------------------
xe2       10    ---   ---   500000   100       up      up

Total number of entries are 1
VTEP1#


VTEP1#
```

```
================================================================================
=======
```

VTEP2#

```
VTEP2#sh nvo vxlan tunnel
VxLAN Network tunnel Entries
Source           Destination      Status           Up/Down        Update
================================================================================
2.2.2.2          1.1.1.1          Installed        00:05:30       00:05:30

Total number of entries are 1
VTEP2#sh nvo vxlan mac-table
================================================================================
============================================================
                                                      VxLAN MAC Entries
================================================================================
============================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI
       Type            Status     MAC move AccessPortDesc
________________________________________________________________________________
_____________________________________________________________

100        xe2      10        ----      0018.2359.69b7 2.2.2.2
       Dynamic Local   -------    0        -------

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         Vlan
-Range   DF-Status    Src-Addr     Dst-Addr
________________________________________________________________________________
_____________________________________________
100      ----         L2       NW   ----      ----                          ----
 ----      2.2.2.2          1.1.1.1
100      ----         --       AC   xe2      --- Single Homed Port ---     10
 ----      ----      ----             ----

Total number of entries are 2


VTEP2#sh nvo vxlan access-if brief

                 Inner   Vlan                    Admin   Link
Interface  Vlan  vlan  Range   Ifindex  Vnid      status  status
-----------------------------------------------------------
xe2       10    ---   ---   500000   100       up      up
```

```
Total number of entries are 1
VTEP2#

================================================================================
=====

host-1#ping 10.10.10.2
Press CTRL+C to exit
PING 10.10.10.2 (10.10.10.2) 56(84) bytes of data.
64 bytes from 10.10.10.2: icmp_seq=1 ttl=64 time=0.392 ms

--- 10.10.10.2 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.392/0.392/0.392/0.000 ms


7015-Host#ping 10.10.10.1
Press CTRL+C to exit
PING 10.10.10.1 (10.10.10.1) 56(84) bytes of data.
64 bytes from 10.10.10.1: icmp_seq=1 ttl=64 time=0.704 ms
64 bytes from 10.10.10.1: icmp_seq=2 ttl=64 time=0.259 ms
```

Intermediate Non-VxLAN Router

This example is same as port mapping but with an intermediate non VxLAN router. It does forwarding based on IP header.

Topology

The procedures in this section use the following topology:

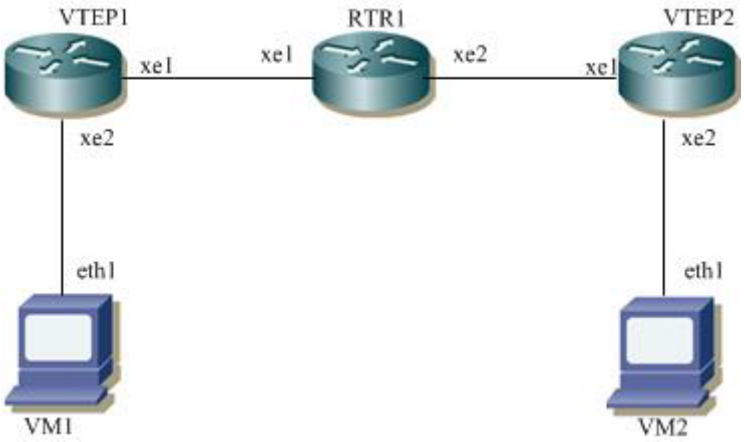

Figure 1-3: VxLAN unicast

VTEP1

`#configure terminal`	Enter the configure mode.
`(config)#interface xe2`	Enter interface mode.
`(config-if)#switchport`	Configure the interface as switchport.
`(config-if)#no shutdown`	Bring the interface into operation with the no shutdown
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe1`	Enter the configure mode.
`(config-if)#ip address 2.2.2.4/24`	Configure IP address on the interface xe1.
`(config-if)#no shutdown`	Bring the interface into operation with the no shutdown
`(config-if)#exit`	Exit interface mode.
`(config)#ip route 3.3.3.0/24 2.2.2.1`	Configuring static route
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN
`(config)#nvo vxlan enable`	Enable VxLAN globally on this vtep.
`(config)#interface tunnel 4`	Specify the interface (Tunnel4)to be configured and Enter interface mode.
`(config-if)#tunnel mode vxlan`	Assign this tunnel mode as VxLAN.
`(config-if)#tunnel source 2.2.2.4`	Configure the source IP of this tunnel as an IP address which is configured on the interface xe1 of VTEP1.
`(config-if)#tunnel destination 3.3.3.4`	Configure the destination IP of this tunnel as an IP address which is configured on the interface xe1 of VTEP2.
`(config-if)#exit`	Exit interface mode.
`(config)#nvo vxlan id 3`	Configure a VNID on this VTEP and enter the nvo mode.
`(config-nvo)#vxlan map-network tunnel Tunnel4`	Map the tunnel 4 with this VNID
`(config-nvo)#vxlan static-entry host-mac 0000.0000.aaaa remote-vtep-ip 3.3.3.4`	Configure a static entry for remote VM with MAC address and IP address.
`(config)# interface xe2 switchport`	Enters sub-interface configuration mode for interface xe2
`(config-if)# encapsulation dot1q`	Configure encapsulation dot1q with vlan-id 3
`(config-if)# access-if-evpn`	Enter access-if-evpn mode
`(config-access-if)# map vpn-id 3`	Map vnid 3 to access interface xe2
`(config-access-if)# exit`	Exit access-if mode
`(config)#commit`	Perform commit operation for the changes to take effect.

RTR1

`(config)#interface xe1`	Enter the configure mode.
`(config-if)#ip address 2.2.2.1/24`	Configure IP address on the interface xe1.
`(config-if)#no shutdown`	Bring the interface into operation with the no shutdown
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter the configure mode.
`(config-if)#ip address 3.3.3.1/24`	Configure IP address on the interface xe2.

(config-if)#no shutdown	Bring the interface into operation with the no shutdown
(config-if)#exit	Exit interface mode.
(config)#commit	Perform commit operation for the changes to take effect.

VTEP2

#configure terminal	Enter the configure mode.
(config)#interface xe2	Enter interface mode.
(config-if)#switchport	Configure the interface as switchport.
(config-if)#no shutdown	Bring the interface into operation with the no shutdown
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter the configure mode.
(config-if)#ip address 3.3.3.4/24	Configure IP address on the interface xe1.
(config-if)#no shutdown	Bring the interface into operation with the no shutdown
(config-if)#exit	Exit interface mode.
(config)#ip route 2.2.2.0/24 3.3.3.1	Configuring static route
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN
(config)#nvo vxlan enable	Enable VxLAN globally on this vtep.
(config)#interface tunnel 4	Specify the interface (Tunnel4)to be configured and Enter interface mode.
(config-if)#tunnel mode vxlan	Assign this tunnel mode as VxLAN.
(config-if)#tunnel source 3.3.3.4	Configure the source IP of this tunnel as an IP address which is configured on the interface xe1 of VTEP1.
(config-if)#tunnel destination 2.2.2.4	Configure the destination IP of this tunnel as an IP address which is configured on the interface xe1 of VTEP2.
(config-if)#exit	Exit interface mode.
(config)#nvo vxlan id 3	Configure a VNID on this VTEP and enter the nvo mode.
(config-nvo)#vxlan map-network tunnel Tunnel4	Map the tunnel 4 with this VNID
(config-nvo)#vxlan static-entry host-mac 0000.0000.bbbb remote-vtep-ip 2.2.2.4	Configure a static entry for remote VM with MAC address and IP address.
(config)# interface xe2 switchport	Enters sub-interface configuration mode for interface xe2
(config-if)# encapsulation dot1q	Configure encapsulation dot1q with vlan-id 3
(config-if)# access-if-evpn	Enter access-if-evpn mode
(config-access-if)# map vpn-id 3	Map vnid 3 to access interface xe2
(config-nvo-acc-if)#exit	Exit the NVO access interface mode
(config)#commit	Perform commit operation for the changes to take effect.

Validation

VTEP1

```
VTEP1#show nvo vxlan tunnel
VxLAN Network tunnel Entries
```

```
 Source           Destination      Status          Up/Down        Update
================================================================================
2.2.2.4          3.3.3.4          Installed     00:37:56     00:37:56
Total number of entries are 1

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   ----          ------                          ---- --
--      2.2.2.4          3.3.3.4
3       ----        --       AC   xe2        --- Single Homed Port ---     ---- ----
----             ----
 Total number of entries are 2
VTEP1#
VTEP1#show nvo vxlan mac-table
================================================================================
=========================================================
======
                                                          VxLAN MAC Entries
================================================================================
=========================================================
======
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                Status              AccessPortDesc
_____________________________________________________________________________
_____________________________________________________________
_______

3         ----       ----   ----         0000.0000.aaaa 3.3.3.4
Static Remote        -------             -------
3         xe2      ----   ----         0000.0000.bbbb                          Dynamic
Local        -------             -------

Total number of entries are : 2

VTEP1#sh nvo vxlan mac-table vnid 3
================================================================================
=========================================================
======
                                                          VxLAN MAC Entries
================================================================================
=========================================================
======
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                Status              AccessPortDesc
```

```
_____________________________________________________________________________________________
_____________________________________________________
______

3         ----       ----   ----         0000.0000.aaaa 3.3.3.4
Static Remote        -------              -------
3         xe2      ----   ----         0000.0000.bbbb                          Dynamic
Local        -------             -------

Total number of entries are : 2
```

VTEP2

```
VTEP2#show nvo vxlan tunnel
VxLAN Network tunnel Entries
 Source            Destination      Status          Up/Down        Update
================================================================================
3.3.3.4          2.2.2.4          Installed     00:34:02     00:34:02
Total number of entries are 1
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   ----          ------                           ---- --
--      3.3.3.4          2.2.2.4
3        ----         --       AC   xe1           --- Single Homed Port ---     ---- --
--      ----             ----
 Total number of entries are 2

VTEP2#sh nvo vxlan mac-table
================================================================================================
=========================================================
======
                                                              VxLAN MAC Entries
================================================================================================
=========================================================
======
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc
_____________________________________________________________________________________________
_____________________________________________________
______

3         xe1        ----   ----         0000.0000.aaaa                          Dynamic
Local        -------             -------
3         ----       ----   ----         0000.0000.bbbb 2.2.2.4
Static Remote        -------              -------
```

```
Total number of entries are : 2

VTEP2#sh nvo vxlan mac-table vnid 3
===========================================================================================
==========================================================
======
                                                          VxLAN MAC Entries
===========================================================================================
==========================================================
======
VNID      Interface  VlanId Inner-VlanId Mac-Addr       VTEP-Ip/ESI
Type                 Status              AccessPortDesc
___________________________________________________________________________________________
__________________________________________________________
______

3         xe1        ----   ----         0000.0000.aaaa                          Dynamic
Local        -------            -------
3         ----       ----   ----         0000.0000.bbbb 2.2.2.4
Static Remote        -------             -------

Total number of entries are : 2
```

VxLAN - Ethernet Virtual Private Network

CHAPTER 1 VXLAN-EVPN Configuration

This section provides basic VXLAN-EVPN configuration examples.

Virtual eXtended LAN (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). VTEPs carry tenant data in L3 tunnels over the network, allowing the network to support multiple tenants. The tenant data is not used in routing or switching, aiding tenant machine movement and allowing tenants to use the same IP/MAC addresses.

In VXLAN, BGP-MP carries crucial information about the VM to reach the VTEP.

MAC VRF supports both manual RT and auto RT modes.

Note: For port-channel/static-channel interfaces, storm control applies to each member port. For example, if interface eth1 and interface eth2 are part of port-channel po1 with storm control set to 2 Mbps for broadcast traffic, then the storm control settings will apply to each member port, rate limiting broadcast traffic to 2 Mbps per port.

Topology

The procedures in this section use the following topology:

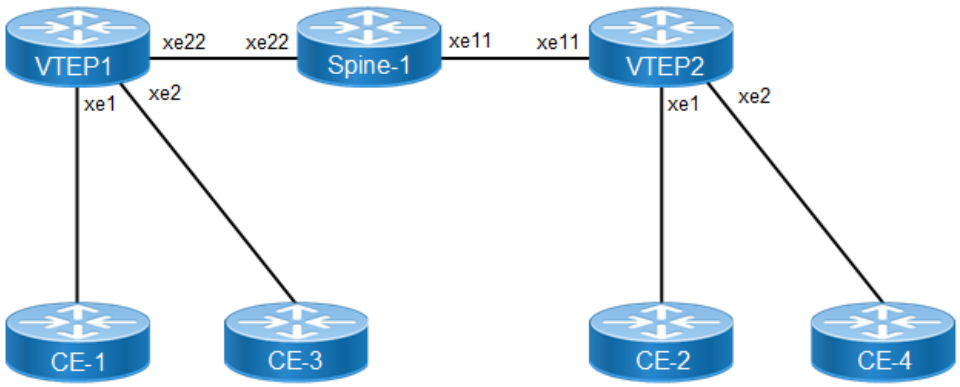

Figure 1-4: VXLAN EVPN

Configure VXLAN mapping to access ports using one of the following methods:

- Enable NVO access-if mode on a physical interface.
- Activate access-if-evpn mode on an L2 sub-interface.

VTEP1

Enable NVO access-if mode on a physical interface.

`#configure terminal`	Enter configure mode
`(config)#interface xe1`	Enter interface mode
`(config-if)#switchport`	Configure the interface as a switch port.
`(config-if)#no shutdown`	Bring the interface into operation
`(config-if)#exit`	Exit interface mode
`(config)#interface xe2`	Enter interface mode

(config-if)#switchport	Configure the interface as switch port.
(config-if)#no shutdown	Bring the interface into operation.
(config-if)#exit	Exit interface mode.
(config)#interface xe22	Enter interface mode.
(config-if)#ip address 10.1.1.1/24	Set an IP address on the interface.
(config-if)#no shutdown	Bring the interface into operation.
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 1.1.1.1/32 secondary	Set an IP address on the interface.
(config-if)#no shutdown	Bring the interface into operation.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance (ABC).
(config-router)#is-type level-1	Configure instance as level-1-only routing.
(config-router)#net 49.0001.1111.1111.1111.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit router mode.
(config)#interface xe22	Enter interface mode
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#mac vrf vrf_evpn_100	Configure a new VRF named vrf_evpn_100.
(config-vrf)#rd 1.1.1.1:1	Assign the Route Distinguisher value.
(config-vrf)#route-target both 100:1	Configure route target to import and export the routes.
(config-vrf)#exit	Exit VRF mode.
(config)#mac vrf vrf_evpn_200	Configure a new VRF named vrf_evpn_200.
(config-vrf)#rd 1.1.1.1:2	Assign the Route Distinguisher value.
(config-vrf)#route-target both evpn-auto-rt	Configure route target to import and export the routes.
(config-vrf)#exit	Exit VRF mode.
(config)#router bgp 1	Define the routing process. The number 1 specifies the AS number of VTEP1.
(config-router)#bgp router-id 1.1.1.1	Configure router-id for this BGP process.
(config-router)#neighbor 2.2.2.2 remote-as 1	Define BGP neighbor: 2.2.2.2 is the IP address of the neighbor (VTEP2), and 1 is the neighbor's AS number.
(config-router)# neighbor 2.2.2.2 update-source 1.1.1.1	Define BGP neighbor: 1.1.1.1 is the peer interface.
(config-router)#address-family l2vpn evpn	Configure address-family L2VPN EVPN.

`(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor in the EVPN address family.
`(config-router-af)#exit-address-family`	Exit the address-family mode.
`(config-router)#exit`	Exit router mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VXLAN.
`(config)#nvo vxlan enable`	Enable VXLAN globally on this VTEP.
`(config)#nvo vxlan vtep-ip-global 1.1.1.1`	Assign a global IP to the VTEP.
`(config)#nvo vxlan id 100 ingress-replication`	Configure a VNID on this VTEP and enter NVO mode.
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_100`	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with `vrf_evpn_100`.
`(config-nvo)#exit`	Exit NVO mode.
`(config)# interface xe1.2 switchport`	Enters sub-interface configuration mode for interface xe1.2
`(config-if)# encapsulation dot1q`	Configure encapsulation dot1q with vlan-id 100
`(config-if)# access-if-evpn`	Enter access-if-evpn mode
`(config-access-if)# map vpn-id 100`	Map vnid 100 to access interface xe1.2
`(config-nvo-acc-if)#exit`	Exit NVO access-if mode.
`(config)#nvo vxlan id 200 ingress-replication`	Configure second VNID on this VTEP and enter NVO mode.
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_200`	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with `vrf_evpn_200`
`(config-nvo)#exit`	Exit NVO mode.
`(config)# interface xe2.3 switchport`	Enters sub-interface configuration mode for interface xe2.3
`(config-if)# encapsulation dot1q`	Configure encapsulation dot1q with vlan-id 200
`(config-if)# access-if-evpn`	Enter access-if-evpn mode
`(config-access-if)# map vpn-id 200`	Map vnid 200 to access interface xe2.3
`(config-nvo-acc-if)#exit`	Exit NVO access-if mode.
`(config)#commit`	Commit the configurations

Activate access-if-evpn mode on an L2 sub-interface.

(config)#interface xe1.1 switchport	Create L2 subinterface xe1.1
(config-if)#encapsulation dot1q 2	Configure encapsulation dot1q with vlan-id 2
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 100	Map vnid 100 to access interface xe1.1
(config-acc-if-evpn)#exit	Exit access-if-evpn mode
(config)#interface xe2.1 switchport	Create L2 subinterface xe2.1
(config-if)#encapsulation dot1q 1	Configure encapsulation dot1q with vlan-id 1
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 200	Map vnid 200 to access interface xe2.1
(config-acc-if-evpn)#exit	Exit access-if-evpn mode
(config-if)#commit	Commit the configurations

Spine-1

#configure terminal	Enter configure mode.
(config)#interface xe22	Enter interface mode.
(config-if)#ip address 10.1.1.2/24	Set an IP address on the interface.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe11	Enter interface mode.
(config-if)#ip address 20.1.1.2/24	Set an IP address on the interface.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance (ABC).
(config-router)#is-type level-1	Configure instance as level-1-only routing.
(config-router)#net 49.0001.3333.3333.3333.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit router mode.
(config)#interface xe22	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface xe11	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the configurations

VTEP2

Enable NVO access-if mode on a physical interface.

`#configure terminal`	Enter configure mode.
`(config)#interface xe1`	Enter interface mode.
`(config-if)#switchport`	Configure the interface as switchport.
`(config-if)#no shutdown`	Bring the interface into operation
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter interface mode.
`(config-if)#switchport`	Configure the interface as switchport.
`(config-if)#no shutdown`	Bring the interface into operation
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe11`	Enter interface mode.
`(config-if)#ip address 20.1.1.1/24`	Set an IP address on the interface.
`(config-if)#no shutdown`	Bring the interface into operation
`(config-if)#exit`	Exit interface mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip address 2.2.2.2/32 secondary`	Set an IP address on the interface.
`(config-if)#no shutdown`	Bring the interface into operation
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance (ABC).
`(config-router)#is-type level-1`	Configure instance as level-1-only routing.
`(config-router)#net 49.0001.2222.2222.2222.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#exit`	Exit router mode.
`(config)#interface xe11`	Enter interface mode.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface (ABC).
`(config-if)#isis circuit-type level-1`	Configure instance as level-1-only routing.
`(config-if)#exit`	Exit interface mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area 49 (ABC).
`(config-if)#isis circuit-type level-1`	Configure instance as level-1-only routing.
`(config-if)#exit`	Exit interface mode.
`(config)#mac vrf vrf_evpn_100`	Configure a new VRF named `vrf_evpn_100`.
`(config-vrf)#rd 2.2.2.2:1`	Assign the Route Distinguisher value.
`(config-vrf)#route-target both 100:1`	Configure route target to import and export the routes.
`(config-vrf)#exit`	Exit VRF mode.
`(config)#mac vrf vrf_evpn_200`	Configure a new VRF named vrf_evpn_200.
`(config-vrf)#rd 2.2.2.2:2`	Assign the Route Distinguisher value.
`(config-vrf)#route-target both evpn-auto-rt`	Configure route target to import and export the routes.
`(config-vrf)#exit`	Exit VRF mode.

`(config)#router bgp 1`	Define the routing process. The number 1 specifies the AS number of VTEP1.
`(config-router)#neighbor 1.1.1.1 remote-as 1`	Define BGP neighbor: 1.1.1.1 is the IP address of the neighbor (VTEP1), and 1 is the neighbor's AS number.
`(config-router)# neighbor 1.1.1.1 update-source 2.2.2.2`	Define BGP neighbor: 2.2.2.2 is the peer interface.
`(config-router)#address-family l2vpn evpn`	Configure address-family L2VPN EVPN.
`(config-router-af)#neighbor 1.1.1.1 activate`	Activate the neighbor in the EVPN address family.
`(config-router-af)#exit-address-family`	Exit address-family mode.
`(config-router)#exit`	Exit router mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VXLAN.
`(config)#nvo vxlan enable`	Enable VXLAN globally on this VTEP.
`(config)#nvo vxlan vtep-ip-global 2.2.2.2`	Assign a global IP to the VTEP.
`(config)#nvo vxlan id 100 ingress-replication`	Configure a VNID on this VTEP and enter NVO mode.
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_100`	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with `vrf_evpn_100`.
`(config-nvo)#exit`	Exit NVO mode.
`(config)#interface xe1.2 switchport`	Create L2 subinterface xe1.2
`(config-if)#encapsulation dot1q 2`	Configure encapsulation dot1q with vlan-id 2
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 100`	Map vnid 100 to access interface xe1.2
`(config-nvo-acc-if)#exit`	Exit NVO access-if mode.
`(config)#nvo vxlan id 200 ingress-replication`	Configure second VNID on this VTEP and enter NVO mode.
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_200`	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with `vrf_evpn_200`.
`(config-nvo)#exit`	Exit NVO mode.
`(config)#interface xe2.3 switchport`	Create L2 subinterface xe2.3
`(config-if)#encapsulation dot1q 3`	Configure encapsulation dot1q with vlan-id 3
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 200`	Map vnid 200 to access interface xe2.3
`(config-nvo-acc-if)#exit`	Exit NVO access-if mode.
`(config)#commit`	Commit the configurations

Activate access-if-evpn mode on an L2 sub-interface.

`(config)#interface xe1.1 switchport`	Create L2 subinterface xe1.1
`(config-if)#encapsulation dot1q 2`	Configure encapsulation dot1q with vlan-id 2
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 100`	Map vnid 100 to access interface xe1.1
`(config-acc-if-evpn)#exit`	Exit access-if-evpn mode
`(config)#interface xe2.1 switchport`	Create L2 subinterface xe2.1

`(config-if)#encapsulation dot1q 3`	Configure encapsulation dot1q with vlan-id 3
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 200`	Map vnid 200 to access interface xe2.1
`(config-acc-if-evpn)#exit`	Exit access-if-evpn mode
`(config-if)#commit`	Commit the configurations

Validation

CE1 and CE2 have hosts configured with MAC addresses, IP addresses, and VLAN identifiers as shown below.

		VLAN	IP Address	Mac Address
VTEP1	CE-1	2	12.12.12.10	0000.0000.abab
VTEP2	CE-2	2	12.12.12.20	0000.0000.cdcd
VTEP1	CE-1	3	13.13.13.10	0000:0b60:25f2
VTEP2	CE-2	3	13.13.13.20	0000:0b60:25f3

Perform a tagged ping of VLAN 2 from CE1 to CE2 and vice-versa. Also perform a tagged ping of VLAN 3 from CE1 to CE2 and vice-versa.

VTEP Tunnel Status

```
VTEP1#show nvo vxlan tunnel
VxLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
1.1.1.1          2.2.2.2          Installed     00:05:53     00:05:53
Total number of entries are 1


VTEP2#show nvo vxlan tunnel
VxLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
2.2.2.2          1.1.1.1          Installed     00:05:46     00:05:46
Total number of entries are 1
```

VTEP ARP Cache

```
VTEP1#sh nvo vxlan arp-cache
VxLAN ARP-CACHE Information
============================
ARP Timeout : 300 sec Random-Jitter-Max : 640
VNID    Ip-Addr          Mac-Addr              Type                    Age-Out
Retries-Left

_____________________________________________________________________________
200      13.13.13.10    0000.0b60.25f2     Dynamic Local        246              2
200      13.13.13.20    0000.0b60.25f3     Dynamic Remote   ------
```

```
100      12.12.12.20    0000.0000.cdcd     Dynamic Remote   ------
100      12.12.12.10   0000.0000.abab     Dynamic Local         246                  2
Total number of entries are 4

VTEP2#sh nvo vxlan arp-cache
VxLAN ARP-CACHE Information
===========================
ARP Timeout : 300 sec Random-Jitter-Max : 640
VNID      Ip-Addr              Mac-Addr                Type                        Age-
Out       Retries-Left

_____________________________________________________________________________
200     13.13.13.10        0000.0b60.25f2       Dynamic Remote      -------
200     13.13.13.20       0000.0b60.25f3     Dynamic Local            257
2
100     12.12.12.10        0000.0000.abab      Dynamic Remote      -------
100     12.12.12.20       0000.0000.cdcd     Dynamic Local            257
2
Total number of entries are 4
```

VTEP MAC Tables

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

_____________________________________________________________________________

100      ----              ----           ----        0000.0000.cdcd      2.2.2.2
Dynamic Remote   -------    -------
100      xe1               2             ----        0000.0000.abab      1.1.1.1
Dynamic Local       -------    -------
200      xe2               3            ----        0000.0b60.25f2       1.1.1.1
Dynamic Local       -------    -------
200      ----             ----           ----        0000.0b60.25f3      2.2.2.2
Dynamic Remote   -------    -------
Total number of entries are : 4

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

_____________________________________________________________________________

100      xe1            2          ----              0000.0000.cdcd     2.2.2.2
Dynamic Local      -------    -------
100      ----            ----       ----              0000.0000.abab     1.1.1.1
Dynamic Remote -------     -------
200      ----            ----       ----              0000.0b60.25f2      1.1.1.1
Dynamic Remote  -------     -------
```

```
200       xe2             3          ----               0000.0b60.25f3      2.2.2.2
Dynamic Local       -------     -------
Total number of entries are : 4
```

MAC VRF Manual RT

```
VTEP1#show mac vrf vrf_evpn_100
VRF vrf_evpn_100, FIB ID 4098
 Router ID is not set
Interfaces:
VRF vrf_evpn_100; default RD 1.1.1.1:1
 RT:100:1
Import VPN route-target communities
 RT:100:1
No import route-map
No export route-map
VPNv4 label allocation mode: per-vrf
VPNv6 label allocation mode: per-vrf
import-vnid: 100
export-vnid: 100
```

MAC VRF Auto RT

```
VTEP1#show mac vrf vrf_evpn_200
VRF vrf_evpn_200, FIB ID 4098
 Router ID is not set
Interfaces:
VRF vrf_evpn_200; default RD 1.1.1.1:2
 Evpn Auto RT:1:268435856
 Evpn Auto RT:1:268435656
Import VPN route-target communities
 Evpn Auto RT:1:268435856
 Evpn Auto RT:1:268435656
No import route-map
No export route-map
VPNv4 label allocation mode: per-vrf
VPNv6 label allocation mode: per-vrf
import-vnid: 200
export-vnid: 200
```

Access IF EVPN Mode on L2 Sub-interface

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

________________________________________________________________________________
```

```
100       ----              ----            ----         0000.0000.cdcd       2.2.2.2
Dynamic Remote    -------     -------
100       xe1.1              2              ----         0000.0000.abab       1.1.1.1
Dynamic Local         -------     -------
200       xe2.1              3             ----         0000.0b60.25f2       1.1.1.1
Dynamic Local         -------     -------
200       ----              ----            ----         0000.0b60.25f3       2.2.2.2
Dynamic Remote    -------     -------
Total number of entries are : 4
```

VTEP MAC-IP BGP EVPN Entries

```
VTEP1#show bgp l2vpn evpn mac-ip
RD[1.1.1.1:1] VRF[vrf_evpn_100]:
ESI               Eth-Tag      Mac-Address           IP-Address     VNID      L3VNID
Nexthop          GW-Type
0                   100        0000:0000:abab     12.12.12.10    100       0
1.1.1.1              --
0                  200        0000:0b60:25f2      13.13.13.10    200       0
1.1.1.1              --

RD[2.2.2.2:1]
ESI          Eth-Tag     Mac-Address         IP-Address    VNID     L3VNID    Nexthop
GW-Type
0                 100        0000:0000:cdcd     12.12.12.20    100           0
2.2.2.2              --
0                  200        0000:0b60:25f3     13.13.13.20    200           0
2.2.2.2               --


VTEP2#show bgp l2vpn evpn mac-ip
RD[1.1.1.1:1]
ESI              Eth-Tag     Mac-Address         IP-Address           VNID      L3VNID
Nexthop          GW-Type
0                 100        0000:0000:abab     12.12.12.10       100         0
1.1.1.1              --
0                 200        0000:0b60:25f2     13.13.13.10       200         0
1.1.1.1              --

RD[2.2.2.2:1] VRF[vrf_evpn_100]:
ESI                Eth-Tag     Mac-Address         IP-Address       VNID      L3VNID
Nexthop          GW-Type
0                   100        0000:0000:cdcd     12.12.12.20     100         0
2.2.2.2                --
0                  100        0000:0b60:25f2     13.13.13.10     100         0
2.2.2.2                --
```

LAG as Access Port with ECMP on the Network Side

This section contains basic VXLAN EVPN configuration with LAG as an access port and ECMP on the network side.

Topology

The procedures in this section use the following topology:

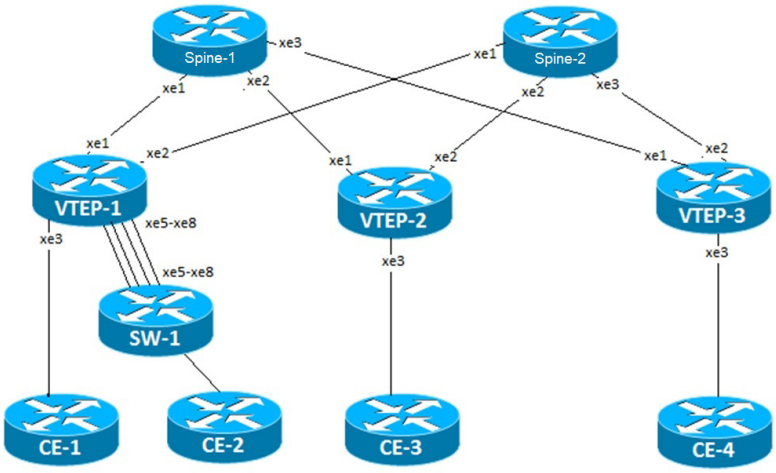

Figure 1-5: VXLAN EVPN with LAG and ECMP

SW-1

`#configure terminal`	Enter configure mode
`(config)#bridge 1 protocol ieee vlan-bridge`	Configure IEEE vlan bridge
`(config)#vlan database`	Enter into the vlan database
`(config-vlan)#vlan 2 bridge 1 state enable`	Configure vlan 2 and associate with bridge 1
`(config-vlan)#vlan 3 bridge 1 state enable`	Configure vlan 3 and associate with bridge 1
`(config-vlan)#exit`	Exit from the vlan database
`(config)#interface xe1`	Enter interface mode
`(config-if)#no shutdown`	Make interface admin up
`(config-if)#switchport`	Set the interface as Layer2 port
`(config-if)#bridge-group 1`	Associate the Interface with bridge-group.
`(config-if)#switchport mode trunk`	Set the switching characteristics of this interface to trunk mode
`(config-if)#switchport trunk allowed vlan add 2`	Configure the VLANs that should be allowed through this interface
`(config-if)#switchport trunk allowed vlan add 3`	Configure the VLANs that should be allowed through this interface
`(config-if)#exit`	Exit interface mode.
`(config)#interface po1`	Enter interface mode
`(config-if)#switchport`	Set the interface as Layer2 port
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe5`	Enter interface mode
`(config-if)#switchport`	Set the interface as Layer2 port
`(config-if)#channel-group 1 mode active`	Configure the interface to be part of port channel 1
`(config-if)#exit`	Exit interface mode.

(config)#interface xe6	Enter interface mode
(config-if)#switchport	Set the interface as Layer2 port
(config-if)#channel-group 1 mode active	Configure the interface to be part of port channel 1
(config-if)#exit	Exit interface mode.
(config)#interface xe7	Enter interface mode
(config-if)#switchport	Set the interface as Layer2 port
(config-if)#channel-group 1 mode active	Configure the interface to be part of port channel 1
(config-if)#exit	Exit interface mode.
(config)#interface xe8	Enter interface mode
(config-if)#switchport	Set the interface as Layer2 port
(config-if)#channel-group 1 mode active	Configure the interface to be part of port channel 1
(config-if)#exit	Exit interface mode.
(config-if)#interface xe2	Enter interface mode
(config-if)#bridge-group 1	Associate the Interface with bridge-group.
(config-if)#switchport mode trunk	Set the switching characteristics of this interface to trunk mode
(config-if)#switchport trunk allowed vlan add 2	Configure the VLANs that should be allowed through this interface
(config-if)#switchport trunk allowed vlan add 3	Configure the VLANs that should be allowed through this interface
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the configurations

VTEP-1

#configure terminal	Enter configure mode
(config)#interface po1	Create interface po1
(config-if)#switchport	Configure the interface as switchport.
(config-if)#exit	Exit interface mode
(config)#interface xe3	Enter interface mode.
(config-if)#switchport	Configure the interface as switchport.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode
(config)#interface xe5	Enter interface mode
(config-if)#switchport	Configure the interface as switchport.
(config-if)#channel-group 1 mode active	Configure the interface to be part of port channel 1
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode
(config)#interface xe6	Enter interface mode
(config-if)#switchport	Configure the interface as switchport.
(config-if)#channel-group 1 mode active	Configure the interface to be part of port channel 1
(config-if)#no shutdown	Bring the interface into operation

(config-if)#exit	Exit interface mode
(config)#interface xe7	Enter interface mode
(config-if)#switchport	Configure the interface as switchport.
(config-if)#channel-group 1 mode active	Configure the interface to be part of port channel 1
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode
(config)#interface xe8	Enter interface mode
(config-if)#switchport	Configure the interface as switchport.
(config-if)#channel-group 1 mode active	Configure the interface to be part of port channel 1
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 10.1.1.0/31	Configure IP address on the interface xe1.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 10.1.1.2/31	Configure IP address on the interface xe2.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 1.1.1.1/32 secondary	Configure IP address on the loopback interface.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance (ABC).
(config-router)#is-type level-1	Configure instance as level-1-only routing.
(config-router)#net	Set a Network Entity Title for this instance, specifying the
49.0001.1111.1111.1111.00	area address and the system ID.
(config-router)#exit	Exit router mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.

(config)#mac vrf vrf_evpn_100	Configure a new VRF named vrf_evpn_100.
(config-vrf)#rd 1.1.1.1:1	Assign the Route Distinguisher value.
(config-vrf)#route-target both 100:1	Configure route target to import and export the routes.
(config-vrf)#exit	Exit VRF mode.
(config)#router bgp 65535	Define the routing process. The number 65535 specifies the AS number of VTEP-1.
(config-router)#bgp router-id 1.1.1.1	Configure router-id for this BGP process.
(config-router)#neighbor 2.2.2.2 remote-as 65535	Define BGP neighbor: 2.2.2.2 is the IP address of the neighbor (VTEP-2), and 65535 is the neighbor's AS number.
(config-router)# neighbor 2.2.2.2 update-source 1.1.1.1	Define BGP neighbor: 1.1.1.1 is the peer interface.
(config-router)#neighbor 3.3.3.3 remote-as 65535	Define BGP neighbor: 3.3.3.3 is the IP address of the neighbor (VTEP-3), and 65535 is the neighbor's AS number.
(config-router)# neighbor 3.3.3.3 update-source 1.1.1.1	Define BGP neighbor: 1.1.1.1 is the peer interface.
(config-router)#address-family l2vpn evpn	Configure address-family L2VPN EVPN.
(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor at VTEP-2 in the EVPN address family.
(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor at VTEP-3 in the EVPN address family.
(config-router-af)#exit-adress-family	Exit address-family mode.
(config-router)#exit	Exit router mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VXLAN.
(config)#nvo vxlan enable	Enable VXLAN globally on this VTEP.
(config)#nvo vxlan vtep-ip-global 1.1.1.1	Assign a global IP to the VTEP.
(config)#nvo vxlan id 100001 ingress-replication	Configure a VNID on this VTEP and enter NVO mode.
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_100	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with VRF vrf_evpn_100.
(config-nvo)#exit	Exit NVO mode.
(config)#interface xe3.2 switchport	Create L2 subinterface xe3.2
(config-if)#encapsulation dot1q 2	Configure encapsulation dot1q with vlan-id 2
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 100001	Map vnid 100001 to access interface xe3.2
(config-nvo-acc-if)#exit	Exit NVO access-if mode.
(config)#interface xe1.2 switchport	Create L2 subinterface xe1.2
(config-if)#encapsulation dot1q 2	Configure encapsulation dot1q with vlan-id 2
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 100001	Map vnid 100001 to access interface xe1.2
(config-nvo-acc-if)#exit	Exit NVO access-if mode.
(config)#nvo vxlan id 200001 ingress-replication	Configure second VNID on this VTEP and enter NVO mode.
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_100	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with VRF vrf_evpn_100.

(config-nvo)#exit	Exit NVO mode.
(config)#interface xe3.3 switchport	Create L2 subinterface xe3.3
(config-if)#encapsulation dot1q 3	Configure encapsulation dot1q with vlan-id 3
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 200001	Map vnid 200001 to access interface xe3.3
(config-nvo-acc-if)#exit	Exit NVO access-if mode.
(config)#interface xe1.3 switchport	Create L2 subinterface xe1.3
(config-if)#encapsulation dot1q 3	Configure encapsulation dot1q with vlan-id 3
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 200001	Map vnid 200001 to access interface xe1.3
(config-nvo-acc-if)#exit	Exit NVO access-if mode.
(config)#commit	Commit the configurations

Spine-1

#configure terminal	Enter configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 12.12.12.12/32 secondary	Set an IP address on the interface.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 10.1.1.1/31	Configure IP address on the interface xe1.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 20.1.1.1/31	Configure IP address on the interface xe2.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe3	Enter intereface mode.
(config-if)#ip address 30.1.1.1/31	Configure IP address on the interface xe3.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance (ABC).
(config-router)#is-type level-1	Configure instance as level-1-only routing.
(config-router)#net 49.0001.4444.4444.4444.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit router mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.

(config)#interface xe1	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface xe3	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the configurations

Spine-2

#configure terminal	Enter configure mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 13.13.13.13/32 secondary	Set an IP address on the interface.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 10.1.1.3/31	Configure IP address on the interface xe1.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 20.1.1.3/31	Configure IP address on the interface xe2.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe3	Enter interface mode.
(config-if)#ip address 30.1.1.3/31	Configure IP address on the interface xe3.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance (ABC).
(config-router)#is-type level-1	Configure instance as level-1-only routing.
(config-router)#net 49.0001.5555.5555.5555.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit router mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.

(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface xe3	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the configurations

VTEP-2

#configure terminal	Enter configure mode
(config)#interface xe3	Enter interface mode
(config-if)#switchport	Configure the interface as switchport.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 20.1.1.0/31	Configure IP address on the interface xe1.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 20.1.1.2/31	Configure IP address on the interface xe2.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 2.2.2.2/32 secondary	Configure IP address on the loopback interface.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance (ABC).
(config-router)#is-type level-1	Configure instance as level-1-only routing.
(config-router)#net 49.0001.2222.2222.2222.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit router mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).

(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#mac vrf vrf_evpn_100	Configure a new VRF named vrf_evpn_100.
(config-vrf)#rd 1.1.1.1:1	Assign the Route Distinguisher value.
(config-vrf)#route-target both 100:1	Configure route target to import and export the routes.
(config-vrf)#exit	Exit VRF mode.
(config)#router bgp 65535	Define the routing process. The number 65535 specifies the AS number of VTEP-1.
(config-router)#neighbor 1.1.1.1 remote-as 65535	Define BGP neighbor: 1.1.1.1 is the IP address of the neighbor (VTEP-1), and 65535 is the neighbor's AS number.
(config-router)#neighbor 1.1.1.1 update-source 2.2.2.2	Define BGP neighbor: 2.2.2.2 is the peer interface.
(config-router)#neighbor 3.3.3.3 remote-as 65535	Define BGP neighbor: 3.3.3.3 is the IP address of the neighbor (VTEP-3), and 65535 is the neighbor's AS number.
(config-router)#neighbor 3.3.3.3 update-source 2.2.2.2	Define BGP neighbor: 2.2.2.2 is the peer interface.
(config-router)#address-family l2vpn evpn	Configure address-family L2VPN EVPN.
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor at VTEP-1 in the EVPN address family.
(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor at VTEP-3 in the EVPN address family.
(config-router-af)#exit-adress-family	Exit address-family mode.
(config-router)#exit	Exit router mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VXLAN.
(config)#nvo vxlan enable	Enable VXLAN globally on this VTEP.
(config)#nvo vxlan vtep-ip-global 2.2.2.2	Assign a global IP to the VTEP.
(config)#nvo vxlan id 100001 ingress-replication	Configure a VNID on this VTEP and enter NVO mode.
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_100	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with VRF vrf_evpn_100.
(config-nvo)#exit	Exit NVO mode.
(config)#interface xe3.2 switchport	Create L2 subinterface xe3.2
(config-if)#encapsulation dot1q 2	Configure encapsulation dot1q with vlan-id 2
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 100001	Map vnid 100001 to access interface xe3.2

(config-nvo-acc-if)#exit	Exit NVO access-if mode.
(config)#nvo vxlan id 200001 ingress-replication	Configure second VNID on this VTEP and enter NVO mode.
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_100	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with VRF vrf_evpn_100.
(config-nvo)#exit	Exit NVO mode.
(config)#interface xe3.3 switchport	Create L2 subinterface xe3.3
(config-if)#encapsulation dot1q 3	Configure encapsulation dot1q with vlan-id 3
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 200001	Map vnid 200001 to access interface xe3.3
(config-nvo-acc-if)#exit-adress-family	Exit NVO access-if mode.
(config)#commit	Commit the configurations

VTEP-3

#configure terminal	Enter configure mode.
(config)#interface xe3	Enter interface mode.
(config-if)#switchport	Configure the interface as switchport.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 30.1.1.0/31	Configure IP address on the interface xe1.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 30.1.1.2/31	Configure IP address on the interface xe2.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 3.3.3.3/32 secondary	Configure IP address on the loopback interface.
(config-if)#no shutdown	Bring the interface into operation
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance (ABC).
(config-router)#is-type level-1	Configure instance as level-1-only routing.
(config-router)#net 49.0001.3333.3333.3333.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit router mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface (ABC).
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode.

`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface (ABC).
`(config-if)#isis circuit-type level-1`	Configure instance as level-1-only routing.
`(config-if)#exit`	Exit interface mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area 49 (ABC).
`(config-if)#isis circuit-type level-1`	Configure instance as level-1-only routing.
`(config-if)#exit`	Exit interface mode.
`(config)#mac vrf vrf_evpn_100`	Configure a new VRF named `vrf_evpn_100`.
`(config-vrf)#rd 1.1.1.1:1`	Assign the Route Distinguisher value.
`(config-vrf)#route-target both 100:1`	Configure route target to import and export the routes.
`(config-vrf)#exit`	Exit VRF mode.
`(config)#router bgp 65535`	Define the routing process. The number `65535` specifies the AS number of VTEP-1.
`(config-router)#neighbor 1.1.1.1 remote-as 65535`	Define BGP neighbor: 1.1.1.1 is the IP address of the neighbor (VTEP-1), and `65535` is the neighbor's AS number.
`(config-router)#neighbor 1.1.1.1 update-source 3.3.3.3`	Define BGP neighbor: 3.3.3.3 is the peer interface.
`(config-router)#neighbor 2.2.2.2 remote-as 65535`	Define BGP neighbor: 2.2.2.2 is the IP address of the neighbor (VTEP-3), and 65535 is the neighbor's AS number.
`(config-router)#neighbor 2.2.2.2 update-source 3.3.3.3`	Define BGP neighbor: 3.3.3.3 is the peer interface.
`(config-router)#address-family l2vpn evpn`	Configure address-family L2VPN EVPN.
`(config-router-af)#neighbor 1.1.1.1 activate`	Activate the neighbor at VTEP-1 in the EVPN address family.
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor at VTEP-2 in the EVPN address family.
`(config-router-af)#exit-adress-family`	Exit address-family mode.
`(config-router)#exit`	Exit router mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VXLAN.
`(config)#nvo vxlan enable`	Enable VXLAN globally on this VTEP.
`(config)#nvo vxlan vtep-ip-global 3.3.3.3`	Assign a global IP to the VTEP.
`(config)#nvo vxlan id 100001 ingress-replication`	Configure a VNID on this VTEP and enter NVO mode.
`(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_100`	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with VRF `vrf_evpn_100`.
`(config-nvo)#exit`	Exit NVO mode.
`(config)#interface xe3.2 switchport`	Create L2 subinterface xe3.2
`(config-if)#encapsulation dot1q 2`	Configure encapsulation dot1q with vlan-id 2
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 100001`	Map vnid 100001 to access interface xe3.2
`(config-nvo-acc-if)#exit-address-family`	Exit NVO access-if mode.
`(config)#nvo vxlan id 200001 ingress-replication`	Configure second VNID on this VTEP and enter NVO mode.

`(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf_evpn_100`	Configure host-reachability-protocol as BGP-EVPN and associate the VNID with VRF `vrf_evpn_100`
`(config-nvo)#exit`	Exit NVO mode.
`(config)#interface xe3.3 switchport`	Create L2 subinterface xe3.3
`(config-if)#encapsulation dot1q 3`	Configure encapsulation dot1q with vlan-id 3
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 200001`	Map vnid 200001 to access interface xe3.3
`(config-nvo-acc-if)#exit`	Exit NVO access-if mode.
`(config)#commit`	Commit the configurations

Validation

CE1, CE2, CE3, and CE4 have hosts configured with MAC addresses, IP addresses, and VLAN identifiers as shown below.

		VLAN	IP Address	MAC Address
VTEP-1	CE-1	2	12.12.12.10	0000.0000.aaaa
VTEP-1	CE-2	2	12.12.12.20	0000.0000.bbbb
VTEP-2	CE-3	2	12.12.12.30	0000.0000.cccc
VTEP-3	CE-4	2	12.12.12.40	0000.0000.dddd
VTEP-1	CE-1	3	14.14.14.10	0000.058e.2181
VTEP-1	CE-2	3	14.14.14.20	0000.058e.2182
VTEP-2	CE-3	3	14.14.14.30	0000.058e.2183
VTEP-3	CE-4	3	14.14.14.40	0000.058e.2184

Perform a tagged ping of VLAN 2 from CE1 to CE2,CE3 and CE4 and vice-versa. Also perform a tagged ping of VLAN 3 from CE1 to CE2, CE3 and CE4 and vice-versa.

VTEP Tunnel Status

```
VTEP-1#show nvo vxlan tunnel
VxLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
1.1.1.1          3.3.3.3          Installed     01:25:20     01:25:20
1.1.1.1          2.2.2.2          Installed     01:35:19     01:35:19
Total number of entries are 2

VTEP-2#show nvo vxlan tunnel
VxLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
2.2.2.2          1.1.1.1          Installed     01:35:42     01:35:42
2.2.2.2          3.3.3.3          Installed     01:25:43     01:25:43
```

```
Total number of entries are 2

VTEP-3#show nvo vxlan tunnel
VxLAN Network tunnel Entries
 Source           Destination      Status          Up/Down        Update
================================================================================
3.3.3.3          2.2.2.2          Installed     01:25:35     01:25:35
3.3.3.3          1.1.1.1          Installed     01:25:35     01:25:35
Total number of entries are 2
```

VTEP ARP Tables

Use the following command to validate the ARP-cache details in all VTEPs.

```
VTEP-3#show nvo vxlan arp-cache
VxLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
100001   12.12.12.30      0000.0000.cccc Dynamic Remote ----
100001   12.12.12.10      0000.0000.aaaa Dynamic Remote ----
100001   12.12.12.20      0000.0000.bbbb Dynamic Remote ----
100001   12.12.12.40      0000.0000.dddd Dynamic Local  ----
200001   14.14.14.20      0000.058e.2182 Dynamic Remote ----
200001   14.14.14.10      0000.058e.2181 Dynamic Remote ----
200001   14.14.14.40      0000.058e.2184 Dynamic Local  ----
200001   14.14.14.30      0000.058e.2183 Dynamic Remote ----
Total number of entries are 8
```

VTEP MAC-IP BGP EVPN Entries

Use the following command to validate the BGP EVPN entries in all VTEPs.

```
VTEP-3#show bgp l2vpn evpn mac-ip

RD[1.1.1.1:1] VRF[vrf_evpn_100]:
ESI                       Eth-Tag    Mac-Address    IP-Address
VNID      L3VNID    Nexthop         GW-Type
0                         100001     0000:0000:dddd --
100001    0         3.3.3.3         --
0                         100001     0000:0000:dddd 12.12.12.40
100001    0         3.3.3.3         --
0                         200001     0000:058e:2184 --
200001    0         3.3.3.3         --
0                         200001     0000:058e:2184 14.14.14.40
200001    0         3.3.3.3         --

RD[1.1.1.1:1]
ESI                       Eth-Tag    Mac-Address    IP-Address
VNID      L3VNID    Nexthop         GW-Type
0                         100001     0000:0000:aaaa --
100001    0         1.1.1.1         --
0                         100001     0000:0000:aaaa 12.12.12.10
100001    0         1.1.1.1         --
```

```
0                          100001     0000:0000:bbbb --
100001    0          1.1.1.1          --
0                          100001     0000:0000:bbbb 12.12.12.20
100001    0          1.1.1.1          --
0                          100001     0000:0000:cccc --
100001    0          2.2.2.2          --
0                          100001     0000:0000:cccc 12.12.12.30
100001    0          2.2.2.2          --
0                          200001     0000:058e:2181 --
200001    0          1.1.1.1          --
0                          200001     0000:058e:2181 14.14.14.10
200001    0          1.1.1.1          --
0                          200001     0000:058e:2182 --
200001    0          1.1.1.1          --
0                          200001     0000:058e:2182 14.14.14.20
200001    0          1.1.1.1          --
0                          200001     0000:058e:2183 --
200001    0          2.2.2.2          --
0                          200001     0000:058e:2183 14.14.14.30
200001    0          2.2.2.2          --
```

CHAPTER 2 VxLAN Multi-Homing Configuration

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

Overview

VxLAN EVPN Multi-homing feature enables connecting a CE/Host node to two VTEPs with all-active redundancy mode. A redundant VTEP device can provide network service to the customer site as soon as a failure is detected. The failure can be link or node 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 a 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 configuration 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

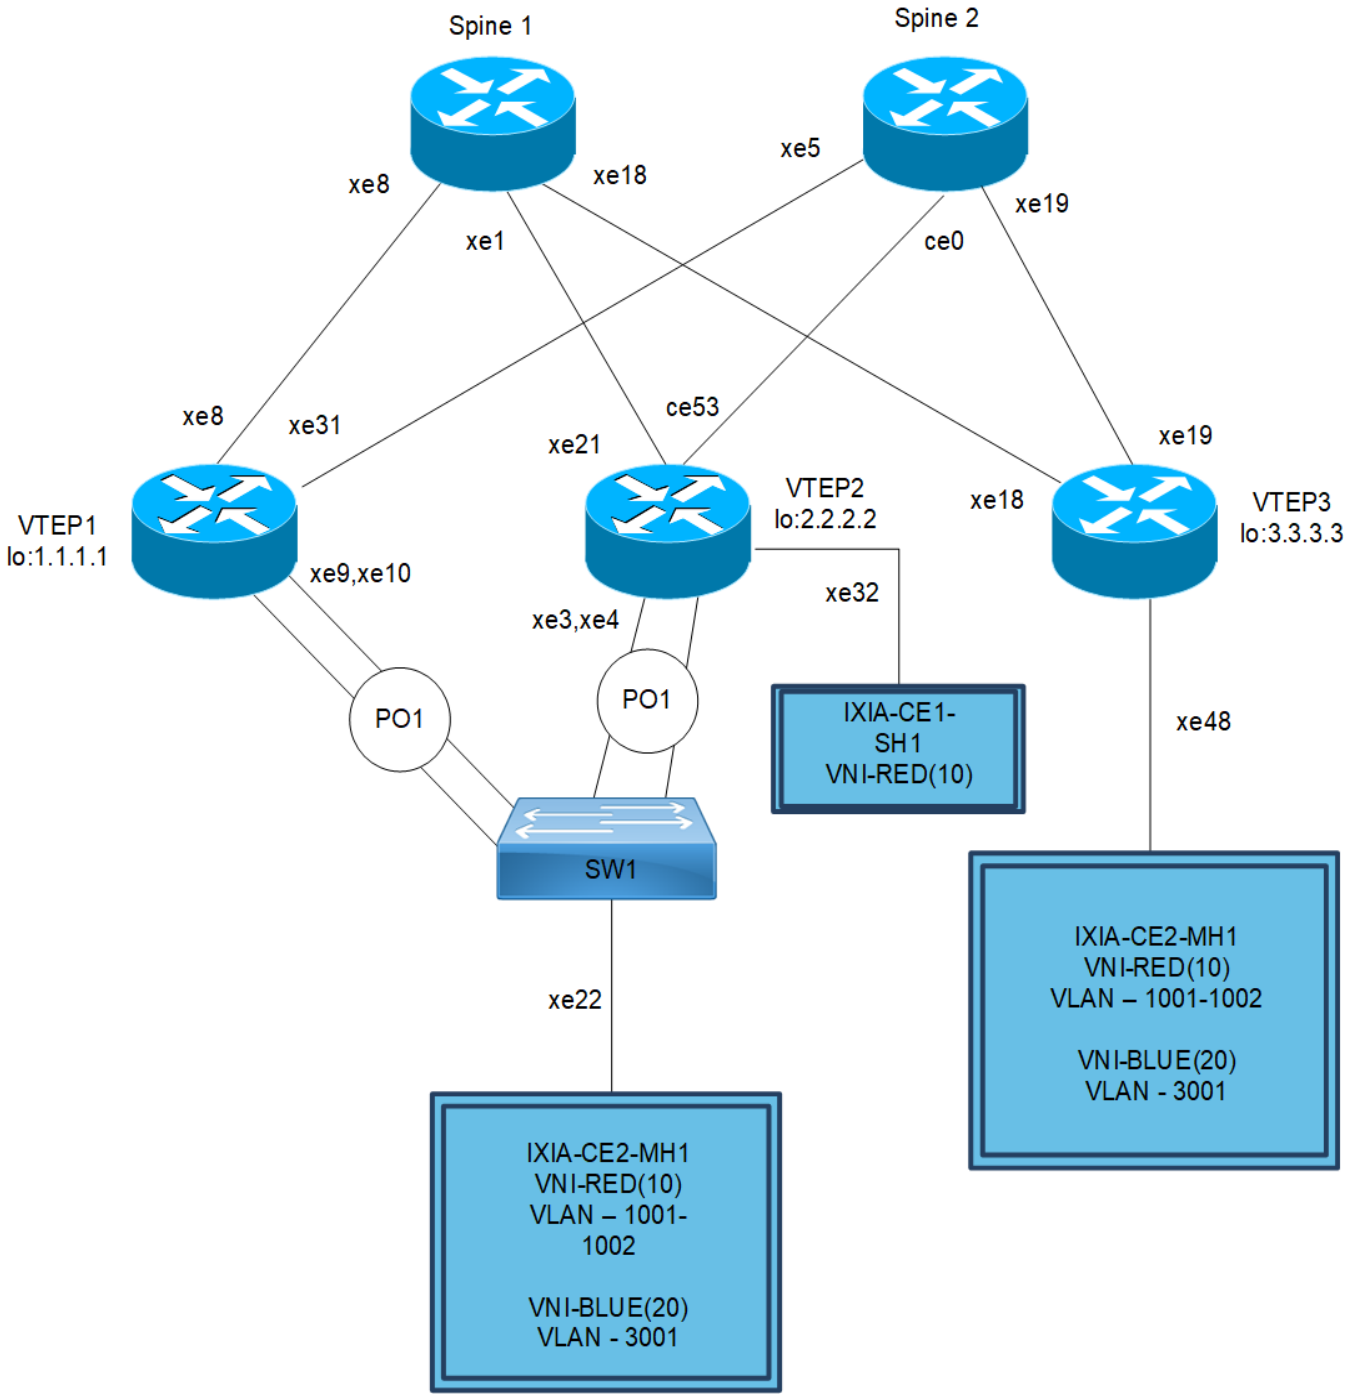

Figure 2-6: VxLAN-Multihoming

Note: Enable VxLAN MUTIHOMING before executing any configurations.

VxLAN-EVPN 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 or Static LAG 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)

Hardware Profile and Generic Configuration

`#configure terminal`	Enter Configure mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware-profile filter for VxLAN multi-homing.
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for VxLAN access-if port counters

(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 60	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)#interface po1.1001 switchport	Create L2 subinterface po1.1001
(config-if)#encapsulation dot1q 1001	Configure encapsulation dot1q with vlan-id 1001
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 (VNI-RED) to access interface po1.1001
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#interface po1.1002 switchport	Create L2 subinterface po1.1002
(config-if)#encapsulation dot1q 1002	Configure encapsulation dot1q with vlan-id 1002
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 (VNI-RED) to access interface po1.1002
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#interface po1.3001 switchport	Create L2 subinterface po1.3001
(config-if)#encapsulation dot1q 3001	Configure encapsulation dot1q with vlan-id 3001
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 20	Map vnid 20 (VNI-RED) to access interface po1.3001
(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 po1 (MH1). And it has xe32 as single home access-if port (SH2)

Hardware Profile and Generic Configuration

`#configure terminal`	Enter Configure mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware-profile filter for VxLAN multi-homing.
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for VxLAN access-if port counters
`(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 60	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)#interface po1.1001 switchport	Create L2 subinterface po1.1001
(config-if)#encapsulation dot1q 1001	Configure encapsulation dot1q with vlan-id 1001
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 (VNI-RED) to access interface po1.1001

(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#interface po1.1002 switchport	Create L2 subinterface po1.1002
(config-if)#encapsulation dot1q 1002	Configure encapsulation dot1q with vlan-id 1002
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 (VNI-RED) to access interface po1.1002
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#interface po1.3001 switchport	Create L2 subinterface po1.3001
(config-if)#encapsulation dot1q 3001	Configure encapsulation dot1q with vlan-id 3001
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 20 (VNI-RED) to access interface po1.3001
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#interface xe32 switchport	Create L2 subinterface xe32
(config-if)#encapsulation dot1q	Configure encapsulation dot1q
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 (VNI-RED) to access interface xe32
(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

VTEP3

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

Hardware Profile and Generic Configuration

#configure terminal	Enter Configure mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#hardware-profile filter vxlan-mh enable	Enable hardware-profile filter for VxLAN multi-homing.
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(config)#evpn vxlan multihoming enable	Enable Multihoming, save configs and reboot the board for multihoming to be effective
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for VxLAN access-if port counters
(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)#interface xe48.1001 switchport	Create L2 subinterface xe48.1001
(config-if)#encapsulation dot1q	Configure encapsulation dot1q with vlan-id 1001
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 (VNI-RED) to access interface xe48.1001
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#interface xe48.1002 switchport	Create L2 subinterface xe48.1002
(config-if)#encapsulation dot1q	Configure encapsulation dot1q with vlan-id 1002
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 (VNI-RED) to access interface xe48.1002
(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode
(config)#interface xe48.3001 switchport	Create L2 subinterface xe48.3001
(config-if)#encapsulation dot1q	Configure encapsulation dot1q with vlan-id 3001
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 20	Map vnid 20 (VNI-RED) to access interface xe48.3001
(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-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 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-if)#exit	Exit Interface mode and return to Configure 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

(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-router)#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-router)#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 Layer 2 port that is part of regular Layer 2 bridge (RSTP/MSTP/STP) and a VxLAN bridge at the same time. A switch port cannot be a part of VxLAN domain and Layer 2 Bridge if the VxLAN port is created with “port only” mapping. 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

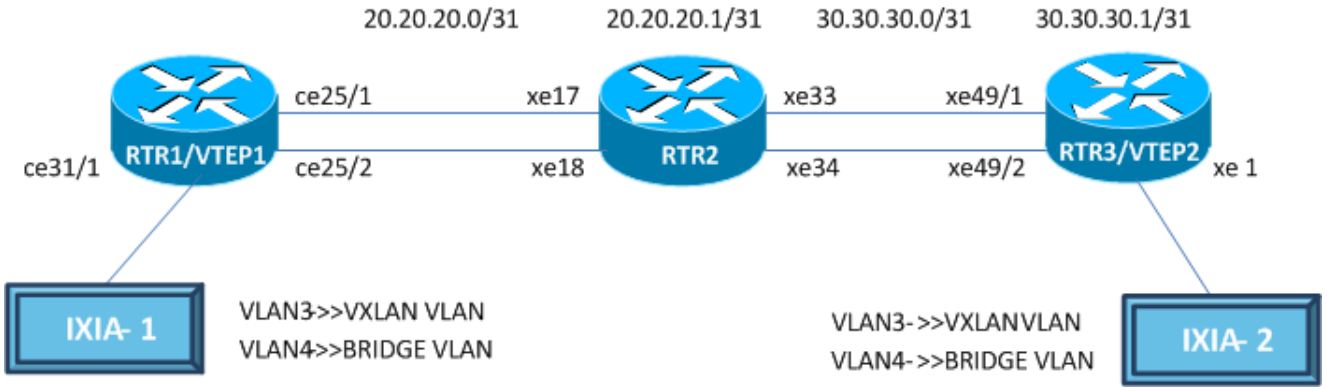

RTR1/VTEP1

`#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 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)#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-af)#exit-address-family	Exit l2vpn address family mode.
(config-router)#exit	Exit BGP router mode.
(config)# hardware-profile filter vxlan enable	Enable hardware profile for vxlan
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for VxLAN access-if port counters
(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)#interface ce31/1 3 switchport	Create L2 subinterface ce31/1 3
(config-if)#encapsulation dot1q 1 3	Configure encapsulation dot1q with vlan-id 1 3

(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 16777215	Map vnid 16777215 to access interface ce31/1 3
(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)# hardware-profile filter vxlan enable	Enable hardware profile for vxlan
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for VxLAN access-if port counters
(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)#interface xe1 3 switchport	Create L2 subinterface xe1 3
(config-if)#encapsulation dot1q 3	Configure encapsulation dot1q with vlan-id 3
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 16777215	Map vnid 16777215 to access interface xe1 3
(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

#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
```

VTEP2

```
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
 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 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 Configuration

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

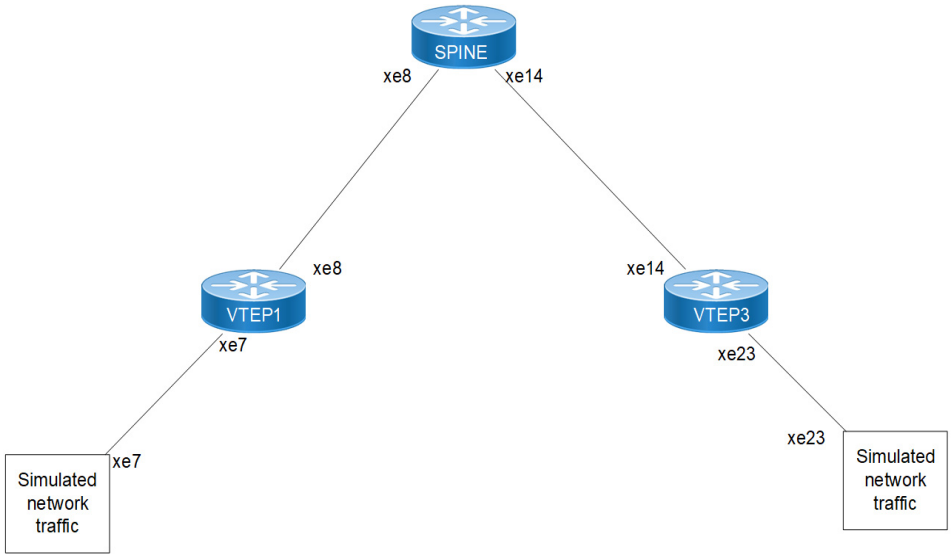

Figure 4-7: VxLAN Trunk Access Port Configuration topology

Base Configuration - L2 VxLAN

VTEP1

Hardware profile and generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for vxlan access-if port counters
`(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)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for vxlan access-if port counters
(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)#interface xe7 switchport	Create L2 subinterface xe7
(config-if)#encapsulation dot1q	Configure encapsulation dot1q
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 100	Map vnid 100 to access interface xe7
(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)#interface xe23 switchport	Create L2 subinterface xe23
(config-if)#encapsulation dot1q	Configure encapsulation dot1q
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 100	Map vnid 100 to access interface xe23
(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 an access port with a specific vlan range configured, all the traffic in that specific range are accepted and forwarded.

VTEP1

`(config)#interface xe7 switchport`	Create L2 subinterface xe7
`(config-if)#encapsulation dot1q`	Configure encapsulation dot1q
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 100`	Map vnid 100 to access interface xe7
`(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)#interface xe23 switchport`	Create L2 subinterface xe23
`(config-if)#encapsulation dot1q`	Configure encapsulation dot1q
`(config-if)#access-if-evpn`	Enter access-if-evpn mode

(config-acc-if-evpn)#map vpn-id 100	Map vnid 100 to access interface xe23
(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 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 Customer Edge (CE) to Provider Edge (PE) models.

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 CE to PE).
- 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.

```
hardware-profile filter vxlan enable
```

Topology

The EVPN MAC Limit feature can be deployed in any standard EVPN-VxLAN topology with CE-PE connectivity.

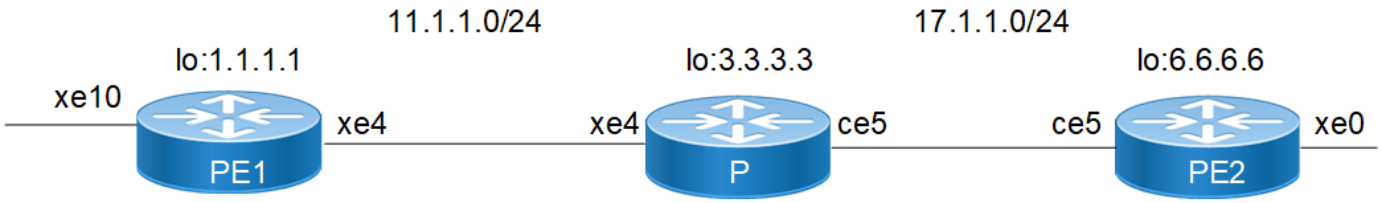

Figure 5-8: VXLAN-EVPN MAC limit enforcement

Configuration

The following configuration steps demonstrate how to enable and configure MAC limit enforcement on PE1 for EVPN-VxLAN.

EVPN prerequisite configurations:

PE1

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.

```
nvo vxlan id 50 ingress-replication
 vxlan host-reachability-protocol evpn-bgp vxlan_l2_elan_sh
```

5. Configure the access interface and map it to the specified VPN ID.

```
interface xe10.50 switchport
 access-if-evpn
  map vpn-id 50
```

PE2

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 ID.

```
nvo vxlan id 50 ingress-replication
 vxlan host-reachability-protocol evpn-bgp vxlan_l2_elan_sh2
```

5. Configure the access interface and map it to the VPN ID.

```
interface xe0.50 switchport
 access-if-evpn
  map vpn-id 50
```

EVPN-MAC-Limit profile configuration:

PE1

Configure a MAC limit profile with the learning-limit parameter using default values for high watermark, low watermark, and action. Attach the profile to the access circuit (AC).

```
mac-limit-profile SH1
learning-limit 10
interface xe10.50 switchport
 access-if-evpn
  learning limit SH1

PE1#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
PE1#
```

Configuration snapshot:

PE1

```
hostname PE1
```

```
!
nvo vxlan enable
!
ip vrf management
!
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
 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 xe4
 description network_to_spine1
 load-interval 30
 ip address 11.1.1.1/24
!
interface xe10
 speed 10g
 switchport
 load-interval 30
!
interface xe10.50 switchport
 encapsulation dot1q 50
 rewrite pop
 access-if-evpn
  map vpn-id 50
!
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
 network 12.1.1.0/24 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
 !
```

P:

```
hostname P
!
interface ce5
 load-interval 30
 ip address 17.1.1.2/24
!
interface xe4
 load-interval 30
 ip address 11.1.1.2/24
!
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
!
```

PE2

```
hardware-profile filter vxlan enable
hardware-profile port-config mode2
!
hostname PE2
!
nvo vxlan enable
!
ip vrf management
!
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
 vxlan host-reachability-protocol evpn-bgp vxlan_l2_elan_sh2
!
interface ce5
 load-interval 30
 ip address 17.1.1.1/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 6.6.6.6/32 secondary
 ipv6 address ::1/128
```

```
!
interface xe0
 speed 10g
 switchport
 load-interval 30
!
interface xe0.50 switchport
 encapsulation dot1q 50
 rewrite pop
 access-if-evpn
  map vpn-id 50
!
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
```

Validation

To verify MAC limit enforcement, send traffic with varying numbers of source MAC addresses and observe the system behavior.

- Send Traffic with 10 Source MACs.
- Generate traffic with 10 unique source MACs on xe10.50.
- When the MAC count reaches the high watermark threshold (default 90% of 10 = 9 MACs), a syslog message is generated.

```
PE1# 2025 Jun 07 03:13:47.291 : PE1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac limit for AC xe10.50, 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.

```
PE1#  2025 Jun 07 03:15:40.963 : PE1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac limit for AC xe10.50, Exceeded MAC count 11 with learning limit of 10
PE1#
```

- Now stop the traffic and check for Low water logs during un learning.

```
PE1#2025 Jun 07 03:16:42.814 : PE1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac limit for AC xe10.50, 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
!
interface xe10.50 switchport
 access-if-evpn
  learning limit SH1
PE1# 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
PE1#
```

- Send Traffic with 10 Source MACs.
- Generate traffic with 10 unique source MACs on xe10.50.
- When the MAC count reaches the high watermark threshold (default 90% of 10 = 9 MACs), a syslog message is generated.

```
PE1#2025 Jun 07 03:19:23.250 : PE1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac
limit for AC xe10.50, High threshold MAC count 8 with high watermark of 8
PE1#
```

- 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 07 03:21:24.749 : PE1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac limit
for AC xe10.50, Exceeded MAC count 11 with learning limit of 10


2025 Jun 07 03:21:24.751 : PE1 : NSM : CRITI : [IFMGR_IF_DOWN_2]: Interface xe10.50
changed state to down


2025 Jun 07 03:21:24.762 : PE1 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID 50: Mac limit
for AC xe10.50, Low threshold MAC count 6 with low watermark of 7


PE1#show nvo vxlan access-if brief

                     Inner                        Admin   Link
Interface  Vlan      vlan  Ifindex    Vnid        status  status
-------------------------------------------------------------
xe10.50    ---       ---   0x9c98032  50          down    down(ED)

Total number of entries are 1
```

Note: Refer sub-interface configuration for VLAN information.

```
PE1#


PE1#show interface brief xe10.50
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
-------------------------------------------------------------------------------
-----------------
Interface           Type             Status  Reason  Speed
-------------------------------------------------------------------------------
-----------------
xe10.50             SUBINTERFACE     down    ED      N/A
PE1#show interface brief xe10

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 #
-------------------------------------------------------------------------------
-----------------
xe10           ETH         --    --       up      none    10g   --       No       No
PE1#
```

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 instance with a mac-limit-profile

```
Configure an EVPN-VXLAN EVI with a mac-limit-profile
PE1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
PE1(config)#mac-limit-profile SH1
PE1(config-mac-limit-profile)#learning-limit 10
PE1(config-mac-limit-profile)#commit
PE1(config-mac-limit-profile)#

PE1#config t
Enter configuration commands, one per line. End with CNTL/Z.
PE1(config)#
PE1(config)#nvo vxlan id 50 ingress-replication
PE1(config-nvo)#learning limit SH1
PE1(config-nvo)#commit
PE1(config-nvo)#end
PE1#
PE1#
PE1#config t
Enter configuration commands, one per line. End with CNTL/Z.
PE1(config)#
PE1(config)#nvo vxlan id 50 ingress-replication
PE1(config-nvo)#no learning limit
PE1(config-nvo)#commit
PE1(config-nvo)#end
PE1#
```

```
onfigure an EVPN-VxLAN access interface with a mac-limit-profile
PE1#config t
Enter configuration commands, one per line. End with CNTL/Z.
PE1(config)#mac-limit-profile SH1
PE1(config-mac-limit-profile)#learning-limit 10
PE1(config-mac-limit-profile)#action log-errdisable
PE1(config-mac-limit-profile)#high-watermark 80
PE1(config-mac-limit-profile)#low-watermark 70
PE1(config-mac-limit-profile)#
PE1(config-mac-limit-profile)#commit

PE1# config t
Enter configuration commands, one per line. End with CNTL/Z.
PE1(config)#interface xe10.50 switchport
PE1(config-if)# access-if-evpn
PE1(config-acc-if-evpn)#learning limit SH1
PE1(config-acc-if-evpn)#commit
PE1(config-acc-if-evpn)#end
PE1#

PE1# config t
Enter configuration commands, one per line. End with CNTL/Z.
PE1(config)#interface xe10.50 switchport
PE1(config-if)# access-if-evpn
PE1(config-acc-if-evpn)#no learning limit
PE1(config-acc-if-evpn)#commit
PE1(config-acc-if-evpn)#end
PE1#
```

VxLAN - EVPN with Integrated Routing and Bridging Deployment Mode

CHAPTER 1 VxLAN-EVPN with IRB

Overview

Integrated Routing and Bridging (IRB) is an EVPN-based 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 1-9.

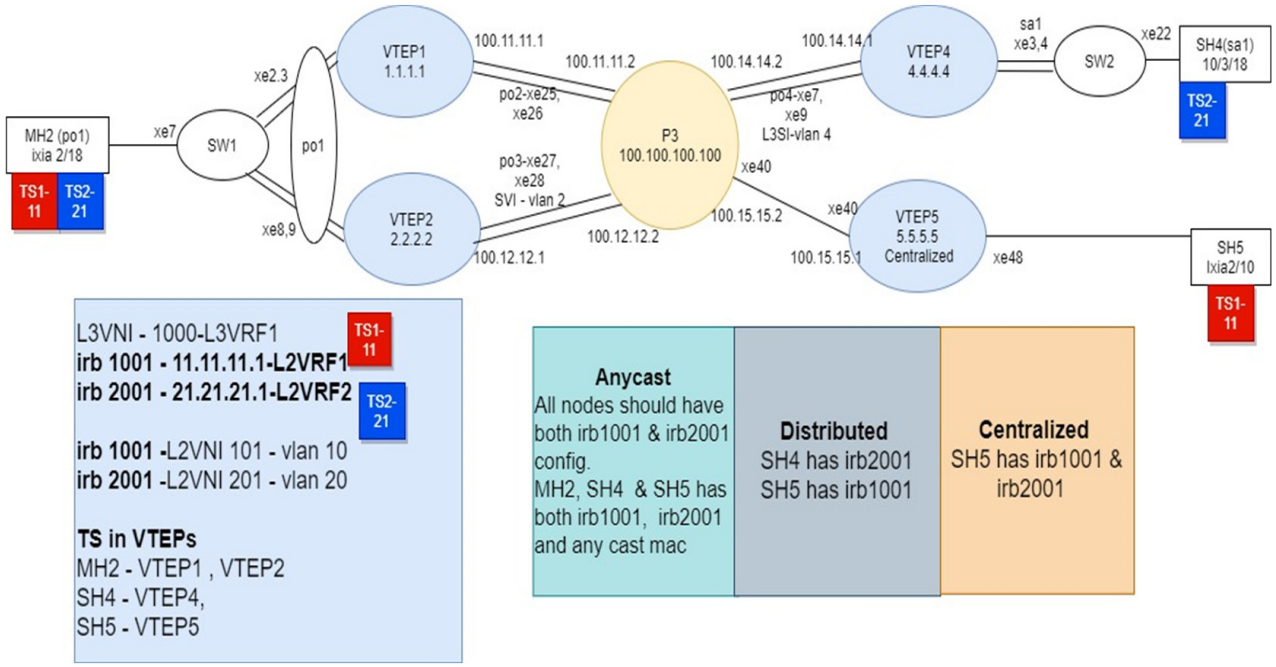

Figure 1-9: VxLAN EVPN IRB

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).

Hardware profile and generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware-profile filter for VxLAN multi-homing.
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for VxLAN access-if port counters
`(config)#qos enable`	Enabling QoS
`(config)#commit`	Commit the candidate configuration to running configuration

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

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

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

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

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)#commit	Commit the candidate configuration to running configuration
(config)#interface po1.10 switchport	Create L2 subinterface po1.10
(config-if)#encapsulation dot1q 10	Configure encapsulation dot1q with vlan id 101
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 101	Map vnid 101 to access interface po1.10
(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)#interface po1.20 switchport	Create L2 subinterface po1.20
(config-if)#encapsulation dot1q 20	Configure encapsulation dot1q with vlan-id 20
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 201	Map vnid 201 to access interface po1.20
(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

VTEP2

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

Hardware profile and generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware-profile filter for VxLAN multi-homing.
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for VxLAN access-if port counters
`(config)#qos enable`	Enabling QoS
`(config)#commit`	Commit the candidate configuration to running configuration

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)#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 100.12.12.1/24	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

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.
(config)#commit	Commit the candidate configuration to running configuration

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

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

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)#commit	Commit the candidate configuration to running configuration
(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)#interface po1.10 switchport	Create L2 subinterface po1.10
(config-if)#encapsulation dot1q 10	Configure encapsulation dot1q with vlan id 101
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 101	Map vnid 101 to access interface po1.10
(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)#interface po1.20 switchport	Create L2 subinterface po1.20
(config-if)#encapsulation dot1q 20	Configure encapsulation dot1q with vlan id 201
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 201	Map vnid 201 to access interface po1.20
(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)#commit	Commit the candidate configuration to running configuration

VTEP4

Single Home -SH5.

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for vxlan access-if port counters
(config)#qos enable	Enabling qos
(config)#commit	Commit the candidate configuration to running configuration

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 po4.4	Enter L3SI po4.4
(config-if)# encapsulation dot1q 4	Encapsulation as dot1q with VLAN 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

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

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

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

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)#commit	Commit the candidate configuration to running configuration
(config)#interface sa1.20 switchport	Create L2 subinterface sa1.20
(config-if)#encapsulation dot1q 20	Configure encapsulation dot1q with vlan id 201
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 201	Map vnid 201 to access interface sa1.20
(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

VTEP5

Single Home -SH3

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for vxlan access-if port counters
(config)#qos enable	Enabling qos
(config)#commit	Commit the candidate configuration to 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 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

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

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

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

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)#interface xe48.10 switchport	Create L2 subinterface xe48.10
(config-if)#encapsulation dot1q 10	Configure encapsulation dot1q with vlan id 101
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 101	Map vnid 101 to access interface xe48.10
(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)#commit	Commit the candidate configuration to running configuration

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

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

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

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 po4.4	Enter L3SI po4.4
(config-if)# encapsulation dot1q 4	Encapsulation as dot1q with VLAN 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

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

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
`(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
`(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)exit`	Exit from interface config mode
`(config)#commit`	Commit the candidate configuration to running configuration
`(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)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-router)#exit`	Exit from bgp router configuration mode
`(config)#commit`	Commit the candidate configuration to running configuration
`(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

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
(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 candidate configuration to running configuration
(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
(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

VTEP2

(config)#nvo vxlan irb	Enable VxLAN IRB
(config)#commit	Commit the candidate configuration to running configuration
(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 candidate configuration to running configuration
(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
(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

VTEP4

(config)#nvo vxlan irb	Enable VxLAN IRB
(config)#commit	Commit the candidate configuration to running configuration
(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)#commit	Commit the candidate configuration to running configuration
(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
(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

VTEP5

(config)#nvo vxlan irb	Enable VxLAN IRB
(config)#commit	Commit the candidate configuration to running configuration
(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)#commit	Commit the candidate configuration to running configuration
(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
(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)#commit`	Commit the candidate configuration to running configuration
`(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

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]:[0]:[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
`(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
(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

VTEP5

Unconfigure vnid 201 from nvo vxlan.

(config)#nvo vxlan irb	Enable VxLAN IRB
(config)#commit	Commit the candidate configuration to running configuration
(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)#commit	Commit the candidate configuration to running configuration
(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

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.

`(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 form address-family
`(config-router-af)#commit`	Commit the transaction

VTEP1 IRB configuration

`(config)#nvo vxlan irb`	Enable VxLAN IRB
`(config)#commit`	Commit the candidate configuration to running configuration
`(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 candidate configuration to running configuration
`(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 candidate configuration to running configuration
`(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 candidate configuration to running configuration

VTEP2

Configure `max-path ibgp 2` on VTEP1 under BGP IPv4 VRF address family.

`(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 form address-family
`(config-router-af)#commit`	Commit the transaction

VTEP2 IRB configuration

`(config)#nvo vxlan irb`	Enable VxLAN irb
`(config)#commit`	Commit the candidate configuration to running configuration
`(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 candidate configuration to running configuration
`(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 candidate configuration to running configuration
(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 candidate configuration to running configuration

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.

(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#hardware-profile filter vxlan-mh enable	Enable hardware-profile filter for VxLAN multi-homing.
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(config)#commit	Commit the transaction

(config)#evpn vxlan multihoming enable	Enable Multihoming, save configs and reboot the board for multihoming to be effective
(config)#commit	Commit the transaction
(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 trasaction
(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 form address-family
(config-router-af)#commit	Commit the trasaction
(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)#commit	Commit the trasaction

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 load shared 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, tunvxlan1001), 00:01:26
                           [200/0] via 2.2.2.2 (recursive is directly connected, tunvxlan1001), 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 VxLAN-EVPN 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 2-10

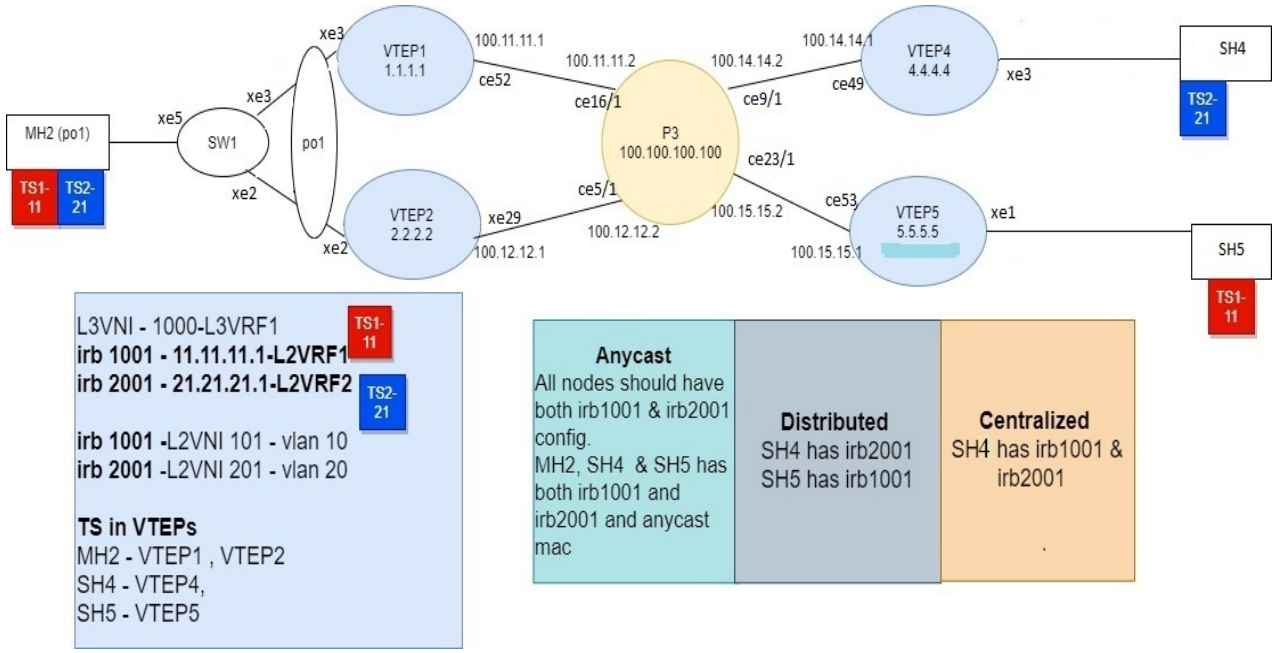

Figure 2-10: VxLAN_EVPN_IRB

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).

Hardware profile and generic configuration:

`#configure terminal`	Enter Configure mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware-profile filter for VxLAN multi-homing.
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for vxlan access-if port counters
`(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)#interface po1.10 switchport`	Create L2 subinterface po1.10
`(config-if)#encapsulation dot1q 10`	Configure encapsulation dot1q with vlan id 101
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 101`	Map vnid 101 to access interface po1.10
`(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)#interface po1.20 switchport`	Create L2 subinterface po1.20
`(config-if)#encapsulation dot1q 20`	Configure encapsulation dot1q with vlan id 201
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 201`	Map vnid 201 to access interface po1.20
`(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)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware-profile filter for VxLAN multi-homing.
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#evpn vxlan multihoming enable`	Enable Multihoming, save configs and reboot the board for multihoming to be effective
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for vxlan access-if port counters
`(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)#interface po1.10 switchport`	Create L2 subinterface po1.10
`(config-if)#encapsulation dot1q 10`	Configure encapsulation dot1q with vlan id 101
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 101`	Map vnid 101 to access interface po1.10
`(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)#interface po1.20 switchport`	Create L2 subinterface po1.20
`(config-if)#encapsulation dot1q 20`	Configure encapsulation dot1q with vlan id 201
`(config-if)#access-if-evpn`	Enter access-if-evpn mode
`(config-acc-if-evpn)#map vpn-id 201`	Map vnid 201 to access interface po1.20
`(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)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware profile mh in SH VTEP also
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for vxlan access-if port counters
`(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)#interface xe3.20 switchport	Create L2 subinterface xe3.20
(config-if)#encapsulation dot1q 20	Configure encapsulation dot1q with vlan id 201
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 201	Map vnid 201 to access interface xe3.20
(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)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#hardware-profile filter vxlan-mh enable	Enable hardware profile mh in SH VTEP also
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for vxlan access-if port counters
(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)#interface xe1.10 switchport	Create L2 subinterface xe1.10
(config-if)#encapsulation dot1q 10	Configure encapsulation dot1q with vlan id 101
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 101	Map vnid 101 to access interface xe1.10
(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 and dscp-encap qos profile should be applied on vxlan tunnel egress.

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)#qos profile dscp-encap DSCP_ENCAP	Configure the QoS profile to remark the overlay dscp value of the traffic in vxlan tunnel egress.
(config-egress-dscp-encap-map)#l3 dscp 20 dscpEncap 56	Egress remarking of the customer dscp packet to overlay dscp. Here classification at the egress vtep is based on customer dscp value
(config-egress-dscp-encap-map)#exit	Exit from qos profile 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)#nvo vxlan tunnel qos-map-mode cos-dscp egress DSCP_ENCAP	Map the qos profile in vxlan tunnel egress
(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
VTEP5(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
!
qos profile dscp-encap DSCP_ENCAP
 l3 dscp 20 dscpEncap 56
!
!
!
!
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 tunnel qos-map-mode cos-dscp egress DSCP_ENCAP
!
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 DSCP_ENCAP
profile name: DSCP_ENCAP
profile type: dscp-encap
profile attached to 1 instances
configured mapping:
 l3 dscp 20 dscpEncap 56
Detailed mapping:
 L3 DSCP to DSCP-ENCAP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  INPUT | OUTPUT   |  INPUT | OUTPUT  |  INPUT | OUTPUT  |  INPUT | OUTPUT
```

```
 -------|--------  | -------|-------- | -------|-------- | -------|--------
  DSCP  | DSCP     |  DSCP  | DSCP    |  DSCP  | DSCP    |  DSCP  | DSCP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  0        0       |  16       16       |  32       32       |  48       48
  1        1       |  17       17       |  33       33       |  49       49
  2        2       |  18       18       |  34       34       |  50       50
  3        3       |  19       19       |  35       35       |  51       51
  4        4       |  20       56       |  36       36       |  52       52
  5        5       |  21       21       |  37       37       |  53       53
  6        6       |  22       22       |  38       38       |  54       54
  7        7       |  23       23       |  39       39       |  55       55
  8        8       |  24       24       |  40       40       |  56       56
  9        9       |  25       25       |  41       41       |  57       57
  10       10       |  26       26       |  42       42       |  58       58
  11       11       |  27       27       |  43       43       |  59       59
  12       12       |  28       28       |  44       44       |  60       60
  13       13       |  29       29       |  45       45       |  61       61
  14       14       |  30       30       |  46       46       |  62       62
  15       15       |  31       31       |  47       47       |  63       63

 L2 Queue + Color to DSCP-ENCAP
 ---------------+--------| ----------------+--------| --------------+--------
      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     8      |  1      yellow    8      |  1      red       8
  2      green     16      |  2      yellow    16      |  2      red       16
  3      green     24      |  3      yellow    24      |  3      red       24
  4      green     32      |  4      yellow    32      |  4      red       32
  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

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

profile name: default
 profile type: dscp-to-dscp (Egress)
 Status: Inactive
mapping:
 ---------------+---------- | ---------------+---------- | ---------------+----------
       INPUT      |  OUTPUT   |       INPUT      |  OUTPUT   |       INPUT      |  OUTPUT
 ---------------+---------- | ---------------+---------- | ---------------+----------
  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP
 ------+---------+---------- | ------+---------+---------- | ------+---------+----------
  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     6         |  6      yellow    6         |  6      red       6
  7      green     7         |  7      yellow    7         |  7      red       7
  8      green     8         |  8      yellow    8         |  8      red       8
  9      green     9         |  9      yellow    9         |  9      red       9
  10     green     10        |  10     yellow    12        |  10     red       14
  11     green     11        |  11     yellow    11        |  11     red       11
  12     green     12        |  12     yellow    12        |  12     red       14
  13     green     13        |  13     yellow    13        |  13     red       13
  14     green     14        |  14     yellow    14        |  14     red       14
  15     green     15        |  15     yellow    15        |  15     red       15
  16     green     16        |  16     yellow    16        |  16     red       16
  17     green     17        |  17     yellow    17        |  17     red       17
  18     green     18        |  18     yellow    20        |  18     red       22
  19     green     19        |  19     yellow    19        |  19     red       19
  20     green     20        |  20     yellow    20        |  20     red       22
  21     green     21        |  21     yellow    21        |  21     red       21
  22     green     22        |  22     yellow    22        |  22     red       22
  23     green     23        |  23     yellow    23        |  23     red       23
  24     green     24        |  24     yellow    24        |  24     red       24
  25     green     25        |  25     yellow    25        |  25     red       25
  26     green     26        |  26     yellow    28        |  26     red       30
  27     green     27        |  27     yellow    27        |  27     red       27
  28     green     28        |  28     yellow    28        |  28     red       30
  29     green     29        |  29     yellow    29        |  29     red       29
  30     green     30        |  30     yellow    30        |  30     red       30
  31     green     31        |  31     yellow    31        |  31     red       31
  32     green     32        |  32     yellow    32        |  32     red       32
  33     green     33        |  33     yellow    33        |  33     red       33
  34     green     34        |  34     yellow    36        |  34     red       38
  35     green     35        |  35     yellow    35        |  35     red       35
  36     green     36        |  36     yellow    36        |  36     red       38
  37     green     37        |  37     yellow    37        |  37     red       37
  38     green     38        |  38     yellow    38        |  38     red       38
  39     green     39        |  39     yellow    39        |  39     red       39
  40     green     40        |  40     yellow    40        |  40     red       40
  41     green     41        |  41     yellow    41        |  41     red       41
  42     green     42        |  42     yellow    42        |  42     red       42
  43     green     43        |  43     yellow    43        |  43     red       43
  44     green     44        |  44     yellow    44        |  44     red       44
  45     green     45        |  45     yellow    45        |  45     red       45
  46     green     46        |  46     yellow    46        |  46     red       46
  47     green     47        |  47     yellow    47        |  47     red       47
  48     green     48        |  48     yellow    48        |  48     red       48
  49     green     49        |  49     yellow    49        |  49     red       49
  50     green     50        |  50     yellow    50        |  50     red       50
  51     green     51        |  51     yellow    51        |  51     red       51
  52     green     52        |  52     yellow    52        |  52     red       52
  53     green     53        |  53     yellow    53        |  53     red       53
  54     green     54        |  54     yellow    54        |  54     red       54
  55     green     55        |  55     yellow    55        |  55     red       55
  56     green     56        |  56     yellow    56        |  56     red       56
  57     green     57        |  57     yellow    57        |  57     red       57
  58     green     58        |  58     yellow    58        |  58     red       58
```

```
  59     green     59        |  59     yellow    59        |  59     red       59
  60     green     60        |  60     yellow    60        |  60     red       60
  61     green     61        |  61     yellow    61        |  61     red       61
  62     green     62        |  62     yellow    62        |  62     red       62
  63     green     63        |  63     yellow    63        |  63     red       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)#qos profile dscp-encap DSCP_ENCAP	Configure the qos profile to map the overlay dscp value at tunnel egress
(config-egress-dscp-encap-map)#l3 dscp 20 dscpEncap 34	Egress remarking of the customer dscp packet to overlay dscp
(config-egress-dscp-encap-map)#exit	Exit from the 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)#nvo vxlan tunnel qos-map-mode cos-dscp egress DSCP_ENCAP	Map the qos profile at vxlan tunnel egress
(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)#qos profile dscp-encap DSCP_ENCAP	Configure the qos profile to map the overlay dscp value at tunnel egress
(config-egress-dscp-encap-map)#l3 dscp 20 dscpEncap 34	Egress remarking of the customer dscp packet to overlay dscp
(config-egress-dscp-encap-map)#exit	Exit from the 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)#nvo vxlan tunnel qos-map-mode cos-dscp egress DSCP_ENCAP	Map the qos profile at vxlan tunnel egress
(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
!
qos profile dscp-encap DSCP_ENCAP
 l3 dscp 20 dscpEncap 34
!
!
!
!
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 tunnel qos-map-mode cos-dscp egress DSCP_ENCAP
!
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 DSCP_ENCAP
profile name: DSCP_ENCAP
profile type: dscp-encap
profile attached to 1 instances
configured mapping:
 l3 dscp 20 dscpEncap 34
Detailed mapping:
 L3 DSCP to DSCP-ENCAP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  INPUT | OUTPUT   |  INPUT | OUTPUT  |  INPUT | OUTPUT  |  INPUT | OUTPUT
 -------|--------  | -------|-------- | -------|-------- | -------|--------
  DSCP  | DSCP     |  DSCP  | DSCP    |  DSCP  | DSCP    |  DSCP  | DSCP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  0        0       |  16       16       |  32       32       |  48       48
  1        1       |  17       17       |  33       33       |  49       49
  2        2       |  18       18       |  34       34       |  50       50
  3        3       |  19       19       |  35       35       |  51       51
  4        4       |  20       34       |  36       36       |  52       52
  5        5       |  21       21       |  37       37       |  53       53
  6        6       |  22       22       |  38       38       |  54       54
  7        7       |  23       23       |  39       39       |  55       55
  8        8       |  24       24       |  40       40       |  56       56
  9        9       |  25       25       |  41       41       |  57       57
  10       10       |  26       26       |  42       42       |  58       58
  11       11       |  27       27       |  43       43       |  59       59
  12       12       |  28       28       |  44       44       |  60       60
  13       13       |  29       29       |  45       45       |  61       61
  14       14       |  30       30       |  46       46       |  62       62
  15       15       |  31       31       |  47       47       |  63       63

 L2 Queue + Color to DSCP-ENCAP
 ---------------+--------| ----------------+--------| --------------+--------
      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     8     |  1      yellow    8     |  1      red       8
  2      green     16     |  2      yellow    16     |  2      red       16
  3      green     24     |  3      yellow    24     |  3      red       24
```

```
  4       green      32      |  4       yellow     32      |  4       red        32
  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

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

profile name: default
 profile type: dscp-to-dscp (Egress)
 Status: Inactive
```

```
mapping:
 ----------------+---------- | ----------------+---------- | ----------------+---------
-
      INPUT      |  OUTPUT   |      INPUT      |  OUTPUT   |      INPUT      |  OUTPUT
 ----------------+---------- | ----------------+---------- | ----------------+---------
-
   DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP
 ------+---------+---------- | ------+---------+---------- | ------+---------+---------
-
   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     6         |  6      yellow    6         |  6      red       6
   7      green     7         |  7      yellow    7         |  7      red       7
   8      green     8         |  8      yellow    8         |  8      red       8
   9      green     9         |  9      yellow    9         |  9      red       9
   10     green     10        |  10     yellow    12        |  10     red       14
   11     green     11        |  11     yellow    11        |  11     red       11
   12     green     12        |  12     yellow    12        |  12     red       14
   13     green     13        |  13     yellow    13        |  13     red       13
   14     green     14        |  14     yellow    14        |  14     red       14
   15     green     15        |  15     yellow    15        |  15     red       15
   16     green     16        |  16     yellow    16        |  16     red       16
   17     green     17        |  17     yellow    17        |  17     red       17
   18     green     18        |  18     yellow    20        |  18     red       22
   19     green     19        |  19     yellow    19        |  19     red       19
   20     green     20        |  20     yellow    20        |  20     red       22
   21     green     21        |  21     yellow    21        |  21     red       21
   22     green     22        |  22     yellow    22        |  22     red       22
   23     green     23        |  23     yellow    23        |  23     red       23
   24     green     24        |  24     yellow    24        |  24     red       24
   25     green     25        |  25     yellow    25        |  25     red       25
   26     green     26        |  26     yellow    28        |  26     red       30
   27     green     27        |  27     yellow    27        |  27     red       27
   28     green     28        |  28     yellow    28        |  28     red       30
   29     green     29        |  29     yellow    29        |  29     red       29
   30     green     30        |  30     yellow    30        |  30     red       30
   31     green     31        |  31     yellow    31        |  31     red       31
   32     green     32        |  32     yellow    32        |  32     red       32
   33     green     33        |  33     yellow    33        |  33     red       33
   34     green     34        |  34     yellow    36        |  34     red       38
   35     green     35        |  35     yellow    35        |  35     red       35
   36     green     36        |  36     yellow    36        |  36     red       38
   37     green     37        |  37     yellow    37        |  37     red       37
   38     green     38        |  38     yellow    38        |  38     red       38
   39     green     39        |  39     yellow    39        |  39     red       39
   40     green     40        |  40     yellow    40        |  40     red       40
   41     green     41        |  41     yellow    41        |  41     red       41
```

```
  42     green     42        |  42     yellow    42        |  42     red       42
  43     green     43        |  43     yellow    43        |  43     red       43
  44     green     44        |  44     yellow    44        |  44     red       44
  45     green     45        |  45     yellow    45        |  45     red       45
  46     green     46        |  46     yellow    46        |  46     red       46
  47     green     47        |  47     yellow    47        |  47     red       47
  48     green     48        |  48     yellow    48        |  48     red       48
  49     green     49        |  49     yellow    49        |  49     red       49
  50     green     50        |  50     yellow    50        |  50     red       50
  51     green     51        |  51     yellow    51        |  51     red       51
  52     green     52        |  52     yellow    52        |  52     red       52
  53     green     53        |  53     yellow    53        |  53     red       53
  54     green     54        |  54     yellow    54        |  54     red       54
  55     green     55        |  55     yellow    55        |  55     red       55
  56     green     56        |  56     yellow    56        |  56     red       56
  57     green     57        |  57     yellow    57        |  57     red       57
  58     green     58        |  58     yellow    58        |  58     red       58
  59     green     59        |  59     yellow    59        |  59     red       59
  60     green     60        |  60     yellow    60        |  60     red       60
  61     green     61        |  61     yellow    61        |  61     red       61
  62     green     62        |  62     yellow    62        |  62     red       62
  63     green     63        |  63     yellow    63        |  63     red       63

VTEP2

VTEP2#show running-config qos
qos enable
!
qos profile dscp-to-queue DSCP_QUEUE
 dscp 20 queue 4
!
qos profile dscp-encap DSCP_ENCAP
 l3 dscp 20 dscpEncap 34
!
!
!
!
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 tunnel qos-map-mode cos-dscp egress DSCP_ENCAP
!
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 DSCP_ENCAP
profile name: DSCP_ENCAP
profile type: dscp-encap
profile attached to 1 instances
configured mapping:
 l3 dscp 20 dscpEncap 34
Detailed mapping:
 L3 DSCP to DSCP-ENCAP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  INPUT | OUTPUT   |  INPUT | OUTPUT  |  INPUT | OUTPUT  |  INPUT | OUTPUT
 -------|--------  | -------|-------- | -------|-------- | -------|--------
  DSCP  | DSCP     |  DSCP  | DSCP    |  DSCP  | DSCP    |  DSCP  | DSCP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  0        0       |  16       16       |  32       32       |  48       48
  1        1       |  17       17       |  33       33       |  49       49
  2        2       |  18       18       |  34       34       |  50       50
  3        3       |  19       19       |  35       35       |  51       51
  4        4       |  20       34       |  36       36       |  52       52
  5        5       |  21       21       |  37       37       |  53       53
```

```
  6        6       |  22       22       |  38       38       |  54       54
  7        7       |  23       23       |  39       39       |  55       55
  8        8       |  24       24       |  40       40       |  56       56
  9        9       |  25       25       |  41       41       |  57       57
  10       10       |  26       26       |  42       42       |  58       58
  11       11       |  27       27       |  43       43       |  59       59
  12       12       |  28       28       |  44       44       |  60       60
  13       13       |  29       29       |  45       45       |  61       61
  14       14       |  30       30       |  46       46       |  62       62
  15       15       |  31       31       |  47       47       |  63       63

 L2 Queue + Color to DSCP-ENCAP
 ---------------+--------| ----------------+--------| --------------+--------
      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     8      |  1       yellow    8      |  1       red        8
  2       green     16      |  2       yellow    16      |  2       red        16
  3       green     24      |  3       yellow    24      |  3       red        24
  4       green     32      |  4       yellow    32      |  4       red        32
  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

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

profile name: default
 profile type: dscp-to-dscp (Egress)
 Status: Inactive
mapping:
 ----------------+---------- | ----------------+---------- | ----------------+---------
-
      INPUT      |  OUTPUT   |      INPUT      |  OUTPUT   |      INPUT      |  OUTPUT
 ----------------+---------- | ----------------+---------- | ----------------+---------
-
  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP
 ------+---------+---------- | ------+---------+---------- | ------+---------+---------
-
  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     6         |  6      yellow    6         |  6      red       6
  7      green     7         |  7      yellow    7         |  7      red       7
  8      green     8         |  8      yellow    8         |  8      red       8
  9      green     9         |  9      yellow    9         |  9      red       9
  10     green     10        |  10     yellow    12        |  10     red       14
  11     green     11        |  11     yellow    11        |  11     red       11
  12     green     12        |  12     yellow    12        |  12     red       14
  13     green     13        |  13     yellow    13        |  13     red       13
  14     green     14        |  14     yellow    14        |  14     red       14
  15     green     15        |  15     yellow    15        |  15     red       15
  16     green     16        |  16     yellow    16        |  16     red       16
  17     green     17        |  17     yellow    17        |  17     red       17
  18     green     18        |  18     yellow    20        |  18     red       22
  19     green     19        |  19     yellow    19        |  19     red       19
```

```
20      green      20          |  20      yellow     20          |  20      red        22
21      green      21          |  21      yellow     21          |  21      red        21
22      green      22          |  22      yellow     22          |  22      red        22
23      green      23          |  23      yellow     23          |  23      red        23
24      green      24          |  24      yellow     24          |  24      red        24
25      green      25          |  25      yellow     25          |  25      red        25
26      green      26          |  26      yellow     28          |  26      red        30
27      green      27          |  27      yellow     27          |  27      red        27
28      green      28          |  28      yellow     28          |  28      red        30
29      green      29          |  29      yellow     29          |  29      red        29
30      green      30          |  30      yellow     30          |  30      red        30
31      green      31          |  31      yellow     31          |  31      red        31
32      green      32          |  32      yellow     32          |  32      red        32
33      green      33          |  33      yellow     33          |  33      red        33
34      green      34          |  34      yellow     36          |  34      red        38
35      green      35          |  35      yellow     35          |  35      red        35
36      green      36          |  36      yellow     36          |  36      red        38
37      green      37          |  37      yellow     37          |  37      red        37
38      green      38          |  38      yellow     38          |  38      red        38
39      green      39          |  39      yellow     39          |  39      red        39
40      green      40          |  40      yellow     40          |  40      red        40
41      green      41          |  41      yellow     41          |  41      red        41
42      green      42          |  42      yellow     42          |  42      red        42
43      green      43          |  43      yellow     43          |  43      red        43
44      green      44          |  44      yellow     44          |  44      red        44
45      green      45          |  45      yellow     45          |  45      red        45
46      green      46          |  46      yellow     46          |  46      red        46
47      green      47          |  47      yellow     47          |  47      red        47
48      green      48          |  48      yellow     48          |  48      red        48
49      green      49          |  49      yellow     49          |  49      red        49
50      green      50          |  50      yellow     50          |  50      red        50
51      green      51          |  51      yellow     51          |  51      red        51
52      green      52          |  52      yellow     52          |  52      red        52
53      green      53          |  53      yellow     53          |  53      red        53
54      green      54          |  54      yellow     54          |  54      red        54
55      green      55          |  55      yellow     55          |  55      red        55
56      green      56          |  56      yellow     56          |  56      red        56
57      green      57          |  57      yellow     57          |  57      red        57
58      green      58          |  58      yellow     58          |  58      red        58
59      green      59          |  59      yellow     59          |  59      red        59
60      green      60          |  60      yellow     60          |  60      red        60
61      green      61          |  61      yellow     61          |  61      red        61
62      green      62          |  62      yellow     62          |  62      red        62
63      green      63          |  63      yellow     63          |  63      red        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)#qos profile dscp-to-dscp DSCP_DSCP	Configure qos profile for dscp value remark if required. Here classification at the egress vtep is based on customer dscp value
(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)#int irb2001	Enter IRB interface mode
(config-irb-if)#qos map-profile dscp-to-dscp DSCP_DSCP	Map Qos profile for dscp remark if required
(config-irb-if)#exit	Exit from interface mode
(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-ingress-dscp-map)#qos profile dscp-encap DSCP_ENCAP	Configure qos profile for overlay dscp remark in vxlan tunnel egress
(config-egress-dscp-encap-map)# l3 dscp 20 dscpEncap 56	Egress remarking of the customer dscp packet to overlay dscp
(config-egress-dscp-encap-map)#exit	Exit from qos profile config mode
(config)#nvo vxlan tunnel qos-map-mode cos-dscp egress DSCP_ENCAP	Map the qos profile in vxlan tunnel egress
(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
!
qos profile dscp-encap DSCP_ENCAP
 l3 dscp 20 dscpEncap 56
!
!
!
!
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 tunnel qos-map-mode cos-dscp egress DSCP_ENCAP
!
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 DSCP_ENCAP
profile name: DSCP_ENCAP
profile type: dscp-encap
profile attached to 1 instances
configured mapping:
 l3 dscp 20 dscpEncap 56
Detailed mapping:
 L3 DSCP to DSCP-ENCAP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  INPUT | OUTPUT   |  INPUT | OUTPUT  |  INPUT | OUTPUT  |  INPUT | OUTPUT
 -------|--------  | -------|-------- | -------|-------- | -------|--------
  DSCP  | DSCP     |  DSCP  | DSCP    |  DSCP  | DSCP    |  DSCP  | DSCP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  0        0       |  16       16       |  32       32       |  48       48
  1        1       |  17       17       |  33       33       |  49       49
  2        2       |  18       18       |  34       34       |  50       50
  3        3       |  19       19       |  35       35       |  51       51
  4        4       |  20       56       |  36       36       |  52       52
  5        5       |  21       21       |  37       37       |  53       53
  6        6       |  22       22       |  38       38       |  54       54
  7        7       |  23       23       |  39       39       |  55       55
  8        8       |  24       24       |  40       40       |  56       56
  9        9       |  25       25       |  41       41       |  57       57
  10       10       |  26       26       |  42       42       |  58       58
  11       11       |  27       27       |  43       43       |  59       59
  12       12       |  28       28       |  44       44       |  60       60
  13       13       |  29       29       |  45       45       |  61       61
  14       14       |  30       30       |  46       46       |  62       62
  15       15       |  31       31       |  47       47       |  63       63

 L2 Queue + Color to DSCP-ENCAP
 ---------------+--------| -----------------+--------| --------------+--------
      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     8     |  1      yellow    8     |  1      red       8
  2      green     16     |  2      yellow    16     |  2      red       16
  3      green     24     |  3      yellow    24     |  3      red       24
  4      green     32     |  4      yellow    32     |  4      red       32
  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

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

profile name: default
 profile type: dscp-to-dscp (Egress)
 Status: Inactive
mapping:
 ---------------+---------- | ---------------+---------- | ---------------+---------
-
      INPUT      |  OUTPUT   |      INPUT     |  OUTPUT   |      INPUT     |  OUTPUT
 ---------------+---------- | ---------------+---------- | ---------------+---------
-
  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP
 ------+---------+---------- | ------+---------+---------- | ------+---------+---------
-
  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	6	6	yellow	6	6	red	6
7	green	7	7	yellow	7	7	red	7
8	green	8	8	yellow	8	8	red	8
9	green	9	9	yellow	9	9	red	9
10	green	10	10	yellow	12	10	red	14
11	green	11	11	yellow	11	11	red	11
12	green	12	12	yellow	12	12	red	14
13	green	13	13	yellow	13	13	red	13
14	green	14	14	yellow	14	14	red	14
15	green	15	15	yellow	15	15	red	15
16	green	16	16	yellow	16	16	red	16
17	green	17	17	yellow	17	17	red	17
18	green	18	18	yellow	20	18	red	22
19	green	19	19	yellow	19	19	red	19
20	green	20	20	yellow	20	20	red	22
21	green	21	21	yellow	21	21	red	21
22	green	22	22	yellow	22	22	red	22
23	green	23	23	yellow	23	23	red	23
24	green	24	24	yellow	24	24	red	24
25	green	25	25	yellow	25	25	red	25
26	green	26	26	yellow	28	26	red	30
27	green	27	27	yellow	27	27	red	27
28	green	28	28	yellow	28	28	red	30
29	green	29	29	yellow	29	29	red	29
30	green	30	30	yellow	30	30	red	30
31	green	31	31	yellow	31	31	red	31
32	green	32	32	yellow	32	32	red	32
33	green	33	33	yellow	33	33	red	33
34	green	34	34	yellow	36	34	red	38
35	green	35	35	yellow	35	35	red	35
36	green	36	36	yellow	36	36	red	38
37	green	37	37	yellow	37	37	red	37
38	green	38	38	yellow	38	38	red	38
39	green	39	39	yellow	39	39	red	39
40	green	40	40	yellow	40	40	red	40
41	green	41	41	yellow	41	41	red	41
42	green	42	42	yellow	42	42	red	42
43	green	43	43	yellow	43	43	red	43
44	green	44	44	yellow	44	44	red	44
45	green	45	45	yellow	45	45	red	45
46	green	46	46	yellow	46	46	red	46
47	green	47	47	yellow	47	47	red	47
48	green	48	48	yellow	48	48	red	48
49	green	49	49	yellow	49	49	red	49
50	green	50	50	yellow	50	50	red	50
51	green	51	51	yellow	51	51	red	51
52	green	52	52	yellow	52	52	red	52

```
  53     green     53        |  53     yellow    53        |  53     red       53
  54     green     54        |  54     yellow    54        |  54     red       54
  55     green     55        |  55     yellow    55        |  55     red       55
  56     green     56        |  56     yellow    56        |  56     red       56
  57     green     57        |  57     yellow    57        |  57     red       57
  58     green     58        |  58     yellow    58        |  58     red       58
  59     green     59        |  59     yellow    59        |  59     red       59
  60     green     60        |  60     yellow    60        |  60     red       60
  61     green     61        |  61     yellow    61        |  61     red       61
  62     green     62        |  62     yellow    62        |  62     red       62
  63     green     63        |  63     yellow    63        |  63     red       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
!
qos profile dscp-to-dscp DSCP_DSCP
 dscp 20 color all dscp 32
!
!
!
!
interface irb2001
 qos map-profile dscp-to-dscp DSCP_DSCP
!
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 DSCP_DSCP
profile name: DSCP_DSCP
profile type: dscp-to-dscp
profile attached to 1 instances
configured mapping:
 dscp 20 color all dscp 32
Detailed mapping:
 ----------------+---------- | ----------------+---------- | ----------------+---------
-
      INPUT      |  OUTPUT   |      INPUT      |  OUTPUT   |      INPUT      |  OUTPUT
 ----------------+---------- | ----------------+---------- | ----------------+---------
-
  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP
 ------+---------+---------- | ------+---------+---------- | ------+---------+---------
-
```

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	6	6	yellow	6	6	red	6
7	green	7	7	yellow	7	7	red	7
8	green	8	8	yellow	8	8	red	8
9	green	9	9	yellow	9	9	red	9
10	green	10	10	yellow	12	10	red	14
11	green	11	11	yellow	11	11	red	11
12	green	12	12	yellow	12	12	red	14
13	green	13	13	yellow	13	13	red	13
14	green	14	14	yellow	14	14	red	14
15	green	15	15	yellow	15	15	red	15
16	green	16	16	yellow	16	16	red	16
17	green	17	17	yellow	17	17	red	17
18	green	18	18	yellow	20	18	red	22
19	green	19	19	yellow	19	19	red	19
20	green	32	20	yellow	32	20	red	32
21	green	21	21	yellow	21	21	red	21
22	green	22	22	yellow	22	22	red	22
23	green	23	23	yellow	23	23	red	23
24	green	24	24	yellow	24	24	red	24
25	green	25	25	yellow	25	25	red	25
26	green	26	26	yellow	28	26	red	30
27	green	27	27	yellow	27	27	red	27
28	green	28	28	yellow	28	28	red	30
29	green	29	29	yellow	29	29	red	29
30	green	30	30	yellow	30	30	red	30
31	green	31	31	yellow	31	31	red	31
32	green	32	32	yellow	32	32	red	32
33	green	33	33	yellow	33	33	red	33
34	green	34	34	yellow	36	34	red	38
35	green	35	35	yellow	35	35	red	35
36	green	36	36	yellow	36	36	red	38
37	green	37	37	yellow	37	37	red	37
38	green	38	38	yellow	38	38	red	38
39	green	39	39	yellow	39	39	red	39
40	green	40	40	yellow	40	40	red	40
41	green	41	41	yellow	41	41	red	41
42	green	42	42	yellow	42	42	red	42
43	green	43	43	yellow	43	43	red	43
44	green	44	44	yellow	44	44	red	44
45	green	45	45	yellow	45	45	red	45
46	green	46	46	yellow	46	46	red	46
47	green	47	47	yellow	47	47	red	47
48	green	48	48	yellow	48	48	red	48
49	green	49	49	yellow	49	49	red	49

```
  50     green     50        |  50     yellow    50        |  50     red       50
  51     green     51        |  51     yellow    51        |  51     red       51
  52     green     52        |  52     yellow    52        |  52     red       52
  53     green     53        |  53     yellow    53        |  53     red       53
  54     green     54        |  54     yellow    54        |  54     red       54
  55     green     55        |  55     yellow    55        |  55     red       55
  56     green     56        |  56     yellow    56        |  56     red       56
  57     green     57        |  57     yellow    57        |  57     red       57
  58     green     58        |  58     yellow    58        |  58     red       58
  59     green     59        |  59     yellow    59        |  59     red       59
  60     green     60        |  60     yellow    60        |  60     red       60
  61     green     61        |  61     yellow    61        |  61     red       61
  62     green     62        |  62     yellow    62        |  62     red       62
  63     green     63        |  63     yellow    63        |  63     red       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

profile name: DSCP_DSCP
 profile type: dscp-to-dscp (Egress)
 Status: Active
mapping:
 ----------------+---------- | ----------------+---------- | ----------------+---------
-
      INPUT      |  OUTPUT   |      INPUT      |  OUTPUT   |      INPUT      |  OUTPUT
 ----------------+---------- | ----------------+---------- | ----------------+---------
-
  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP  |  DSCP |  Color  | Out DSCP
 ------+---------+---------- | ------+---------+---------- | ------+---------+---------
-
  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     6         |  6      yellow    6         |  6      red       6
  7      green     7         |  7      yellow    7         |  7      red       7
  8      green     8         |  8      yellow    8         |  8      red       8
  9      green     9         |  9      yellow    9         |  9      red       9
  10     green     10        |  10     yellow    12        |  10     red       14
  11     green     11        |  11     yellow    11        |  11     red       11
  12     green     12        |  12     yellow    12        |  12     red       14
  13     green     13        |  13     yellow    13        |  13     red       13
  14     green     14        |  14     yellow    14        |  14     red       14
  15     green     15        |  15     yellow    15        |  15     red       15
  16     green     16        |  16     yellow    16        |  16     red       16
  17     green     17        |  17     yellow    17        |  17     red       17
  18     green     18        |  18     yellow    20        |  18     red       22
  19     green     19        |  19     yellow    19        |  19     red       19
  20     green     32        |  20     yellow    32        |  20     red       32
  21     green     21        |  21     yellow    21        |  21     red       21
  22     green     22        |  22     yellow    22        |  22     red       22
  23     green     23        |  23     yellow    23        |  23     red       23
  24     green     24        |  24     yellow    24        |  24     red       24
  25     green     25        |  25     yellow    25        |  25     red       25
  26     green     26        |  26     yellow    28        |  26     red       30
  27     green     27        |  27     yellow    27        |  27     red       27
  28     green     28        |  28     yellow    28        |  28     red       30
  29     green     29        |  29     yellow    29        |  29     red       29
  30     green     30        |  30     yellow    30        |  30     red       30
```

31	green	31	31	yellow	31	31	red	31
32	green	32	32	yellow	32	32	red	32
33	green	33	33	yellow	33	33	red	33
34	green	34	34	yellow	36	34	red	38
35	green	35	35	yellow	35	35	red	35
36	green	36	36	yellow	36	36	red	38
37	green	37	37	yellow	37	37	red	37
38	green	38	38	yellow	38	38	red	38
39	green	39	39	yellow	39	39	red	39
40	green	40	40	yellow	40	40	red	40
41	green	41	41	yellow	41	41	red	41
42	green	42	42	yellow	42	42	red	42
43	green	43	43	yellow	43	43	red	43
44	green	44	44	yellow	44	44	red	44
45	green	45	45	yellow	45	45	red	45
46	green	46	46	yellow	46	46	red	46
47	green	47	47	yellow	47	47	red	47
48	green	48	48	yellow	48	48	red	48
49	green	49	49	yellow	49	49	red	49
50	green	50	50	yellow	50	50	red	50
51	green	51	51	yellow	51	51	red	51
52	green	52	52	yellow	52	52	red	52
53	green	53	53	yellow	53	53	red	53
54	green	54	54	yellow	54	54	red	54
55	green	55	55	yellow	55	55	red	55
56	green	56	56	yellow	56	56	red	56
57	green	57	57	yellow	57	57	red	57
58	green	58	58	yellow	58	58	red	58
59	green	59	59	yellow	59	59	red	59
60	green	60	60	yellow	60	60	red	60
61	green	61	61	yellow	61	61	red	61
62	green	62	62	yellow	62	62	red	62
63	green	63	63	yellow	63	63	red	63

CHAPTER 3 VxLAN Eline xConnect Configuration

This chapter contains examples of VxLAN Eline xConnect using single-homed and multi-homed topologies.

Vxlan Eline xConnect is a mechanism for a point-to-point tunnel for data and control packets from one leaf to another. It helps in achieving pseudowire between two endpoints. Since there are only two endpoints in this, there is no need for MAC learning and hence the differentiation of Unicast/Broadcast-Unicast-Multicast traffic is not seen. Inner Tags are preserved and VxLAN encapsulated within the outer VNID which is specified as the Xconnect VNID.

Single-Homed VxLAN Eline xConnect

Topology

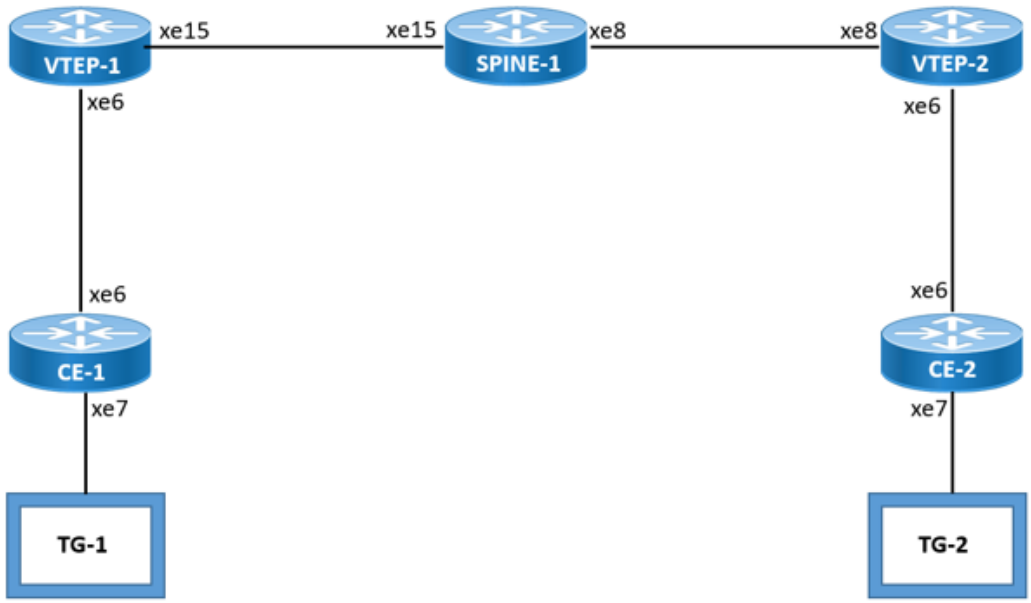

Figure 3-11: Single-homed VxLAN Eline xConnect topology

L2, IPv4, and IPv6 traffic streams are sent from TG-1 and TG-2.

CE1

`(config)#hostname CE1`	Configure host name
`(config)#bridge 1 protocol ieee vlan-bridge`	Configure bridge as IEEE VLAN bridge
`(config)#interface xe6`	Enter interface mode
`(config-if)#switchport`	Configure the interface as switchport
`(config-if)#exit`	Exit interface mode
`(config)#interface xe7`	Enter interface mode
`(config-if)#switchport`	Configure the interface as switchport

(config-if)#bridge-group 1	Associate bridge to an interface.
(config-if)#switchport mode trunk	Configure port as a trunk.
(config-if)#switchport trunk allowed vlan all	Allow all VLANs on the interface.
(config-if)#commit	Commit the candidate configuration to the running configuration.

VTEP-1

#con t	Enter configure mode
(config)#hostname VTEP-1	Configure host name
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN
(config)#hardware-profile statistics ingress-acl enable	Configure hardware profile statistics ingress-acl
(config)#nvo vxlan enable	Enable VxLAN globally on this vtep
(config)#mac vrf evpn_sh	Configure a new VRF
(config-vrf)#rd 1.1.1.1:1	Assign the Route Distinguisher value.
(config-vrf)#route-target both 10:1	Configure route target to import and export the routes
(config-vrf)#exit	Exit vrf mode
(config)#bridge 1 protocol ieee vlan-bridge	Configure bridge as IEEE VLAN bridge
(config)#interface xe6	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#exit	Exit interface mode
(config)#interface xe15	Enter interface mode
(config-if)#ip address 10.10.10.1/24	Configure IP address on the interface
(config-if)#exit	Exit interface mode
(config)#interface lo	Enter interface mode
(config-if)#ip address 1.1.1.1/32 secondary	Configure IP address on the interface
(config-if)#exit	Exit interface mode
(config)#router ospf 100	Enter router mode for OSPF
(config-router)#bfd all-interfaces	Enable BFD for all-interface into OSPF
(config-router)#network 1.1.1.1/32 area 0.0.0.0	Add 1.1.1.1 network into area 0
(config-router)#network 10.10.10.0/24 area 0.0.0.0	Add 10.10.10.0 network into area 0
(config)#router bgp 100	Enter BGP router mode
(config-router)#bgp router-id 1.1.1.1	Assign BGP router ID
(config-router)#neighbor 2.2.2.2 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 2.2.2.2 update-source lo	Specify the neighbor to use loopback address as source
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 2.2.2.2 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 vtep-ip-global 1.1.1.1	Configure Source vtep-ip-global configuration
(config)#nvo vxlan id 10 xconnect target-vxlan-id 20	add a tenant and the type of VPN. This creates an ELAN with source and target identifier for ELINE XConnect
(config-nvo)#vxlan host-reachability-protocol evpn-bgp evpn_sh	Set the host reachable protocol to Ethernet-
(config)#interface xe6.2 switchport	Create L2 subinterface xe6.2
(config-if)#encapsulation dot1q 2	Configure encapsulation dot1q with vlan id 10
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 to access interface xe6.2
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration.

SPINE-1

(config)#hostname SPINE-1	Configure hostname
(config)#hardware-profile statistics ingress-acl enable	Configure hardware profile statistics ingress-acl
(config)#interface xe8	Enter interface mode
(config-if)#ip address 20.20.20.1/24	Configure IP address on the interface
(config-if)#exit	Exit interface mode
(config)#interface xe15	Enter interface mode
(config-if)#ip address 10.10.10.2/24	Configure IP address on the interface
(config-if)#exit	Exit interface mode
(config)#router ospf 100	Enter OSPF router mode
(config-router)#bfd all-interfaces	Enable BFD for all-interface into OSPF
(config-router)#network 10.10.10.0/24 area 0.0.0.0	Add 10.10.10.0 network into area 0
(config-router)#network 20.20.20.0/24 area 0.0.0.0	Add 20.20.20.0 network into area 0
(config-router)#commit	Commit the candidate configuration to the running configuration.

VTEP-2

#con t	Enter configure mode
(config)#hostname VTEP-2	Configure hostname
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN
(config)#hardware-profile statistics ingress-acl enable	Configure hardware profile statistics ingress-acl
(config)#nvo vxlan enable	Enable VxLAN globally on this vtep
(config)#mac vrf evpn_sh	Configure a new VRF
(config-vrf)#rd 2.2.2.2:1	Assign the Route Distinguisher value.

(config-vrf)#route-target both 10:1	Configure route target to import and export the routes
(config-vrf)#exit	Exit vrf mode
(config)#bridge 1 protocol ieee vlan-bridge	Configure bridge as IEEE VLAN bridge
(config)#interface xe6	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#exit	Exit interface mode
(config)#interface xe8	Enter interface mode
(config-if)#ip address 20.20.20.2/24	Configure IP address on the interface
(config-if)#exit	Exit interface mode
(config)#interface lo	Enter interface mode
(config-if)#ip address 2.2.2.2/32 secondary	Configure IP address on the interface
(config-if)#exit	Exit interface mode
(config)#router ospf 100	Enter into router ospf mode
(config-router)#bfd all-interfaces	Enable BFD for all-interface into OSPF
(config-router)#network 2.2.2.2/32 area 0.0.0.0	Add lo ip address 2.2.2.2 as network into area 0
(config-router)#network 20.20.20.0/24 area 0.0.0.0	Add 20.20.20.0 network into area 0
(config-router)#exit	Exit router mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#bgp router-id 2.2.2.2	Assign BGP router ID
(config-router)#neighbor 1.1.1.1 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Specify the neighbor to use loopback address as source
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 1.1.1.1 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 vtep-ip-global 2.2.2.2	Configure Source vtep-ip-global configuration
(config)#nvo vxlan id 20 xconnect target-vxlan-id 10	add a tenant and the type of VPN. This creates an ELAN with source and target identifier for ELINE XConnect
(config-nvo)#vxlan host-reachability-protocol evpn-bgp evpn_sh	Set the host reachable protocol to Ethernet-
(config)#interface xe6.2 switchport	Create L2 subinterface xe6.2
(config-if)#encapsulation dot1q 2	Configure encapsulation dot1q with vlan id 20
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 20	Map vnid 20 to access interface xe6.2
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration.

CE2

(config)#hostname CE1	Configure host name
(config)#bridge 1 protocol ieee vlan-bridge	Configure bridge as IEEE VLAN bridge
(config)#interface xe6	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#exit	Exit interface mode
(config)#interface xe7	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#bridge-group 1	Associate bridge to an interface.
(config-if)#switchport mode trunk	Configure port as a trunk.
(config-if)#switchport trunk allowed vlan all	Allow all VLANs on the interface.
(config-if)#commit	Commit the candidate configuration to the running configuration.

Validation

```
VTEP-1#sh nvo vxlan xconnect
EVPN Xconnect Info
========================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                              Remote       Connection-Details
================================== ============
================================================================================
VPN-ID       EVI-Name      MTU   VPN-ID       Source       Destination
PE-IP           MTU   Type   NW-Status
================================== ============
================================================================================
10           ----          1500  20           xe6 2        --- Single Homed Port ---
2.2.2.2         1500  AC-NW  NW-SET

Total number of entries are 1
VTEP-1#sh nvo vxlan tunnel
VxLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================
1.1.1.1          2.2.2.2          Installed     00:02:01      00:02:01

Total number of entries are 1
VTEP-1#
```

```
VTEP-2#sh nvo vxlan xconnect
EVPN Xconnect Info
========================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                               Remote       Connection-Details
================================== ============
================================================================================
VPN-ID       EVI-Name      MTU   VPN-ID       Source       Destination
PE-IP           MTU   Type   NW-Status
================================== ============
================================================================================
20           ----          1500  10           xe6 2        --- Single Homed Port ---
1.1.1.1         1500  AC-NW  NW-SET

Total number of entries are 1
VTEP-2#sh nvo vxlan tunnel
VxLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=====================================================================
2.2.2.2          1.1.1.1          Installed     00:01:16      00:01:16

Total number of entries are 1
VTEP-2#
```

Multi-Homed VxLAN Eline xConnect

Topology

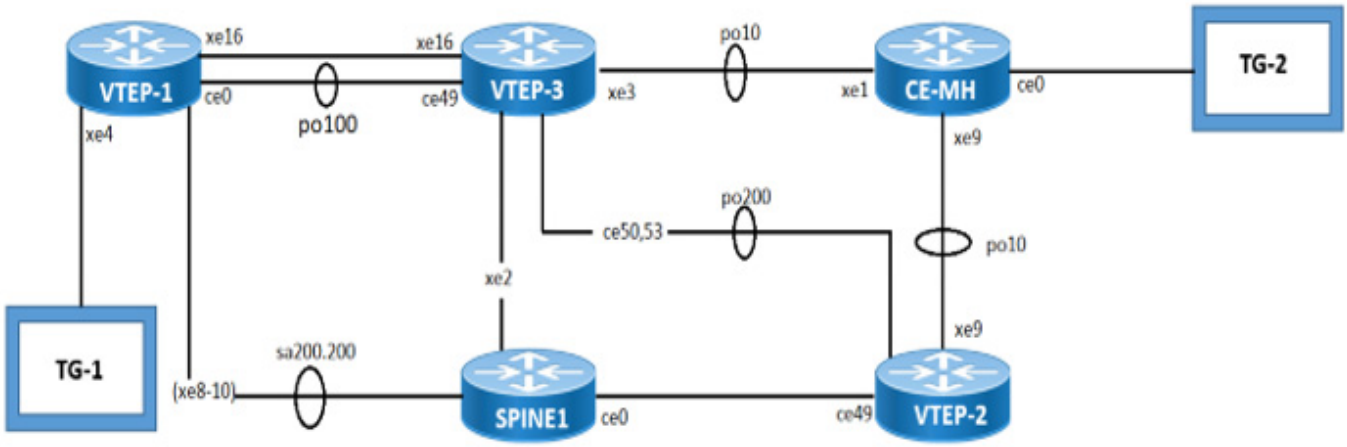

Figure 3-12: Multi-homed VxLAN Eline xConnect topology

L2, IPv4, and IPv6 traffic streams are sent from TG-1 and TG-2.

VTEP-1

`#con t`	Enter configure mode
`(config)#hostname VTEP1`	Configure hostname
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware filter for egress ipv4
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware-profile filter for multi-homed VxLAN
`(config)#hardware-profile statistics ingress-acl enable`	Configure hardware profile statistics ingress-acl
`(config)#nvo vxlan enable`	Enable VxLAN globally on this vtep
`(config)#mac vrf evpn_mh_eline`	Configure a new VRF
`(config-vrf)#rd 1.1.1.1:2`	Assign the Route Distinguisher value.
`(config-vrf)#route-target both 2:2`	Configure route target to import and export the routes
`(config-vrf)#evpn vxlan multihoming enable`	Enable evpn vxlan multihoming
`(config-vrf)#exit`	Exit vrf mode
`(config)#interface po100`	Enter interface mode
`(config-if)#ip address 51.1.1.1/24`	Configure IP address on the interface
`(config-if)#isis circuit-type level-1`	Configure instance as level-1-only routing.
`(config-if)#ip router isis VxLAN`	Enable IS-IS routing on the interface
`(config-if)#exit`	Exit interface mode
`(config)#interface sa200`	Enter interface mode

(config-if)#interface sa200.200	Enter interface mode
(config-if)#encapsulation dot1q 200	Configure the encapsulation as dot1q matching vlan 200
(config-if)#ip address 61.1.1.1/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface ce0	Enter interface mode
(config-if)#channel-group 100 mode active	Add this interface to channel group 100 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface lo	Enter interface mode
(config-if)#ip address 1.1.1.1/32 secondary	Configure IP address on the interface
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface xe4	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#interface xe8	Enter interface mode
(config-if)#static-channel-group 200	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface xe9	Enter interface mode
(config-if)#static-channel-group 200	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface xe10	Enter interface mode
(config-if)#static-channel-group 200	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface xe16.100	Enter interface mode
(config-if)#encapsulation dot1q 100	Configure the encapsulation as dot1q matching vlan 100
(config-if)#ip address 50.1.1.1/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#exit	Exit interface mode
(config)#router isis VxLAN	Create an IS-IS routing instance
(config-router)#is-type level-1	Configure instance as level-1-only routing
(config-router)#spf-interval-exp 0 0	Set the minimum and maximum hold intervals between Shortest Path First (SPF) calculations
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#mpls traffic-eng level-1	Enable MPLS-TE in is-type Level-1.
(config-router)#dynamic-hostname	Configure a hostname to use for the Dynamic Hostname Exchange Mechanism and System-ID to hostname translation
(config-router)#bfd all-interfaces	Enable BFD for all-interface on ISIS
(config-router)#net 49.0000.0000.0001.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit router 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 3.3.3.3 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 4.4.4.4 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 3.3.3.3 update-source lo	Specify the neighbor to use loopback address as source
(config-router)#neighbor 4.4.4.4 update-source lo	Specify the neighbor to use loopback address as source
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 3.3.3.3 activate	Activate the peer into address family mode
(config-router-af)#neighbor 4.4.4.4 activate	Activate the peer into address family mode
(config-router-af)#exit-address-family	Exit l2vpn address family mode
(config-router)#nvo vxlan vtep-ip-global 1.1.1.1	Configure Source vtep-ip-global configuration
(config)#nvo vxlan id 2 xconnect target-vxlan-id 1002	add a tenant and the type of VPN. This creates an ELAN with source and target identifier for ELINE XConnect
(config-nvo)#vxlan host-reachability-protocol evpn-bgp evpn_mh_eline	Set the host reachable protocol to Ethernet-
(config)#interface xe4.2 switchport	Create L2 subinterface xe4.2
(config-if)#encapsulation dot1q 2	Configure encapsulation dot1q with vlan id 2
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 2	Map vnid 2 to access interface xe4.2
(config)#interface xe4.2000 switchport	Create L2 subinterface xe4.2000
(config-if)#encapsulation dot1q 2000	Configure encapsulation dot1q with vlan id 2000
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 2000	Map vnid 2000 to access interface xe4.2000
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration.

VTEP-2

#con t	Enter configure mode
(config)#hostname VTEP2	Configure hostname
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware filter for egress ipv4
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN
(config)#hardware-profile filter vxlan-mh enable	Enable hardware-profile filter for multi-homed VxLAN
(config)#hardware-profile statistics ingress-acl enable	Configure hardware profile statistics ingress-acl
(config)#nvo vxlan enable	Enable VxLAN globally on this vtep
(config)#mac vrf evpn_mh_eline	Configure a new VRF
(config-vrf)#rd 4.4.4.4:2	Assign the Route Distinguisher value.

(config-vrf)#route-target both 2:2	Configure route target to import and export the routes
(config-vrf)#evpn vxlan multihoming enable	Enable evpn vxlan multihoming
(config)#vlan database	Enter VLAN configure mode
(config-vlan)#vlan 200 bridge 1	Configure a VLAN and add it to the bridge.
(config-vlan)#interface po10	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#evpn multi-homed system-mac 0000.0000.1111	Configure system mac as ESI value for the interface
(config-if)#interface po200	Enter interface mode
(config-if)#ip address 63.1.1.2/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface ce49	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#bridge-group 1	Associate bridge to an interface.
(config-if)#switchport mode trunk	Configure port as a trunk.
(config-if)#switchport trunk allowed vlan add 200	Allow VLAN 200 on the interface.
(config-if)#interface ce50	Enter interface mode
(config-if)#channel-group 200 mode active	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface ce53	Enter interface mode
(config-if)#channel-group 200 mode active	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface lo	Enter interface mode
(config-if)#ip address 4.4.4.4/32 secondary	Configure IP address on the interface
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface vlan1.200	Enter interface mode
(config-if)#ip address 20.1.1.2/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface xe9	Enter interface mode
(config-if)#channel-group 10 mode active	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#exit	Exit interface mode
(config)#router isis VxLAN	Create an IS-IS routing instance
(config-router)#is-type level-1	Configure instance as level-1-only routing
(config-router)#spf-interval-exp 0 0	Set the minimum and maximum hold intervals between Shortest Path First (SPF) calculations
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#mpls traffic-eng level-1	Enable MPLS-TE in is-type Level-1.

(config-router)#dynamic-hostname	Configure a hostname to use for the Dynamic Hostname Exchange Mechanism and System-ID to hostname translation
(config-router)#bfd all-interfaces	Enable BFD for all-interface on ISIS
(config-router)#net 49.0000.0000.0004.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit isis interface mode
(config-router)#router bgp 100	Enter BGP router mode
(config-router)#bgp router-id 4.4.4.4	Assign BGP router ID
(config-router)#neighbor 1.1.1.1 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 3.3.3.3 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Specify the neighbor to use loopback address as source
(config-router)#neighbor 3.3.3.3 update-source lo	Specify the neighbor to use loopback address as source
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate the peer into address family mode
(config-router-af)#neighbor 3.3.3.3 activate	Activate the peer into address family mode
(config-router-af)#exit-address-family	Exit l2vpn address family mode
(config-router)#nvo vxlan vtep-ip-global 4.4.4.4	Configure Source vtep-ip-global configuration
(config)#nvo vxlan id 3000 xconnect target-vxlan-id 2000	add a tenant and the type of VPN. This creates an ELAN with source and target identifier for ELINE XConnect
(config-nvo)#vxlan host-reachability-protocol evpn-bgp evpn_mh_eline	Set the host reachable protocol to Ethernet-
(config)#interface po10.2000 switchport	Create L2 subinterface po10.2000
(config-if)#encapsulation dot1q 2000	Configure encapsulation dot1q with vlan id 3000
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 3000	Map vnid 3000 to access interface po10.2000
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration.

VTEP-3

#con t	Enter configure mode
(config)#hostname VTEP3	Configure host name
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware filter for egress ipv4
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN
(config)#hardware-profile filter vxlan-mh enable	Enable hardware-profile filter for multi-homed VxLAN
(config)#hardware-profile statistics ingress-acl enable	Configure hardware profile statistics ingress-acl
(config)#nvo vxlan enable	Enable VxLAN globally on this vtep

(config)#mac vrf evpn_mh_eline	Configure a new VRF
(config-vrf)#rd 3.3.3.3:2	Assign the Route Distinguisher value.
(config-vrf)#route-target both 2:2	Configure route target to import and export the routes
(config-vrf)#evpn vxlan multihoming enable	Enable evpn vxlan multihoming
(config)#bridge 1 protocol ieee vlan-bridge	Configure bridge as IEEE VLAN bridge
(config)#vlan database	Enter VLAN configure mode
(config-vlan)#vlan 300 bridge 1	Configure a VLAN and add it to the bridge.
(config-vlan)#interface po10	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#evpn multi-homed system-mac 0000.0000.1111	Configure system mac as ESI value for the interface
(config-if)#interface po100	Enter interface mode
(config-if)#ip address 51.1.1.2/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface po200	Enter interface mode
(config-if)#ip address 63.1.1.1/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface ce49	Enter interface mode
(config-if)#channel-group 100 mode active	Add this interface to channel group 100 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface ce50	Enter interface mode
(config-if)#channel-group 200 mode active	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface ce53	Enter interface mode
(config-if)#channel-group 200 mode active	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface lo	Enter interface mode
(config-if)#ip address 3.3.3.3/32 secondary	Configure IP address on the interface
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface vlan1.300	Enter interface mode
(config-if)#ip address 40.1.1.2/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface xe2	Enter interface mode
(config-if)#ip address 62.1.1.1/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface xe3	Enter interface mode

(config-if)#channel-group 10 mode active	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface xe16.100	Enter interface mode
(config-if)#encapsulation dot1q 100	Configure the encapsulation as dot1q matching vlan 100
(config-if)#ip address 50.1.1.2/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#exit	Exit interface mode
(config)#router isis VxLAN	Create an IS-IS routing instance
(config-router)#is-type level-1	Configure instance as level-1-only routing
(config-router)#spf-interval-exp 0 0	Set the minimum and maximum hold intervals between Shortest Path First (SPF) calculations
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#mpls traffic-eng level-1	Enable MPLS-TE in is-type Level-1.
(config-router)#dynamic-hostname	Configure a host name to use for the Dynamic Hostname Exchange Mechanism and System-ID to host name translation
(config-router)#bfd all-interfaces	Enable BFD for all-interface on ISIS
(config-router)#net 49.0000.0000.0006.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit isis router mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#bgp router-id 3.3.3.3	Assign BGP router ID
(config-router)#neighbor 1.1.1.1 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 4.4.4.4 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#neighbor 1.1.1.1 update-source lo	Specify the neighbor to use loopback address as source
(config-router)#neighbor 4.4.4.4 update-source lo	Specify the neighbor to use loopback address as source
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate the peer into address family mode
(config-router-af)#neighbor 4.4.4.4 activate	Activate the peer into address family mode
(config-router-af)#exit-address-family	Exit l2vpn address family mode
(config-router)#nvo vxlan vtep-ip-global 3.3.3.3	Configure Source vtep-ip-global configuration
(config)#nvo vxlan id 3000 xconnect target-vxlan-id 2000	add a tenant and the type of VPN. This creates an ELAN with source and target identifier for ELINE XConnect
(config-nvo)#vxlan host-reachability-protocol evpn-bgp evpn_mh_eline	Set the host reachable protocol to Ethernet-
(config)#interface po10.2000 switchport	Create L2 subinterface po10.2000
(config-if)#encapsulation dot1q 2000	Configure encapsulation dot1q with vlan id 3000
(config-if)#access-if-evpn	Enter access-if-evpn mode

(config-acc-if-evpn)#map vpn-id 3000	Map vnid 3000 to access interface po10.2000
(config-nvo-acc-if)#commit	Commit the candidate configuration to the running configuration.

SPINE

#con t	Enter configure mode
(config)#hostname SPINE	Configure host name
(config)#hardware-profile statistics ingress-acl enable	Configure hardware profile statistics ingress-acl
(config)#bridge 1 protocol ieee vlan-bridge	Configure bridge as IEEE VLAN bridge
(config)#vlan database	Enter VLAN configure mode
(config-vlan)#vlan 200 bridge 1	Configure a VLAN and add it to the bridge.
(config-if)#interface sa200	Enter interface mode
(config-if)#interface sa200.200	Enter interface mode
(config-if)#encapsulation dot1q 200	Configure the encapsulation as dot1q matching vlan 200
(config-if)#ip address 61.1.1.2/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface ce0	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#bridge-group 1	Associate bridge to an interface.
(config-if)#switchport mode trunk	Configure port as a trunk.
(config-if)#switchport trunk allowed vlan add 200	Allow VLAN 200 on the interface.
(config-if)#interface lo	Enter interface mode
(config-if)#ip address 11.11.11.11/32 secondary	Configure IP address on the interface
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface vlan1.200	Enter interface mode
(config-if)#ip address 20.1.1.1/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface xe2	Enter interface mode
(config-if)#ip address 62.1.1.2/24	Configure IP address on the interface
(config-if)#isis circuit-type level-1	Configure instance as level-1-only routing.
(config-if)#ip router isis VxLAN	Enable IS-IS routing on the interface
(config-if)#interface xe8	Enter interface mode
(config-if)#static-channel-group 200	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface xe9	Enter interface mode

(config-if)#static-channel-group 200	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface xe10	Enter interface mode
(config-if)#static-channel-group 200	Add this interface to channel group 200 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#exit	Exit interface mode
(config-if)#router isis VxLAN	Create an IS-IS routing instance
(config-router)#is-type level-1	Configure instance as level-1-only routing
(config-router)#spf-interval-exp 0 0	Set the minimum and maximum hold intervals between Shortest Path First (SPF) calculations
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#mpls traffic-eng level-1	Enable MPLS-TE in is-type Level-1.
(config-router)#dynamic-hostname	Configure a host name to use for the Dynamic Hostname Exchange Mechanism and System-ID to host name translation
(config-router)#bfd all-interfaces	Enable BFD for all-interface on ISIS
(config-router)#net 49.0000.0000.0003.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit isis interface mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#neighbor 1.1.1.1 remote-as 100	Assign BGP router ID
(config-router)#neighbor 5.5.5.5 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
(config-router)#address-family l2vpn evpn	Enter into l2vpn address family mode
(config-router-af)#neighbor 1.1.1.1 activate	Activate the peer into address family mode
(config-router-af)#neighbor 1.1.1.1 route-reflector-client	Configure this node as the route reflector with the mentioned peer as its client
(config-router-af)#neighbor 5.5.5.5 activate	Activate the peer into address family mode
(config-router-af)#neighbor 5.5.5.5 route-reflector-client	Configure this node as the route reflector with the mentioned peer as its client
(config-router-af)#exit-address-family	Exit l2vpn address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

CE-MH

#con t	Enter configure mode
(config)#hostname CE3-MH	Configure host name
(config)#hardware-profile statistics ingress-acl enable	Configure hardware profile statistics ingress-acl
(config)#interface po10	Enter interface mode
(config-if)#switchport	Configure the interface as switchport
(config-if)#interface po10.100 switchport	Enter interface mode

(config-if)#encapsulation default	Configure the encapsulation as default
(config-if)#interface ce0	Enter interface mode
(config-if)#interface ce0.100 switchport	Enter interface mode
(config-if)#encapsulation default	Configure the encapsulation as default
(config-if)#interface xe1	Enter interface mode
(config-if)#channel-group 10 mode active	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#interface xe9	Enter interface mode
(config-if)#channel-group 10 mode active	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
(config-if)#cross-connect MH-CE3	Create cross-connect with name MH-CE3
(config-xc)#interface po10.100	Attach interface po10.100
(config-xc)#interface ce0.100	Attach interface ce0.100
(config-xc)#commit	Commit the candidate configuration to the running configuration.

Validation

```
VTEP1#show nvo vxlan xconnect
EVPN Xconnect Info
========================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                               Remote       Connection-Details
================================= ============
=============================================================================
VPN-ID       EVI-Name      MTU   VPN-ID       Source       Destination
PE-IP           MTU   Type   NW-Status
================================= ============
=============================================================================
2000         ----          1500  3000         xe4 2000     00:00:00:00:00:11:11:00:00:00
4.4.4.4         1500  AC-NW  NW-SET

3.3.3.3         1500  ----   ----

Total number of entries are 1
VTEP1-7012#
VTEP1-7012#show nvo vxlan tunnel
VxLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================
1.1.1.1          4.4.4.4          Installed     00:03:30      00:03:30
```

```
1.1.1.1          3.3.3.3          Installed     00:03:27      00:03:27

Total number of entries are 2
VTEP1#


VTEP3#show nvo vxlan xconnect
EVPN Xconnect Info
========================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                              Remote       Connection-Details
================================== ============
================================================================================
VPN-ID       EVI-Name      MTU   VPN-ID       Source       Destination
PE-IP           MTU   Type   NW-Status
================================== ============
================================================================================
3000         ----          1500  2000         po10 2000    --- Single Homed Port ---
1.1.1.1         1500  AC-NW  NW-SET

Total number of entries are 1
VTEP3#


VTEP2#show nvo vxlan xconnect
EVPN Xconnect Info
========================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                              Remote       Connection-Details
================================== ============
================================================================================
VPN-ID       EVI-Name      MTU   VPN-ID       Source       Destination
PE-IP           MTU   Type   NW-Status
================================== ============
================================================================================
3000         ----          1500  2000         po10 2000    --- Single Homed Port ---
1.1.1.1         1500  AC-NW  NW-SET
```

```
Total number of entries are 1
VTEP2#
```

CHAPTER 4 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 following topology:

Figure 4-13: DHCP Relay over IRB

ROUTER-1

`#configure terminal`	Enter Configure mode.
`(config)#interface lo`	Enter Interface mode for loopback.
`(config-if)#ip address 1.1.1.1/32 secondary`	Assign secondary IP address.
`(config-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#nvo vxlan enable`	Enable VxLAN
`(config)#nvo vxlan irb`	Enable VxLAN IRB
`(config)#ip vrf vrf1`	Create routing/forwarding instance with VRF1 name and enter into VRF mode
`(config-vrf)#rd 200:1`	Assign RD value
`(config-vrf)#route-target both 200:1`	Assign route target value
`(config-vrf)#ip dhcp relay address 40.40.40.1`	The relay address configured should be server interface address connected to DUT machine
`(config-vrf)#ip dhcp relay uplink evpn`	Configure the uplink interface as L3 VNI interface for specific VRF
`(config-vrf)#l3vni 45001`	Configure L3VNI as 45001 for VRF1

(config-vrf)#exit	Exit IP VRF mode
(config)#mac vrf vrfred	Create MAC VRF instance with vrfred name and enter into VRF mode
(config-vrf)#rd 1.1.1.1:1	Assign RD value
(config-vrf)#route-target both 1.1.1.1:1	Assign route target value
(config-vrf)#exit	Exit MAC VRF mode
(config)#interface irb 1	Configure IRB interface
(config-irb-if)#ip vrf forwarding vrf1	Configure IP VRF forwarding
(config-irb-if)#ip address 11.1.1.1/24	Assign IP address on IRB interface.
(config-irb-if)#ip dhcp relay	Relay should be configured on the interface connecting to the relay
(config-irb-if)#exit	Exit IRB interface mode
(config)#interface irb 2	Configure irb interface
(config-irb-if)#ip vrf forwarding vrf1	Configure IP VRF forwarding
(config-irb-if)#ip address 70.70.70.1/24	Assign IP address on IRB interface.
(config-irb-if)#exit	Exit IRB interface mode
(config)#interface ce49	Enter Interface mode for ce49.
(config-if)#ip address 10.1.1.0/31	Assign IP address on ce49 interface.
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe5	Enter Interface mode for xe5.
(config-if)#switchport	Configure interface as L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#router ospf	Enter the Router OSPF mode
(config-router)#network 1.1.1.1/32 area 0.0.0.0	Advertise loopback address in OSPF
(config-router)#network 10.1.1.0/31 area 0.0.0.0	Advertise network address in OSPF
(config-router)#exit	Exit from Router OSPF mode and enter into config mode
(config)#router bgp 1	Enter into BGP router mode
(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 1.1.1.1	Configure update as loopback 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 neighbor in L2VPN mode
(config-router-af)#exit-address-family	Exit from Address family mode
(config-router)#address-family ipv4 vrf vrf1	Enter into address-family mode for VRF1
(config-router-af)#redistribute connected	Configure Redistribute connected
(config-router-af)#exit-address-family	Exit from Address family mode
(config-router)#exit	Exit from router BGP mode and enter into config mode
(config)#nvo vxlan vtep-ip-global 1.1.1.1	Configure Source VTEP-IP-global configuration. Use loopback IP address
(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 vrfred	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#evpn irb1	Configure IRB1 under VxLAN ID 10
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(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
(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrfred	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#evpn irb2	Configure IRB2 under VxLAN ID 30
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#interface xe5.2 switchport	Create L2 subinterface xe5.2
(config-if)#encapsulation dot1q 2	Configure encapsulation dot1q with vlan id 10
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 10	Map vnid 10 to access interface xe5.2
(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

ROUTER-2

#configure terminal	Enter Configure mode.
(config)#interface ce0	Enter Interface mode for ce0.
(config-if)#ip address 10.1.1.1/31	Assign IP address on ce0 interface.
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe13	Enter Interface mode for xe13.
(config-if)#ip address 30.1.1.1/31	Assign IP address on xe13 interface.
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#router ospf	Enter the Router OSPF mode
(config-router)#network 10.1.1.0/24 area 0.0.0.0	Advertise network address in OSPF
(config-router)#network 30.30.30.0/24 area 0.0.0.0	Advertise network address in OSPF
(config-router)#exit	Exit from Router OSPF mode and enter into config mode
(config)#commit	commit the candidate configuration to the running configuration
(config)#exit	Exit from configuration mode

ROUTER-3

#configure terminal	Enter Configure mode.
(config)#interface lo	Enter Interface mode for loopback.

(config-if)#ip address 2.2.2.2/32 secondary	Assign secondary IP address.
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#nvo vxlan enable	Enable VxLAN
(config)#nvo vxlan irb	Enable VxLAN IRB
(config)#ip vrf vrf1	Create routing/forwarding instance with VRF1 name and enter into VRF mode
(config-vrf)#rd 300:1	Assign RD value
(config-vrf)#route-target both 200:1	Assign route target value
(config-vrf)#ip dhcp relay uplink evpn	Configure the uplink interface as L3 VNI interface for specific VRF
(config-vrf)#l3vni 45001	Configure L3VNI as 45001 for VRF1
(config-vrf)#exit	Exit IP VRF mode
(config)#mac vrf vrfred	Create MAC VRF instance with vrfred name and enter into VRF mode
(config-vrf)#rd 2.2.2.1:1	Assign RD value
(config-vrf)#route-target both 1.1.1.1:1	Assign route target value
(config-vrf)#exit	Exit MAC VRF mode
(config)#interface irb 2	Configure IRB interface
(config-irb-if)#ip vrf forwarding vrf1	Configure IP VRF forwarding
(config-irb-if)#ip address 40.40.40.2/24	Assign IP address on IRB interface.
(config-irb-if)#exit	Exit IRB interface mode
(config)#interface xe13	Enter Interface mode for xe13.
(config-if)#ip address 30.1.1.0/31	Assign IP address on xe13 interface.
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#interface xe19	Enter Interface mode for xe19.
(config-if)#switchport	Configure interface as L2 interface
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#router ospf	Enter the Router OSPF mode
(config-router)#network 2.2.2.2/32 area 0.0.0.0	Advertise loopback address in OSPF
(config-router)#network 30.1.1.0/24 area 0.0.0.0	Advertise network address in OSPF
(config-router)#network 40.1.1.0/24 area 0.0.0.0	Advertise network address in OSPF
(config-router)#exit	Exit from Router OSPF mode and enter into config mode
(config)#router bgp 1	Enter into BGP router mode
(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 2.2.2.2	Configure update as loopback 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 neighbor in L2VPN mode
(config-router-af)#exit-address-family	Exit from Address family mode

(config-router)#address-family ipv4 vrf vrf1	Enter into address-family mode for VRF1
(config-router-af)#redistribute connected	Configure Redistribute connected
(config-router-af)#exit-address-family	Exit from Address family mode
(config-router)#exit	Exit from router BGP mode and enter into config mode
(config)#nvo vxlan vtep-ip-global 2.2.2.2	Configure Source VTEP-IP-global configuration. Use loopback IP address
(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 vrfred	Assign VRF for EVPN-BGP to carry EVPN route
(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 vrfred	Assign VRF for EVPN-BGP to carry EVPN route
(config-nvo)#evpn irb2	Configure IRB2 under VxLAN ID 20
(config-nvo)#exit	Exit from VxLAN tenant mode and enter into configuration mode.
(config)#interface xe19 switchport	Create L2 subinterface xe19
(config-if)#encapsulation dot1q	Configure encapsulation dot1q with vlan id 20
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 20	Map vnid 20 to access interface xe19
(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

Validation

ROUTER-1

```
VTEP1#show running-config nvo vxlan
!
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 5 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 5-14.

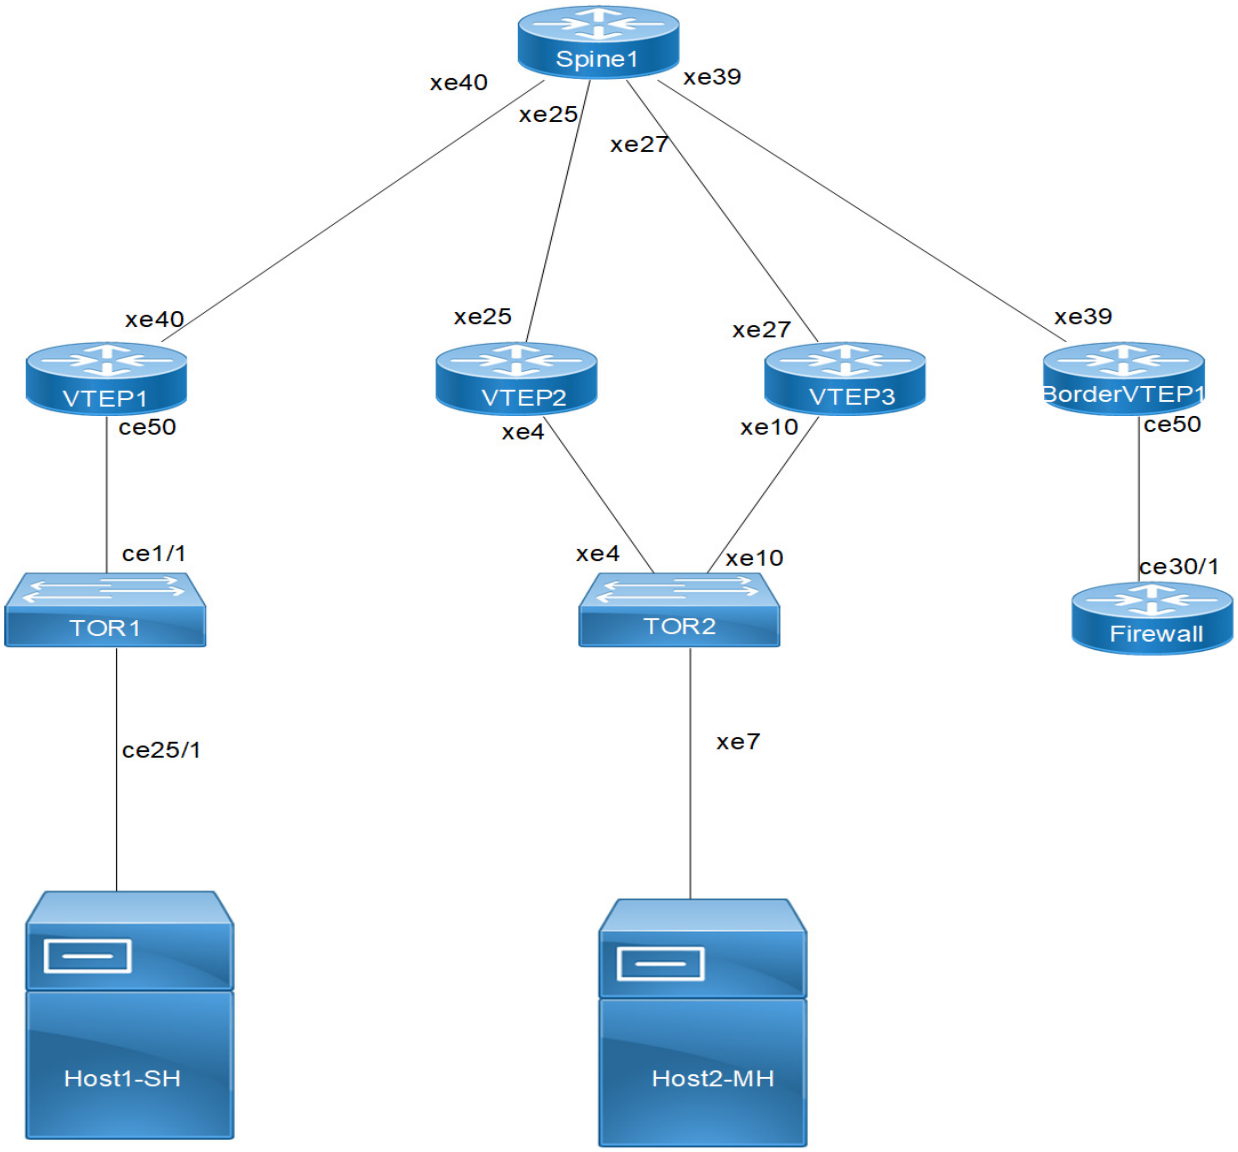

Figure 5-14: VxLAN_EVPN_IVRF

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)#hardware-profile filter vxlan enable`	Enable hardware-profile filter for VxLAN.
`(config)#hardware-profile filter vxlan-mh enable`	Enable hardware-profile filter for VxLAN multi-homing.
`(config)#nvo vxlan enable`	Enable VxLAN
`(config)#nvo vxlan irb`	Enable VxLAN irb
`(config)#hardware-profile filter egress-ipv4 enable`	Enable hardware-profile filter for egress IPv4.
`(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for vxlan access-if port counters
`(config)#qos enable`	Enabling qos
`(Config)# bfd interval 3 minrx 3 multiplier 3`	Configure bfd
`(config)#commit`	Commit the candidate configuration to running configuration

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-if)#exit`	Exit Interface mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to running configuration

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

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)# 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 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

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

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

VxLAN configuration:

(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)#interface ce50.1050 switchport	Create L2 subinterface ce50.1050
(config-if)#encapsulation dot1q 1050	Configure encapsulation dot1q with vlan id 1050
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 1050	Map vnid 1050 to access interface ce50.1050
(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

VTEP2

(Multi-homed group) - Part of both Multi-homed with po1000(MH).

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#hardware-profile filter vxlan-mh enable	Enable hardware-profile filter for VxLAN multi-homing.
(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)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(Config)#hardware-profile statistics ac-lif enable	Enable ac-lif for vxlan access-if port counters
(Config)#qos enable	Enabling qos
(Config)# bfd interval 3 minrx 3 multiplier 3	Configure bfd
(config)#commit	Commit the candidate configuration to running configuration

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

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

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)# 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 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

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

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
(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

VxLAN configuration:

(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)#interface po1000.1060 switchport	Create L2 subinterface po1000.1060
(config-if)#encapsulation dot1q 1060	Configure encapsulation dot1q with vlan id 1060
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 1060	Map vnid 1060 to access interface po1000.1060
(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

VTEP3

(Multi-homed group) - Part of both Multi-homed with po1000(MH).

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.
(config)#hardware-profile filter vxlan-mh enable	Enable hardware-profile filter for VxLAN multi-homing.
(config)#nvo vxlan enable	Enable VxLAN
(config)#nvo vxlan irb	Enable VxLAN irb

(config)#commit	Commit the candidate configuration to running configuration
(config)# evpn irb-forwarding anycast-gateway-mac 0000.2222.3333	Configure Anycast gateway mac
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for vxlan access-if port counters
(config)#qos enable	Enabling qos
(config)# bfd interval 3 minrx 3 multiplier 3	Configure bfd
(config)#commit	Commit the candidate configuration to running configuration

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	Configure loopback ip address
(config-if)#exit	Exit Interface mode and return to Configure mode.
(config)#commit	Commit the candidate configuration to running configuration
(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

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

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)# 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 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

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

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
(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

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)#interface po1000.1060 switchport	Create L2 subinterface po1000.1060
(config-if)#encapsulation dot1q 1060	Configure encapsulation dot1q with vlan id 1060
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 1060	Map vnid 1060 to access interface po1000.1060
(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

BorderVTEP1

Hardware profile and generic configuration:

#configure terminal	Enter Configure mode.
(config)#hardware-profile filter vxlan enable	Enable hardware-profile filter for VxLAN.

(config)#hardware-profile filter vxlan-mh enable	Enable hardware-profile filter for VxLAN multi-homing.
(config)#nvo vxlan enable	Enable VxLAN
(config)#nvo vxlan irb	Enable VxLAN irb
(config)#hardware-profile filter egress-ipv4 enable	Enable hardware-profile filter for egress IPv4.
(Config)#hardware-profile statistics ac-lif enable	Enable ac-lif for vxlan access-if port counters
(Config)#qos enable	Enabling qos
(Config)# bfd interval 3 minrx 3 multiplier 3	Configure bfd
(config)#commit	Commit the candidate configuration to running configuration

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

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 10.10.24.0/24 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

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

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

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

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)#interface ce50.1067 switchport	Create L2 subinterface ce50.1067
(config-if)#encapsulation dot1q 1067	Configure encapsulation dot1q with vlan id 1067
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 1067	Map vnid 1067 to access interface ce50.1067
(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

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

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

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 default-originate	Do default originate
(config-router-af)# exit-address-family	Exit from IPV4 unicast address family
(config-router)#exit	
(config)#commit	Commit the candidate configuration to running configuration

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

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

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

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

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)#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

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 6 Single-Home for 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.

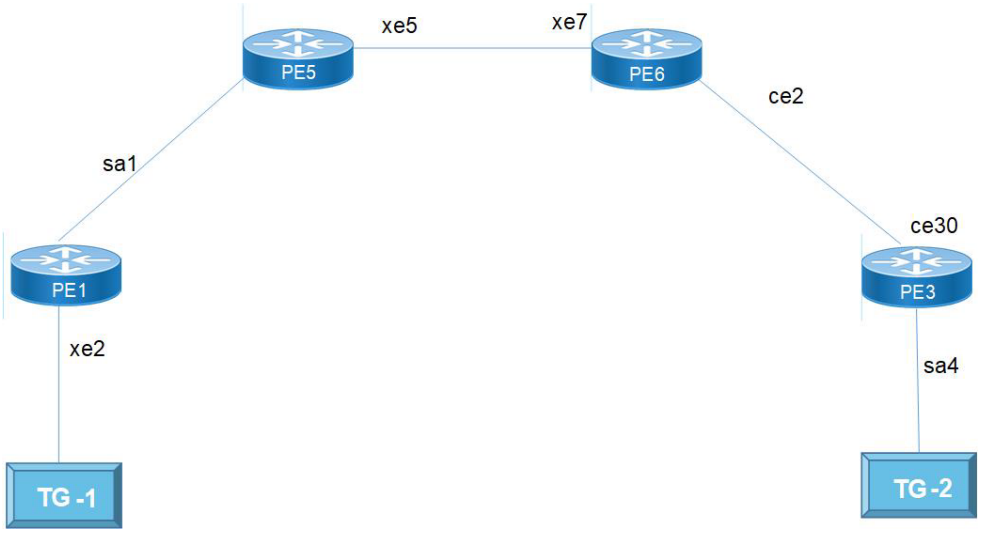

Single Home VxLAN IRB with OSPF

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.

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)#interface xe1`	Enters the interface xe1 mode.
`PE1(config-if)# static-channel-group 1`	Assigns the static channel group 1 to the xe2 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)#hardware-profile filter vxlan enable`	Enables the hardware profile filter for VXLAN.
`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)#interface xe2.100 switchport`	Configures a Layer 2 access interface subinterface
`PE1(config-if)# encapsulation dot1q 100`	Sets the encapsulation to be a single-tag with VLAN ID 100
`PE1(config-if)# rewrite pop`	Configures the rewrite pop action.
`PE1(config-if)# access-if-evpn`	Configures the port as an access interface for EVPN.
`PE1(config-acc-if-evpn)#  map vpn-id 101`	Maps the VPN ID 101 to the interface.
`PE1(config)#interface xe2.2001 switchport`	Configures a Layer 2 access interface subinterface.
`PE1(config-if)# rewrite pop`	Configures the rewrite pop action.
`PE1(config-if)# access-if-evpn`	Configures the port as an access interface for EVPN.
`PE1(config-acc-if-evpn)#  map vpn-id 2001`	Maps the VPN ID 2001 to the interface.
`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)#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)#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)#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)#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)#hardware-profile filter vxlan enable`	Enables hardware profile filter for vxlan
`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)# 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-if)#interface sa4.100 switchport`	Configures subinterface sa4.100 as a switchport with VLAN 100.
`PE3(config-if)# encapsulation dot1q 100`	Configures the encapsulation with VLAN ID 100.v
`PE3(config-if)# rewrite pop`	Configures the rewrite behavior for popping VLAN tags.
`PE3(config-if)# access-if-evpn`	Configures the interface as an access interface for EVPN.
`PE3(config-acc-if-evpn)# map vpn-id 102`	Maps the access interface to VPN ID 102 for EVPN routing.
`PE3(config)#interface sa4.2002 switchport`	Configures subinterface sa4.2002 as a switchport with VLAN 2002.
`PE3(config-if)# encapsulation dot1q 2002`	Configures the encapsulation with VLAN ID 2002.

`PE3(config-if)# rewrite pop`	Configures the rewrite behavior for popping VLAN tags.
`PE3(config-if)# access-if-evpn`	Configures the interface as an access interface for EVPN.
`PE3(config-acc-if-evpn)# map vpn-id 2002`	Maps the access interface to VPN ID 2002 for EVPN routing.
`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)#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)#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.

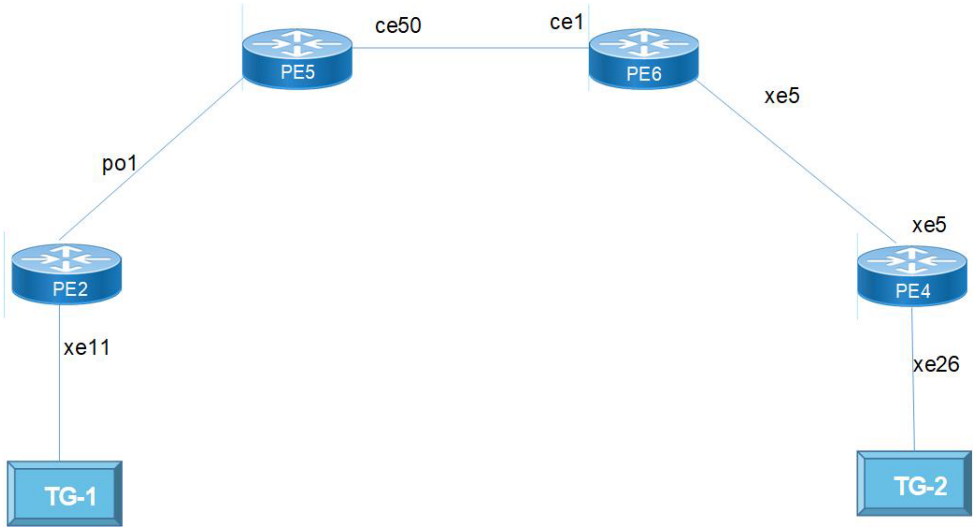

Single Home VxLAN IRB with ISIS

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)#hardware-profile filter vxlan enable`	Enables the hardware profile filter for VxLAN on the device.
`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)#interface xe11.200 switchport`	Enters the configuration mode for subinterface xe11.200 and configures it as a Layer 2 switchport.
`PE2(config-if)#encapsulation dot1q 200`	Sets the IEEE 802.1Q VLAN encapsulation for subinterface xe11.200 with VLAN ID 200.
`PE2(config-if)#rewrite pop`	Configures the subinterface to rewrite the outer header for provider edge devices.
`PE2(config-if)#access-if-evpn`	Configures the subinterface as an access interface for EVPN (Ethernet VPN).
`PE2(config-if)#map vpn-id 201`	Maps the VPN ID 201 to the subinterface for EVPN.
`PE2(config-if)#interface xe11.3001 switchport`	Configures interface xe11.3001.
`PE2(config-if)#encapsulation dot1q 200`	Sets the IEEE 802.1Q VLAN encapsulation for subinterface xe11.200 with VLAN ID 200.
`PE2(config-if)#rewrite pop`	Configures the subinterface to rewrite the outer header for provider edge devices.
`PE2(config-if)#access-if-evpn`	Configures the subinterface as an access interface for EVPN.
`PE2(config-if)#map vpn-id 3001`	Maps the VPN ID 3001 to the subinterface for EVPN.
`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)#mpls traffic-eng router-id 2.2.2.2`	Sets the MPLS Traffic Engineering router ID to 2.2.2.2.
`PE2(config-if)#mpls traffic-eng level-1`	Enables MPLS Traffic Engineering for Level 1 ISIS.
`PE2(config-if)#mpls traffic-eng level-2`	Enables MPLS Traffic Engineering for Level 2 ISIS.
`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)#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)#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 xe5`	Enters configuration mode for interface xe5.

`PE6(config-if)#ip address 60.1.1.1/24`	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)#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)#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 xe26`	Enters configuration mode for xe26.
`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)#ip router isis 1`	Enables ISIS routing protocol on the interface with process ID 1.
`PE4(config)#hardware-profile filter vxlan enable`	Enables the hardware profile filter for VXLAN.
`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 xe26`	Enters the configuration mode for the interface 26.
`PE4(config-if)#switchport`	Configures the interface as a L2 switchport.
`PE4(config)#hardware-profile filter vxlan enable`	Enables the hardware profile filter for VxLAN on the device.
`PE4(config-if)#interface xe26.200 switchport`	Enters the configuration mode for subinterface xe26.200 and configures it as a Layer 2 switchport.
`PE4(config-if)#encapsulation dot1q 200`	Sets the IEEE 802.1Q VLAN encapsulation for subinterface xe11.200 with VLAN ID 200.
`PE4(config-if)#rewrite pop`	Configures the subinterface to rewrite the outer header for provider edge devices.
`PE4(config-if)#access-if-evpn`	Configures the subinterface as an access interface for EVPN (Ethernet VPN).
`PE4(config-if)#map vpn-id 201`	Maps the VPN ID 201 to the subinterface for EVPN.

`PE4(config-if)#interface xe26.3001 switchport`	Configures xe26.3001 interface.
`PE4(config-if)#encapsulation dot1q 200`	Sets the IEEE 802.1Q VLAN encapsulation for subinterface xe11.200 with VLAN ID 200.
`PE4(config-if)#rewrite pop`	Configures the subinterface to rewrite the outer header for provider edge devices.
`PE4(config-if)#access-if-evpn`	Configures the subinterface as an access interface for EVPN.
`PE4(config-if)#map vpn-id 3002`	Maps the VPN ID 3002 to the subinterface for EVPN.
`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)#mpls traffic-eng router-id 2.2.2.2`	Sets the MPLS Traffic Engineering router ID to 2.2.2.2.
`PE4(config-if)#mpls traffic-eng level-1`	Enables MPLS Traffic Engineering for Level 1 ISIS.
`PE4(config-if)#mpls traffic-eng level-2`	Enables MPLS Traffic Engineering for Level 2 ISIS.
`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 2.2.2.32 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.

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 Netwrok
MPLS	Multiprotocol Label Switching
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 7 Single-Home for VxLAN EVPN IRB with OSPF or ISIS

Overview

Single Home EVPN-MPLS IRB with OSPF and ISIS feature provides streamline routing and bridging operations within Service Provider (SP) networks. It seamlessly integrates EVPN-MPLS architecture and handles both Layer 3 routing and Layer2 bridging, making it an ideal choice for SP networks. This feature utilizes routing protocols, OSPF and ISIS, known for their efficiency and reliability. It also supports Integrated Routing and Bridging (IRB), enabling the integration of Layer3 IP routing and Layer2 MAC address bridging within the same interface. It simplifies network operations, enhances efficiency, and ensures network reliability for SP in dynamic environments.

Feature Characteristics

The Single Home EVPN-MPLS IRB with OSPF and ISIS feature has the following characteristics:

- Integrates with the EVPN-MPLS architecture, providing a comprehensive solution for Layer 3 routing and Layer 2 bridging within the network.
- Supports IRB, allowing for the integration of Layer 3 routing and Layer 2 bridging within the same interface.

Benefits

The Single Home EVPN-MPLS IRB with OSPF and ISIS feature has the following benefits:

- With OSPF and ISIS support, the network can effortlessly scale to accommodate increased traffic and new network elements.
- The use of robust routing protocols like OSPF and ISIS ensures reliable and fault-tolerant network communication, minimizing downtime and enhancing network stability.
- The integration of multiple functionalities into a single feature simplifies network management, reducing operational complexity and lowering the risk of configuration errors.

Prerequisites

- Router must be up and running.

Topology for OSPF

The network topology includes various network elements such as routers, customer edge (CE) 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.

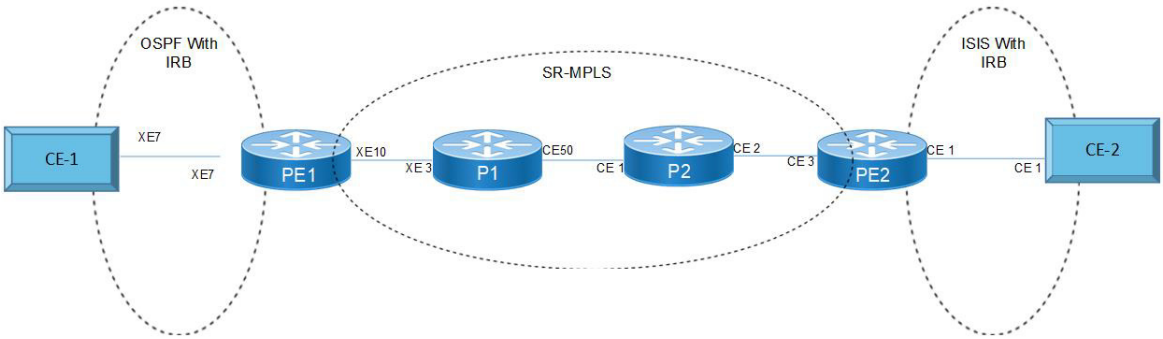

Single Home EVPN MPLS IRB with OSPF and ISIS

Configuration

Perform the following configurations to set up interfaces, assign IP addresses, configure VLAN encapsulation, establish OSPF routing, and prepare the network for routing and forwarding operations. Please ensure that the subnet mask values are accurate and suitable for your network requirements.

CE1

`CE1#configure terminal`	Enters the configuration mode on the OcNOS device.
`CE1(config)#interface po1`	Enters the configuration mode for a Port-Channel interface named po1.
`CE1(config-if)# load-interval 30`	Sets the load interval for the Port-Channel interface to 30 seconds.
`CE1(config-if)#interface po1.101`	Enters the configuration mode for a subinterface of Port-Channel 1, specifically subinterface 101.
`CE1(config-if)# encapsulation dot1q 101`	Configures the subinterface to use IEEE 802.1Q encapsulation with VLAN ID 101.
`CE1(config-if)# ip address 101.101.101.2/24`	Assigns the IP address 101.101.101.2 with a subnet mask of /24 to the subinterface.
`CE1(config-if)#interface po1.1001`	Enters the configuration mode for another subinterface of Port-Channel 1, specifically subinterface 1001.
`CE1(config-if)# encapsulation dot1q 1001`	Configures the subinterface with IEEE 802.1Q encapsulation using VLAN ID 1001.
`CE1(config-if)# ipv6 address 1001::2/64`	Assigns an IPv6 address 1001::2 with a subnet prefix length of /64 to the subinterface.
`CE1(config-if)# ipv6 router ospf area 0.0.0.0 tag 1001 instance-id 0`	Configures OSPF for IPv6 on the subinterface, specifying area 0.0.0.0 with a tag of 1001 and instance ID 0.

`CE1(config-if)# ip address 99.99.99.1/243`	Assigns the IP address 99.99.99.1 with a subnet mask of /243 to the Gigabit Ethernet interface. Note: The subnet mask value seems to be incorrect; it should typically be in the format /24.
`CE1(config-if)#interface lo`	Enters the configuration mode for a loopback interface.
`CE1(config-if)# ip address 127.0.0.1/8`	Assigns the primary IP address 127.0.0.1 with a subnet mask of /8 to the loopback interface.
`CE1(config-if)# ip address 1.1.1.1/32 secondary`	Adds a secondary IP address, 1.1.1.1 with a /32 subnet mask, to the loopback interface.
`CE1(config-if)# ipv6 address ::1/128`	Assigns the IPv6 address ::1 with a subnet prefix length of /128 to the loopback interface.
`CE1(config-if)#interface xe7`	Enters the configuration mode for Ten Gigabit Ethernet interface 7.
`CE1(config-if)# channel-group 1 mode active`	Configures the interface to be part of an EtherChannel group with group number 1 and sets the mode to active.
`CE1(config-if)#exit`	Exits the interface configuration mode.
`CE1(config)#router ospf 101`	Enters the configuration mode for OSPF (routing process 101.
`CE1(config-router)# ospf router-id 1.1.1.1`	Sets the OSPF router ID to 1.1.1.1.
`CE1(config-router)# redistribute connected`	Configures OSPF to redistribute connected routes into the OSPF routing process.
`CE1(config-router)# network 101.101.101.0/24 area 0.0.0.0`	Specifies that the network 101.101.101.0/24 is part of OSPF Area 0.0.0.0.
`CE1(config-router)#router ipv6 ospf 1001`	Enters the configuration mode for OSPFv3 (OSPF for IPv6) routing process 1001.
`CE1(config-router)# router-id 1.1.1.1`	Sets the OSPFv3 router ID to 1.1.1.1.
`CE1(config-router)#Exit`	Exists from the configurations.

Configure MTU over IRB

`CE1(config)#int po1`	Enters interface configuration mode for port-channel 1 (po1).
`CE1(config-if)#mtu 9216`	Sets the Maximum Transmission Unit (MTU) for the interface po1 to 9216 bytes, which is a jumbo frame MTU size.
`CE1(config-if)#commit`	Commit the changes made to the interface configuration.
`CE1((config-if)#int po1.101`	Enters interface configuration mode for subinterface po1.101.
`CE1(config-if)#mtu 9216`	Sets the MTU to 9216 bytes.
`CE1(config-if)#int po1.1001`	Enters interface configuration mode for subinterface po1.1001.
`CE1(config-if)#mtu 9216`	Sets the MTU to 9216 bytes.
`CE1(config-if)#Exit`	Exits the interface configuration mode.

PE1

`PE1#configure terminal`	Enters the global configuration mode.
`PE1(config)#evpn mpls enable`	Enables EVPN-MPLS within the configuration.
`PE1(config)#evpn mpls irb`	Configures EVPN-MPLS IRB.
`PE1(config)#mac vrf MacVrf1`	Creates a MAC VRF named MacVrf1.

`PE1(config-vrf)# rd 2.2.2.2:100`	Defines a RD for MacVrf1 as 2.2.2.2:100.
`PE1(config-vrf)# route-target both 100:1`	Sets the route target for MacVrf1 to 100:1 for both import and export.
`PE1(config-vrf)#ip vrf IpVrf1`	Creates an IP VRF named IpVrf1.
`PE1(config-vrf)# rd 2.2.2.2:200`	Defines an RD for IpVrf1 as 2.2.2.2:200.
`PE1(config-vrf)# route-target both 100:100`	Sets the route target for IpVrf1 to 100:100 for both import and export.
`PE1(config-vrf)# l3vni 100`	Configures Layer 3 Virtual Network Identifier (L3VNI) as 100 for IpVrf1.
`PE1(config-vrf)#evpn mpls vtep-ip-global 2.2.2.2`	Configures the EVPN-MPLS Virtual Tunnel EndPoint (VTEP) IP address as 2.2.2.2.
`PE1(config)#evpn mpls id 101`	Defines EVPN-MPLS ID as 101.
`PE1(config-evpn-mpls)# host-reachability-protocol evpn-bgp MacVrf1`	Specifies BGP as the host reachability protocol for MacVrf1.
`PE1(config-evpn-mpls)# evpn irb irb101`	Associates IRB 101 with EVPN-MPLS.
`PE1(config-evpn-mpls)#evpn mpls id 1001`	Defines EVPN-MPLS ID as 1001.
`PE1(config-evpn-mpls)# host-reachability-protocol evpn-bgp MacVrf1`	Specifies BGP as the host reachability protocol for MacVrf1.
`PE1(config-evpn-mpls)# evpn irb irb1001`	Associates IRB 1001 with EVPN-MPLS.
`PE1(config-evpn-mpls)#qos enable`	Enables QoS features.
`PE1(config)#qos statistics`	Configures QoS statistics collection.
`PE1(config)#hostname PE1`	Sets the hostname of the device to PE1.
`PE1(config)#segment-routing`	Enters the configuration mode for segment routing.
`PE1(config-sr)# mpls sr-prefer`	Configures MPLS as the preferred segment routing method.
`PE1(config-sr)#exit`	Exits the segment routing configuration mode.
`PE1(config)#interface po1`	Enters the configuration mode for Port-Channel 1.
`PE1(config-if)# switchport`	Configures the interface as a switchport.
`PE1(config-if)# load-interval 30`	Sets the load interval to 30 seconds for the interface.
`PE1(config-if)#interface po1.101 switchport`	Enters the configuration mode for subinterface 101 of Port-Channel 1 and configures it as a switchport.
`PE1(config-if)# encapsulation dot1q 101`	Configures IEEE 802.1Q encapsulation with VLAN ID 101 for the subinterface.
`PE1(config-if)# rewrite pop`	Configures packet rewriting for the subinterface.
`PE1(config-if)# access-if-evpn`	Specifies that this interface is associated with EVPN.
`PE1(config-acc-if-evpn)# map vpn-id 101`	Maps this interface to VPN ID 101.
`PE1(config-acc-if-evpn)#interface po1.1001 switchport`	Enters the configuration mode for subinterface 1001 of Port-Channel 1 and configures it as a switchport.
`PE1(config-if)# encapsulation dot1q 1001`	Configures IEEE 802.1Q encapsulation with VLAN ID 1001 for the subinterface.
`PE1(config-if)# rewrite pop`	Configures packet rewriting for the subinterface.
`PE1(config-if)# access-if-evpn`	Specifies that this interface is associated with E
`PE1(config-acc-if-evpn)# map vpn-id 1001`	Maps this interface to VPN ID 1001.
`PE1(config-acc-if-evpn)#exit`	Exits the access-if-evpn configuration mode.
`PE1(config-if)#interface irb101`	Enters the configuration mode for IRB interface 101.

`PE1(config-irb-if)# ip vrf forwarding IpVrf1`	Associates the IRB interface with the IP VRF IpVrf1.
`PE1(config-irb-if)# ip address 101.101.101.1/24`	Configures the IP address 101.101.101.1 with a subnet mask of /24 on the IRB interface.
`PE1(config-irb-if)#interface irb1001`	Enters the configuration mode for IRB interface 1001.
`PE1(config-irb-if)# ip vrf forwarding IpVrf1`	Associates the IRB interface with the IP VRF IpVrf1.
`PE1(config-irb-if)# ipv6 address 1001::1/64`	Configures an IPv6 address 1001::1 with a subnet prefix length of /64 on the IRB interface.
`PE1(config-irb-if)# ipv6 router ospf area 0.0.0.0 tag IpVrf1 instance-id 0`	Configures OSPF for IPv6 on the IRB interface, specifying area 0.0.0.0 with a tag and instance ID.
`PE1(config-irb-if)#interface lo`	Enters the configuration mode for loopback interface.
`PE1(config-if)# ip address 127.0.0.1/8`	Configures the primary IP address of the loopback interface as 127.0.0.1 with a subnet mask of /8.
`PE1(config-if)# ip address 2.2.2.2/32 secondary`	Adds a secondary IP address 2.2.2.2 with a /32 subnet mask to the loopback interface.
`PE1(config-if)# ipv6 address ::1/128`	Configures an IPv6 address ::1 with a subnet prefix length of /128 on the loopback interface.
`PE1(config-if)# prefix-sid index 2`	Configures a prefix SID with an index of 2.
`PE1(config-if)#interface xe7`	Enters the configuration mode for Ten Gigabit Ethernet interface 7.
`PE1(config-if)# channel-group 1 mode active`	Configures the interface to be part of an EtherChannel group with group number 1 and sets the mode to active.
`PE1(config-if)#interface xe10`	Enters the configuration mode for Ten Gigabit Ethernet interface 10.
`PE1(config-if)# speed 10g`	Sets the speed of the interface to 10 gigabits per second.
`PE1(config-if)# ip address 10.1.1.1/24`	Configures an IP address 10.1.1.1 with a subnet mask of /24 on the interface.
`PE1(config-if)# label-switching`	Enables label switching on the interface.
`PE1(config-if)# exit`	Exits the interface configuration mode.
`PE1(config)#router ospf 100`	Enters the configuration mode for OSPF routing process 100.
`PE1(config-router)# ospf router-id 2.2.2.2`	Sets the OSPF router ID to 2.2.2.2.
`PE1(config-router)# network 2.2.2.2/32 area 0.0.0.0`	Configures the network 2.2.2.2/32 within OSPF Area 0.0.0.0.
`PE1(config-router)# network 10.1.1.0/24 area 0.0.0.0`	Configures the network 10.1.1.0/24 within OSPF Area 0.0.0.0.
`PE1(config-router)# segment-routing mpls`	Configures MPLS for segment routing.
`PE1(config-router)#router ospf 101 IpVrf1`	Enters the configuration mode for OSPF routing process 101 within VRF IpVrf.
`PE1(config-router)# network 10.1.1.0/24 area 0.0.0.0`	Specifies the network 10.1.1.0 with a /24 subnet mask to be included in OSPF Area 0.0.0.0.
`PE1(config-router)# segment-routing mpls`	Enables MPLS-based segment routing in the OSPF configuration.
`PE1(config-router)#router ospf 101 IpVrf1`	Enters the configuration mode for OSPF routing process 101 associated with VRF IpVrf1.
`PE1(config-router)# ospf router-id 2.2.2.2`	Sets the OSPF router ID to 2.2.2.2.
`PE1(config-router)# redistribute bgp`	Configures the redistribution of BGP routes into OSPF.
`PE1(config-router)# network 101.101.101.0/24 area 0.0.0.0`	Specifies the network 101.101.101.0 with a /24 subnet mask to be included in OSPF Area 0.0.0.0 for VRF IpVrf1.

PE1(config-router)#router ipv6 vrf ospf IpVrf1	Enters the configuration mode for OSPFv3 (OSPF for IPv6) routing process associated with VRF IpVrf1.
PE1(config-router)# router-id 2.2.2.2	Sets the OSPFv3 router ID to 2.2.2.2.
PE1(config-router)#router bgp 100	Enters the configuration mode for BGP routing process 100.
PE1(config-router)# bgp router-id 2.2.2.2	Sets the BGP router ID to 2.2.2.2.
PE1(config-router)# neighbor 5.5.5.5 remote-as 100	Specifies a BGP neighbor with IP address 5.5.5.5 and assigns it the remote autonomous system number 100.
PE1(config-router)# neighbor 5.5.5.5 update-source lo	Sets the loopback interface as the source for BGP updates to the neighbor.
PE1(config-router)# neighbor 5.5.5.5 advertisement-interval 0	Configures the BGP neighbor's advertisement interval to 0.
PE1(config-router)# address-family l2vpn evpn	Enters the configuration mode for the BGP address family l2vpn evpn.
PE1(config-router-af)# neighbor 5.5.5.5 activate	Activates the BGP neighbor within the l2vpn evpn address family.
PE1(config-router-af)# exit-address-family	Exits the l2vpn evpn address family configuration.
PE1(config-router)# address-family ipv4 vrf IpVrf1	Enters the configuration mode for the BGP address family ipv4 vrf IpVrf1.
PE1(config-router-af)# redistribute ospf 101	Configures the redistribution of OSPF routes from OSPF process 101 into the ipv4 vrf IpVrf1 address family.
PE1(config-router-af)# exit-address-family	Exits the ipv4 vrf IpVrf1 address family configuration.
PE1(config-router)# address-family ipv6 vrf IpVrf1	Enters the configuration mode for the BGP address family ipv6 vrf IpVrf1.
PE1(config-router-af)# redistribute ospf	Configures the redistribution of OSPFv3 routes into the ipv6 vrf IpVrf1 address family.
PE1(config-router-af)# exit-address-family	Exits the ipv6 vrf IpVrf1 address family configuration.
PE1(config-router)#line console 0	Enters the configuration mode for the console line.
PE1(config-line)# exec-timeout 0 0	Sets the console line's execution timeout values to 0 seconds.
PE1(config-line)#line vty 0 871	Enters the configuration mode for virtual terminal lines 0 to 871.
PE1(config-line)# exec-timeout 0 0	Sets the execution timeout values for virtual terminal lines to 0 seconds.
PE1(config-line)# privilege level 16	Specifies the privilege level for users accessing the virtual terminal lines.
PE1(config-line)#commit	Commits the configuration changes.
PE1(config-line)#end	Exits the configuration mode.

Configure MTU over IRB

PE1(config)#interface irb101	Enters interface configuration mode for IRB interface 101.
PE1(config-irb-if)#mtu 9216	Sets the MTU to 9216 bytes.
PE1(config-irb-if)#commit	Commit the changes made to the interface configuration.
PE1(config-irb-if)#exit	Exits the IRB interface configuration mode and returns to the global configuration mode.
PE1(config)#interface irb1001	Enters interface configuration mode for IRB interface 1001.
PE1(config-irb-if)#mtu 9216	Sets the MTU to 9216 bytes.
PE1(config-irb-if)#commit	Commit the changes made to the interface configuration.

`PE1(config-irb-if)#exit`	Exits the IRB interface configuration mode and returns to the global configuration mode.
`PE1(config)#interface int po1`	Enters interface configuration mode for IRB int po1.
`PE1(config-irb-if)#mtu 9216`	Sets the MTU to 9216 bytes.
`PE1(config-irb-if)#commit`	Commit the changes made to the interface configuration.
`PE1(config-irb-if)#exit`	Exits the IRB interface configuration mode and returns to the global configuration mode.

P1

`P1#configure terminal`	Enters the global configuration mode.
`P1(config)#segment-routing`	Enters the configuration mode for Segment Routing.
`P1(config-sr)# mpls sr-prefer`	Configures the preference for MPLS as part of the Segment Routing setup.
`P1(config-if)#interface ce50`	Enters the configuration mode for interface ce50.
`P1(config-if)# ip address 20.1.1.1/24`	Assigns an IP address of 20.1.1.1 with a /24 subnet mask to interface ce50.
`P1(config-if)# label-switching`	Enables label switching for the interface ce50.
`P1(config-if)#interface lo`	Enters the configuration mode for a loopback interface.
`P1(config-if)# ip address 127.0.0.1/8`	Assigns an IP address of 127.0.0.1 with a /8 subnet mask to the loopback interface.
`P1(config-if)# ip address 3.3.3.3/32 secondary`	Adds a secondary IP address of 3.3.3.3 with a /32 subnet mask to the loopback interface.
`P1(config-if)# ipv6 address ::1/128`	Assigns an IPv6 address of ::1/128 to the loopback interface.
`P1(config-if)# prefix-sid index 3`	Configures a prefix SID with an index of 3 for the loopback interface.
`P1(config-if)#interface xe3`	Enters the configuration mode for Ten Gigabit Ethernet interface 3.
`P1(config-if)# ip address 10.1.1.2/24`	Assigns an IP address of 10.1.1.2 with a /24 subnet mask to interface xe3.
`P1(config-if)# label-switching`	Enables label switching for interface xe3.
`P1(config-if)# exit`	Exits the interface configuration.
`P1(config-router)#router ospf 100`	Enters the configuration mode for OSPF routing process 100.
`P1(config-router)# ospf router-id 3.3.3.3`	Sets the OSPF router ID to 3.3.3.3.
`P1(config-router)# network 3.3.3.3/32 area 0.0.0.0`	Specifies the network 3.3.3.3 with a /32 subnet mask to be included in OSPF Area 0.0.0.0.
`P1(config-router)# network 10.1.1.0/24 area 0.0.0.0`	Specifies the network 10.1.1.0 with a /24 subnet mask to be included in OSPF Area 0.0.0.0.
`P1(config-router)# network 20.1.1.0/24 area 0.0.0.0`	Specifies the network 20.1.1.0 with a /24 subnet mask to be included in OSPF Area 0.0.0.0.
`P1(config-router)# segment-routing mpls`	Enables MPLS-based segment routing in the OSPF configuration.

`P1(config-router)#line console 0`	Enters the configuration mode for the console line.
`P1(config-line)# exec-timeout 0 0`	Sets the console line's execution timeout values to 0 seconds.
`P1(config-line)#line vty 0 871`	Enters the configuration mode for virtual terminal lines 0 to 871.
`P1(config-line)# exec-timeout 0 0`	Sets the execution timeout values for virtual terminal lines to 0 seconds.
`P1(config-line)# privilege level 16`	Specifies the privilege level for users accessing the virtual terminal lines.
`P1(config-line)# Exit`	Exits from configuration mode.

P2

`P2#configure terminal`	Enters the global configuration mode.
`P2(config)#segment-routing`	Enters the configuration mode for Segment Routing.
`P2(config-sr)# mpls sr-prefer`	Configures the preference for MPLS as part of the Segment Routing setup.
`P2(config-sr)#interface ce1`	Enters the configuration mode for interface ce1.
`P2(config-if)# ip address 20.1.1.2/24`	Assigns an IP address of 20.1.1.2 with a /24 subnet mask to interface ce1.
`P2(config-if)# label-switching`	Enables label switching for interface ce1.
`P2(config-if)#interface ce2`	Enters the configuration mode for interface ce2.
`P2(config-if)# speed 40g`	Sets the speed of interface ce2 to 40Gbps.
`P2(config-if)# ip address 30.1.1.1/24`	Assigns an IP address of 30.1.1.1 with a /24 subnet mask to interface ce2.
`P2(config-if)# label-switching`	Enables label switching for interface ce2.
`P2(config-if)#interface lo`	Enters the configuration mode for a loopback interface.
`P2(config-if)# ip address 127.0.0.1/8`	Assigns an IP address of 127.0.0.1 with a /8 subnet mask to the loopback interface.
`P2(config-if)# ip address 4.4.4.4/32 secondary`	Adds a secondary IP address of 4.4.4.4 with a /32 subnet mask to the loopback interface.
`P2(config-if)# ipv6 address ::1/128`	Assigns an IPv6 address of ::1/128 to the loopback interface.
`P2(config-if)# prefix-sid index 4`	Configures a prefix SID with an index of 4 for the loopback interface.
`P2(config)#router ospf 100`	Enters the configuration mode for OSPF routing process 100.
`P2(config-router)# ospf router-id 4.4.4.4`	Sets the OSPF router ID to 4.4.4.4.
`P2(config-router)# network 4.4.4.4/32 area 0.0.0.0`	Specifies the network 4.4.4.4 with a /32 subnet mask to be included in OSPF Area 0.0.0.0.
`P2(config-router)# network 20.1.1.0/24 area 0.0.0.0`	Specifies the network 20.1.1.0 with a /24 subnet mask to be included in OSPF Area 0.0.0.0.
`P2(config-router)# network 30.1.1.0/24 area 0.0.0.0`	Specifies the network 30.1.1.0 with a /24 subnet mask to be included in OSPF Area 0.0.0.0.
`P2(config-router)# segment-routing mpls`	Enables MPLS-based segment routing in the OSPF configuration.
`P2(config-router)#line console 0`	Enters the configuration mode for the console line.
`P2(config-line)# exec-timeout 0 0`	Sets the console line's execution timeout values to 0 seconds.

`P2(config-line)#line vty 0 871`	Enters the configuration mode for virtual terminal lines 0 to 871.
`P2(config-line)# exec-timeout 0 0`	Sets the execution timeout values for virtual terminal lines to 0 seconds.
`P2(config-line)# privilege level 16`	Specifies the privilege level for users accessing the virtual terminal lines.
`P2(config-line)#Exit`	Exits the configuration.

PE2

`PE2#configure terminal`	Enters the global configuration mode.
`PE2(config)#evpn mpls enable`	Enables EVPN MPLS.
`PE2(config)#evpn mpls irb`	Enables EVPN MPLS IRB.
`PE2(config)mac vrf MacVrf2`	Enters the configuration mode for MAC VRF with the name MacVrf2.
`PE2(config-vrf)# rd 5.5.5.5:100`	Sets the route distinguisher for the MAC VRF to 5.5.5.5:100.
`PE2(config-vrf)# route-target both 100:1`	Specifies route targets for import and export for the MAC VRF.
`PE2(config-vrf)#ip vrf IpVrf1`	Enters the configuration mode for IP VRF with the name IpVrf1.
`PE2(config-vrf)# rd 5.5.5.5:200`	Sets the route distinguisher for the IP VRF to 5.5.5.5:200.
`PE2(config-vrf)# route-target both 100:100`	Specifies route targets for import and export for the IP VRF.
`PE2(config-vrf)# l3vni 100`	Configures an L3VNI with ID 100 for the IP VRF.
`PE2(config-vrf)#evpn mpls vtep-ip-global 5.5.5.5`	Sets the global VTEP IP address for EVPN MPLS to 5.5.5.5.
`PE2(config)#evpn mpls id 201`	Configures an EVPN MPLS instance with ID 201.
`PE2(config-evpn-mpls)# host-reachability-protocol evpn-bgp MacVrf2`	Specifies BGP as the host reachability protocol for EVPN MPLS using MacVrf2.
`PE2(config-evpn-mpls)# evpn irb irb101`	Configures EVPN MPLS IRB with the IRB interface named irb101.
`PE2(config-evpn-mpls)#evpn mpls id 2001`	Configures another EVPN MPLS instance with ID 2001.
`PE2(config-evpn-mpls)# host-reachability-protocol evpn-bgp MacVrf2`	Specifies BGP as the host reachability protocol for EVPN MPLS using MacVrf2.
`PE2(config-evpn-mpls)# evpn irb irb1001`	Configures EVPN MPLS IRB with the IRB interface named irb1001.
`PE2(config-evpn-mpls)#qos enable`	Enables QoS features for EVPN MPLS.
`PE2(config)qos statistics`	Configures statistics for QoS.
`PE2(config)#hostname PE2`	Sets the hostname of the device to PE2.
`PE2(config)#segment-routing`	Enters the configuration mode for Segment Routing.
`PE2(config-sr)# mpls sr-prefer`	Configures the preference for MPLS in Segment Routing.
`PE2(config)#interface ce1`	Enters the configuration mode for interface ce1.
`PE2(config-if)# switchport`	Enables the switchport mode for the interface.
`PE2(config-if)# load-interval 30`	Sets the load interval to 30 seconds for the interface.
`PE2(config)#interface ce1.101 switchport`	Enters the configuration mode for subinterface ce1.101 and enables the switchport mode.
`PE2(config-if)# encapsulation dot1q 101`	Specifies VLAN encapsulation with VLAN ID 101 for the subinterface.

`PE2(config-if)# rewrite pop`	Configures the rewrite operation as "pop" for the subinterface.
`PE2(config-if)# access-if-evpn`	Configures the subinterface as an access interface for EVPN.
`PE2(config-acc-if-evpn)# map vpn-id 201`	Maps the subinterface to EVPN with VPN ID 201.
`PE2(config-acc-if-evpn)#exit`	Exits the access interface configuration.
`PE2(config)#interface ce3`	Enters the configuration mode for interface ce3.
`PE2(config)# ip address 30.1.1.2/24`	Assigns the IP address 30.1.1.2 with a subnet mask of 255.255.255.0 to interface ce3.
`PE2(config)# label-switching`	Enables label-switching on interface ce3.
`PE2(config-if)#interface irb101`	Enters the configuration mode for interface irb101.
`PE2(config-irb-if)# ip vrf forwarding IpVrf1`	Associates interface irb101 with the VRF (Virtual Routing and Forwarding) instance IpVrf1.
`PE2(config-irb-if)# ip address 201.201.201.1/24`	Assigns the IP address 201.201.201.1 with a subnet mask of 255.255.255.0 to interface irb101.
`PE2(config-irb-if)# ip router isis 101`	Enables ISIS routing protocol on interface irb101.
`PE2(config-irb-if)#interface irb1001`	Enters the configuration mode for interface irb1001.
`PE2(config-irb-if)# ip vrf forwarding IpVrf1`	Associates interface irb1001 with the VRF instance IpVrf1.
`PE2(config-irb-if)# ipv6 address 2001::1/64`	Assigns the IPv6 address 2001::1 with a prefix length of 64 to interface irb1001.
`PE2(config-irb-if)# ipv6 router isis 101`	Enables ISIS routing protocol for IPv6 on interface irb1001.
`PE2(config-irb-if)#interface lo`	Enters the configuration mode for the loopback interface.
`PE2(config-if)# ip address 127.0.0.1/8`	Assigns the IP address 127.0.0.1 with a subnet mask of 255.0.0.0 to the loopback interface.
`PE2(config-if)# ip address 5.5.5.5/32 secondary`	Assigns the secondary IP address 5.5.5.5 with a subnet mask of 255.255.255.255 to the loopback interface.
`PE2(config-if)# ipv6 address ::1/128`	Assigns the IPv6 address ::1 with a prefix length of 128 to the loopback interface.
`PE2(config-if)# prefix-sid index 5`	Configures a prefix SID with an index of 5.
`PE2(config-if)#exit`	Exits the interface configuration mode.
`PE2(config)#router ospf 100`	Enters the configuration mode for OSPF with process ID 100.
`PE2(config-router)# ospf router-id 5.5.5.5`	Sets the OSPF router ID to 5.5.5.5.
`PE2(config-router)# network 5.5.5.5/32 area 0.0.0.0`	Advertises the network 5.5.5.5/32 into OSPF area 0.0.0.0.
`PE2(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.
`PE2(config-router)# segment-routing mpls`	Enables MPLS for segment routing.
`PE2(config-router)#router isis 101 IpVrf1`	Enters the configuration mode for ISIS with process ID 101 and VRF instance IpVrf1.
`PE2(config-irb-if)#interface lo`	Enters the configuration mode for the loopback interface.
`PE2(config-if)# ip address 127.0.0.1/8`	Assigns the IP address 127.0.0.1 with a subnet mask of 255.0.0.0 to the loopback interface.
`PE2(config-if)# ip address 5.5.5.5/32 secondary`	Assigns the secondary IP address 5.5.5.5 with a subnet mask of 255.255.255.255 to the loopback interface.
`PE2(config-if)# ipv6 address ::1/128`	Assigns the IPv6 address ::1 with a prefix length of 128 to the loopback interface.
`PE2(config-if)# prefix-sid index 5`	Configures a prefix SID with an index of 5.

`PE2(config)#router ospf 100`	Enters the configuration mode for OSPF with process ID 100.
`PE2(config-router)# ospf router-id 5.5.5.5`	Sets the OSPF router ID to 5.5.5.5.
`PE2(config-router)# network 5.5.5.5/32 area 0.0.0.0`	Advertises the network 5.5.5.5/32 into OSPF area 0.0.0.0.
`PE2(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.
`PE2(config-router)# segment-routing mpls`	Enables MPLS for segment routing.
`PE2(config-router)#router isis 101 IpVrf1`	Enters the configuration mode for ISIS with process ID 101 and VRF instance IpVrf1.
`PE2(config-router)# is-type level-2-only`	Configures ISIS as a level-2-only routing protocol.
`PE2(config-router)# net 49.0001.0000.0000.0005.00`	Sets the ISIS NET for this router.
`PE2(config-router)# redistribute bgp`	Configures redistribution of BGP routes into ISIS.
`PE2(config-router)#router bgp 100`	Enters the configuration mode for BGP with AS number 100.
`PE2(config-router)# bgp router-id 5.5.5.5`	Sets the BGP router ID to 5.5.5.5.
`PE2(config-router)# neighbor 2.2.2.2 remote-as 100`	Configures a BGP neighbor with the remote AS number 100 and IP address 2.2.2.2.
`PE2(config-router)# neighbor 2.2.2.2 update-source lo`	Specifies the loopback interface as the source for BGP updates to the neighbor.
`PE2(config-router)# neighbor 2.2.2.2 advertisement-interval 0`	Sets the BGP advertisement interval to 0.
`PE2(config-router)# address-family l2vpn evpn`	Enters the address-family configuration mode for L2VPN EVPN.
`PE2(config-router-af)# neighbor 2.2.2.2 activate`	Activates the BGP neighbor for the L2VPN EVPN address family.
`PE2(config-router-af)# exit-address-family`	Exits the address-family configuration.
`PE2(config-router)# address-family ipv4 vrf IpVrf1`	Enters the address-family configuration mode for IPv4 with VRF instance IpVrf1.
`PE2(config-router-af)# redistribute isis`	Configures redistribution of ISIS routes into the IPv4 address family.
`PE2(config-router-af)# exit-address-family`	Exits the address-family configuration.

Configure MTU over IRB

`PE2(config)#interface irb101`	Enters interface configuration mode for IRB interface 101.
`PE2(config-irb-if)#mtu 9216`	Sets the MTU to 9216 bytes.
`PE2(config-irb-if)#commit`	Commit the changes made to the interface configuration.
`PE2(config)#interface irb1001`	Enters interface configuration mode for IRB interface 1001.
`PE2(config-irb-if)#mtu 9216`	Sets the MTU to 9216 bytes.
`PE2(config-irb-if)#commit`	Commit the changes made to the interface configuration.
`PE2(config-irb-if)#exit`	Exits the IRB interface configuration mode and returns to the global configuration mode.
`PE2(config)#interface ce1`	Enters interface configuration mode for IRB interface ce1.
`PE2(config-irb-if)#mtu 9216`	Sets the MTU to 9216 bytes.

`PE2(config-irb-if)#commit`	Commit the changes made to the interface configuration.
`PE2(config-irb-if)#exit`	Exits the IRB interface configuration mode and returns to the global configuration mode.

CE2

`CE2#Configure Terminal`	Enters the configuration mode.
`CE2(config)#interface ce1`	Enters the configuration mode for interface ce1.
`CE2(config-if)# load-interval 30`	Sets the load interval for the interface to 30 seconds.
`CE2(config-if)#interface ce1.101`	Enters the configuration mode for subinterface ce1.101.
`CE2(config-if)# encapsulation dot1q 101`	Configures 802.1Q encapsulation with VLAN ID 101 for the subinterface.
`CE2(config-if)# ip address 201.201.201.2/24`	Assigns the IPv4 address 201.201.201.2 with a subnet mask of /24 to the subinterface.
`CE2(config-if)# ip router isis 101`	Specifies ISIS as the routing protocol for IPv4.
`CE2(config-if)#interface ce1.1001`	Enters the configuration mode for subinterface ce1.1001.
`CE2(config-if)# encapsulation dot1q 1001`	Configures 802.1Q encapsulation with VLAN ID 1001 for the subinterface.
`CE2(config-if)# ipv6 address 2001::2/64`	Assigns the IPv6 address 2001::2 with a /64 prefix length to the subinterface.
`CE2(config-if)# ipv6 router isis 101`	Specifies ISIS as the routing protocol for IPv6.
`CE2(config-if)#interface lo`	Enters the configuration mode for the loopback interface.
`CE2(config-if)# ip address 127.0.0.1/8`	Assigns the primary IPv4 address 127.0.0.1 to the loopback interface with a subnet mask of /8.
`CE2(config-if)# ip address 6.6.6.6/32 secondary`	Adds a secondary IPv4 address, 6.6.6.6 with a /32 subnet mask, to the loopback interface.
`CE2(config-if)# ipv6 address ::1/128`	Assigns the primary IPv6 address ::1 to the loopback interface with a /128 prefix length.
`CE2(config)#router isis 101`	Enters the configuration mode for ISIS with process ID 101.
`CE2(config-router)# is-type level-2-only`	Configures ISIS as a level-2-only routing protocol.
`CE2(config-router)# net 49.0001.0000.0000.0006.00`	Sets the ISIS NET for this router.
`CE2(config-router)# address-family ipv6`	Enters the address-family configuration mode for IPv6.
`CE2(config-router-af)# exit-address-family`	Exits the address-family configuration for IPv6.
`CE2(config)#line console 0`	Enters the configuration mode for the console line.
`CE2(config-line)# exec-timeout 0 0`	Sets the console line's exec-timeout to 0.
`CE2(config-line)#line vty 0 871`	Enters the configuration mode for VTY lines 0 through 871.
`CE2(config-router-af)# exit`	Exits the from configuration mode.

Configure MTU over IRB

`CE2(config)#int ce1`	Enters interface configuration mode for int ce1.
`CE2(config-irb-if)#mtu 9216`	Sets the MTU to 9216 bytes.
`CE2(config-irb-if)#commit`	Commit the changes made to the interface configuration.
`CE2(config-irb-if)#exit`	Exits the IRB interface configuration mode and returns to the global configuration mode.

`CE2(config)#int  ce1.101`	Enters interface configuration mode for int ce1.101.
`CE2(config-irb-if)#mtu 9216`	Sets the MTU to 9216 bytes.
`CE2(config-irb-if)#commit`	Commit the changes made to the interface configuration.
`CE2(config-irb-if)#exit`	Exits the IRB interface configuration mode and returns to the global configuration mode.
`CE2(config)#int  ce1.1001`	Enters interface configuration mode for int ce1.1001.
`CE2(config-irb-if)#mtu 9216`	Sets the MTU to 9216 bytes.
`CE2(config-irb-if)#commit`	Commit the changes made to the interface configuration.
`CE2(config-irb-if)#exit`	Exits the IRB interface configuration mode and returns to the global configuration mode.

Implementation Examples

Scenario: Configure OSPF and ISIS protocols on an IRB interface with an assigned IP address.

Validation

```
CE1#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 101 VRF(default):
Neighbor ID     Pri   State            Dead Time   Address         Interface
Instance ID
2.2.2.2           1   Full/DR         00:00:37    101.101.101.1   po1.101
0
CE1#
CE1#
CE1#show ipv6 ospf neighbor

Total number of full neighbors: 1
OSPFv3 Process (1001)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
2.2.2.2           1   Full/DR         00:00:34    po1.1001   0
CE1#
CE1#
CE1#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            1.1.1.1/32 is directly connected, lo, 00:16:58
```

```
C            99.99.99.0/24 is directly connected, ge2, 00:16:59
C            101.101.101.0/24 is directly connected, po1.101, 00:14:10
C            127.0.0.0/8 is directly connected, lo, 00:18:23
O E2         201.201.201.0/24 [110/1] via 101.101.101.1, po1.101, 00:06:14

Gateway of last resort is not set
CE1#
CE1#
CE1#show ipv6 route
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, 00:18:30
C       1001::/64 via ::, po1.1001, 00:14:17
O E2    2001::/64 [110/20] via fe80::eac5:7aff:fea8:7cb3, po1.1001, 00:03:50
C       fe80::/64 via ::, po1.1001, 00:14:17
PE1  :
PE1#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
3.3.3.3           1   Full/DR         00:00:38    10.1.1.2       xe10                    0

Total number of full neighbors: 1
OSPF process 101 VRF(IpVrf1):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
1.1.1.1           1   Full/Backup     00:00:32    101.101.101.2   irb101
0

PE1#show ipv6 ospf neighbor

Total number of full neighbors: 1
OSPFv3 Process (IpVrf1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
1.1.1.1           1   Full/Backup     00:00:33    irb1001    0

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status        Up/Down       Update        evpn-id
================================================================================
2.2.2.2          5.5.5.5          Installed     00:24:13      00:24:13      100
```

```
Total number of entries are 1

PE1#show bgp l2vpn evpn summary
BGP router identifier 2.2.2.2, local AS number 100
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
5.5.5.5                 4   100   20        21      1     0     0  00:04:09
9      0      5      2      0      2

Total number of neighbors 1

Total number of Established sessions 1

PE1#show evpn mpls
EVPN-MPLS Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VPN-ID   EVI-Name     EVI-Type Type Interface ESI                          VLAN      DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
100      ----         L3       NW   ----      ----                          ----      -
---      2.2.2.2          5.5.5.5
101      ----         --       AC   po1.101   --- Single Homed Port ---     ----      -
---      ----             ----
1001     ----         --       AC   po1.1001  --- Single Homed Port ---     ----      -
---      ----             ----

Total number of entries are 3

Note: Refer sub-interface config for VLAN information.
PE1#
PE1#
PE1#show bgp l2vpn evpn prefix-route

RD[5.5.5.5:200]
ESI             Eth-Tag Prefix-Length    IP-Address      GW-IPAddress    L3VNID/LABEL
Nexthop          Encap     Router-Mac
0               0       24               201.201.201.0   0.0.0.0         16
5.5.5.5          MPLS      e49d:73b1:c301
0               0       64               2001::          ::              16
5.5.5.5          MPLS      e49d:73b1:c301
PE1#
PE1#show bgp l2vpn evpn mac-ip

RD[2.2.2.2:100] VRF[MacVrf1]:
```

```
ESI                           Eth-Tag    Mac-Address    IP-Address
VNID/LABEL     L3VNID    Nexthop         GW-Type          Encap
0                          101        9819:2ccd:9320 --
17             0         2.2.2.2         --               MPLS
0                          101        9819:2ccd:9320 101.101.101.2
17             0         2.2.2.2         --               MPLS
0                          101        e8c5:7aa8:7cb3 101.101.101.1
17             0         2.2.2.2         --               MPLS
0                          201        e49d:73b1:c301 201.201.201.1
17             0         5.5.5.5         --               MPLS
0                          201        e8c5:7a76:583b --
17             0         5.5.5.5         --               MPLS
0                          201        e8c5:7a76:583b 201.201.201.2
17             0         5.5.5.5         --               MPLS
0                          1001       9819:2ccd:9320 --
18             0         2.2.2.2         --               MPLS
0                          1001       e8c5:7aa8:7cb3 1001::1
18             0         2.2.2.2         --               MPLS
0                          2001       e49d:73b1:c301 2001::1
18             0         5.5.5.5         --               MPLS
0                          2001       e8c5:7a76:583b --
18             0         5.5.5.5         --               MPLS

RD[5.5.5.5:100]
ESI                           Eth-Tag    Mac-Address    IP-Address
VNID/LABEL     L3VNID    Nexthop         GW-Type          Encap
0                          201        e49d:73b1:c301 201.201.201.1
17             0         5.5.5.5         --               MPLS
0                          201        e8c5:7a76:583b --
17             0         5.5.5.5         --               MPLS
0                          201        e8c5:7a76:583b 201.201.201.2
17             0         5.5.5.5         --               MPLS
0                          2001       e49d:73b1:c301 2001::1
18             0         5.5.5.5         --               MPLS
0                          2001       e8c5:7a76:583b --
18             0         5.5.5.5         --               MPLS

PE1#show mpls ftn-table
 Primary FTN entry with FEC: 3.3.3.3/32, id: 1, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: OSPF-SR, distance: 110, Action-type: Redirect to Tunnel, Exp-bits: 0x0,
Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, QoS Resource id: 0, Description:
N/A, , Color: 0
     Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 2
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe10, out label: 3
    Nexthop addr: 10.1.1.2        cross connect ix: 3, op code: Push


 Primary FTN entry with FEC: 4.4.4.4/32, id: 2, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: OSPF-SR, distance: 110, Action-type: Redirect to Tunnel, Exp-bits: 0x0,
Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, QoS Resource id: 0, Description:
N/A, , Color: 0
```

```
     Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
      Owner: OSPF-SR, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 4, owner: OSPF-SR, Stale: NO, out intf: xe10, out label:
16004
    Nexthop addr: 10.1.1.2        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: OSPF-SR, distance: 110, Action-type: Redirect to Tunnel, Exp-bits: 0x0,
Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, QoS Resource id: 0, Description:
N/A, , Color: 0
     Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 6
      Owner: OSPF-SR, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 6, owner: OSPF-SR, Stale: NO, out intf: xe10, out label:
16005
    Nexthop addr: 10.1.1.2        cross connect ix: 5, op code: Push

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled
Code    FEC               FTN-ID    Nhlfe-ID  Tunnel-id   Pri   LSP-Type     Out-Label
Out-Intf    ELC   Nexthop
   O>   3.3.3.3/32         1         3         0          Yes   LSP_DEFAULT  3
xe10        No    10.1.1.2
   O>   4.4.4.4/32         2         5         0          Yes   LSP_DEFAULT  16004
xe10        No    10.1.1.2
   O>   5.5.5.5/32         3         7         0          Yes   LSP_DEFAULT  16005
xe10        No    10.1.1.2

PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

ILM-ECMP LDP: Disabled
Code    FEC/VRF/L2CKT    ILM-ID     In-Label    Out-Label   In-Intf    Out-Intf/VRF
Nexthop                  pri    LSP-Type
   B>   evpn:1001          5          18          Nolabel     N/A        N/A
127.0.0.1                Yes    LSP_DEFAULT
   B>   evpn:101           3          17          Nolabel     N/A        N/A
127.0.0.1                Yes    LSP_DEFAULT
   B>   IpVrf1             1          16          Nolabel     N/A        N/A             N/
A                      Yes    LSP_DEFAULT
```

```
   B>   evpn:101          2          26880       Nolabel     N/A        N/A
127.0.0.1               Yes    LSP_DEFAULT
   O>   4.4.4.4/32        8          16004       16004       N/A        xe10
10.1.1.2                Yes    LSP_DEFAULT
   O>   3.3.3.3/32        6          16003       3           N/A        xe10
10.1.1.2                Yes    LSP_DEFAULT
   O>   5.5.5.5/32        9          16005       16005       N/A        xe10
10.1.1.2                Yes    LSP_DEFAULT
   B>   evpn:1001         4          26881       Nolabel     N/A        N/A
127.0.0.1               Yes    LSP_DEFAULT
   O>   10.1.1.2/32       7          27520       3           N/A        xe10
10.1.1.2                Yes    LSP_DEFAULT

PE2   :
=========

PE2#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State            Dead Time   Address         Interface
Instance ID
4.4.4.4           1   Full/DR         00:00:28    30.1.1.1        ce3                    0
PE2#
PE2#
PE2#
PE2#show clns is-neighbors

Tag 101:  VRF : IpVrf1
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0006 irb101      Up     L2   64        0000.0000.0006.01
0000.0000.0006 irb1001     Up     L2   64        0000.0000.0006.02
PE2#
PE2#
PE2#
PE2#
PE2#show clns is-neighbors detail

Tag 101:  VRF : IpVrf1
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0006 irb101      Up     L2   64        0000.0000.0006.01
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 00:11:21
  Area Address(es): 49.0001
  IP Address(es): 201.201.201.2
  Level-2 Protocols Supported: IPv4, IPv6
  Adjacency advertisement: Advertise

0000.0000.0006 irb1001     Up     L2   64        0000.0000.0006.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
```

```
  Uptime: 00:11:21
  Area Address(es): 49.0001
  IPv6 Address(es): fe80::eac5:7aff:fe76:583b
  Level-2 Protocols Supported: IPv4, IPv6
  Adjacency advertisement: Advertise

PE2#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status        Up/Down       Update        evpn-id
================================================================================
5.5.5.5          2.2.2.2          Installed     00:24:13      00:24:13      100

Total number of entries are 1
PE2#
PE2#
PE2#
PE2#show bgp l2vpn evpn summary
BGP router identifier 5.5.5.5, 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   93        94      3      0      0  00:14:47
12      0      5      2      0      5

Total number of neighbors 1

Total number of Established sessions 1
PE2#
PE2#
PE2#
PE2#show bgp l2vpn evpn prefix-route

RD[2.2.2.2:200]
ESI             Eth-Tag Prefix-Length    IP-Address      GW-IPAddress    L3VNID/LABEL
Nexthop         Encap     Router-Mac
0               0       24               99.99.99.0      0.0.0.0         16
2.2.2.2         MPLS      e8c5:7aa8:7cb3
0               0       24               101.101.101.0   0.0.0.0         16
2.2.2.2         MPLS      e8c5:7aa8:7cb3
0               0       32               1.1.1.1         0.0.0.0         16
2.2.2.2         MPLS      e8c5:7aa8:7cb3
0               0       64               1001::          ::              16
2.2.2.2         MPLS      e8c5:7aa8:7cb3
0               0       64               9001::          ::              16
2.2.2.2         MPLS      e8c5:7aa8:7cb3
PE2#show bgp l2vpn evpn mac-ip

RD[2.2.2.2:100]
```

```
ESI                          Eth-Tag    Mac-Address    IP-Address
VNID/LABEL     L3VNID    Nexthop         GW-Type          Encap
0                            101        9819:2ccd:9320 --
17             0         2.2.2.2          --               MPLS
0                            101        9819:2ccd:9320 101.101.101.2
17             0         2.2.2.2          --               MPLS
0                            101        e8c5:7aa8:7cb3 101.101.101.1
17             0         2.2.2.2          --               MPLS
0                            1001       9819:2ccd:9320 --
18             0         2.2.2.2          --               MPLS
0                            1001       e8c5:7aa8:7cb3 1001::1
18             0         2.2.2.2          --               MPLS

RD[5.5.5.5:100] VRF[MacVrf2]:
ESI                          Eth-Tag    Mac-Address    IP-Address
VNID/LABEL     L3VNID    Nexthop         GW-Type          Encap
0                            101        9819:2ccd:9320 --
17             0         2.2.2.2          --               MPLS
0                            101        9819:2ccd:9320 101.101.101.2
17             0         2.2.2.2          --               MPLS
0                            101        e8c5:7aa8:7cb3 101.101.101.1
17             0         2.2.2.2          --               MPLS
0                            201        e49d:73b1:c301 201.201.201.1
17             0         5.5.5.5          --               MPLS
0                            201        e8c5:7a76:583b --
17             0         5.5.5.5          --               MPLS
0                            201        e8c5:7a76:583b 201.201.201.2
17             0         5.5.5.5          --               MPLS
0                            1001       9819:2ccd:9320 --
18             0         2.2.2.2          --               MPLS
0                            1001       e8c5:7aa8:7cb3 1001::1
18             0         2.2.2.2          --               MPLS
0                            2001       e49d:73b1:c301 2001::1
18             0         5.5.5.5          --               MPLS
0                            2001       e8c5:7a76:583b --
18             0         5.5.5.5          --               MPLS
PE2#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled
Code    FEC               FTN-ID    Nhlfe-ID  Tunnel-id   Pri   LSP-Type     Out-Label
Out-Intf    ELC   Nexthop
   O>   2.2.2.2/32          1         3         0          Yes   LSP_DEFAULT  16002
ce3         No    30.1.1.1
   O>   3.3.3.3/32          2         5         0          Yes   LSP_DEFAULT  16003
ce3         No    30.1.1.1
   O>   4.4.4.4/32          3         7         0          Yes   LSP_DEFAULT  3
ce3         No    30.1.1.1
PE2#
PE2#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
```

```
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

ILM-ECMP LDP: Disabled
Code    FEC/VRF/L2CKT    ILM-ID     In-Label    Out-Label   In-Intf    Out-Intf/VRF
Nexthop                  pri    LSP-Type
   B>   evpn:2001          5          18          Nolabel     N/A        N/A
127.0.0.1                Yes    LSP_DEFAULT
   B>   evpn:201           3          17          Nolabel     N/A        N/A
127.0.0.1                Yes    LSP_DEFAULT
   B>   IpVrf1            1          16          Nolabel     N/A        N/A             N/
A                       Yes    LSP_DEFAULT
   B>   evpn:201           2          26880       Nolabel     N/A        N/A
127.0.0.1                Yes    LSP_DEFAULT
   O>   3.3.3.3/32        7          16003       16003       N/A        ce3
30.1.1.1                 Yes    LSP_DEFAULT
   O>   2.2.2.2/32        6          16002       16002       N/A        ce3
30.1.1.1                 Yes    LSP_DEFAULT
   O>   4.4.4.4/32        8          16004       3           N/A        ce3
30.1.1.1                 Yes    LSP_DEFAULT
   B>   evpn:2001          4          26881       Nolabel     N/A        N/A
127.0.0.1                Yes    LSP_DEFAULT
   O>   30.1.1.1/32       9          27520       3           N/A        ce3
30.1.1.1                 Yes    LSP_DEFAULT
PE2#
PE2#
PE2#show mpls ftn-table
 Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: OSPF-SR, distance: 110, Action-type: Redirect to Tunnel, Exp-bits: 0x0,
Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, QoS Resource id: 0, Description:
N/A, , Color: 0
     Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 2
      Owner: OSPF-SR, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: OSPF-SR, Stale: NO, out intf: ce3, out label: 16002
    Nexthop addr: 30.1.1.1        cross connect ix: 3, op code: Push


 Primary FTN entry with FEC: 3.3.3.3/32, id: 2, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: OSPF-SR, distance: 110, Action-type: Redirect to Tunnel, Exp-bits: 0x0,
Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, QoS Resource id: 0, Description:
N/A, , Color: 0
     Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
      Owner: OSPF-SR, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 4, owner: OSPF-SR, Stale: NO, out intf: ce3, out label: 16003
    Nexthop addr: 30.1.1.1        cross connect ix: 4, op code: Push
```

```
 Primary FTN entry with FEC: 4.4.4.4/32, id: 3, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: OSPF-SR, distance: 110, Action-type: Redirect to Tunnel, Exp-bits: 0x0,
Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, QoS Resource id: 0, Description:
N/A, , Color: 0
     Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 6
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: ce3, out label: 3
    Nexthop addr: 30.1.1.1        cross connect ix: 5, op code: Push


CE2#show clns is-neighbors

Tag 101:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0005 ce1.101     Up     L2   64        0000.0000.0006.01
0000.0000.0005 ce1.1001    Up     L2   64        0000.0000.0006.02

CE2#show clns is-neighbors detail

Tag 101:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0005 ce1.101     Up     L2   64        0000.0000.0006.01
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 00:13:41
  Area Address(es): 49.0001
  IP Address(es): 201.201.201.1
  Level-2 Protocols Supported: IPv4, IPv6
  Adjacency advertisement: Advertise

0000.0000.0005 ce1.1001    Up     L2   64        0000.0000.0006.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 00:13:41
  Area Address(es): 49.0001
  IPv6 Address(es): fe80::e69d:73ff:feb1:c301
  Level-2 Protocols Supported: IPv4, IPv6
  Adjacency advertisement: Advertise

CE2#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"
i L2         1.1.1.1/32 [115/10] via 201.201.201.1, ce1.101, 00:10:22
C            6.6.6.6/32 is directly connected, lo, 00:13:54
i L2         99.99.99.0/24 [115/10] via 201.201.201.1, ce1.101, 00:10:22
i L2         101.101.101.0/24 [115/10] via 201.201.201.1, ce1.101, 00:10:22
C            127.0.0.0/8 is directly connected, lo, 00:18:45
C            201.201.201.0/24 is directly connected, ce1.101, 00:13:54

Gateway of last resort is not set
CE2#show ipv6 route
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, 00:18:50
i L2    1001::/64 [115/10] via fe80::e69d:73ff:feb1:c301, ce1.1001, 00:10:27
C       2001::/64 via ::, ce1.1001, 00:13:59
C       fe80::/64 via ::, ce1.1001, 00:13:59
```

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
MPLS	Multiprotocol Label Switching
MPLS	Multiprotocol Label Switching
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 EVPN-MPLS	A network architecture that combines EVPN and MPLS to provide efficient and scalable Layer 2 and Layer 3 services within a network, particularly in SP environments.
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 VxLAN-EVPN 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 VxLAN-EVPN 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 VxLAN-EVPN 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 8-1

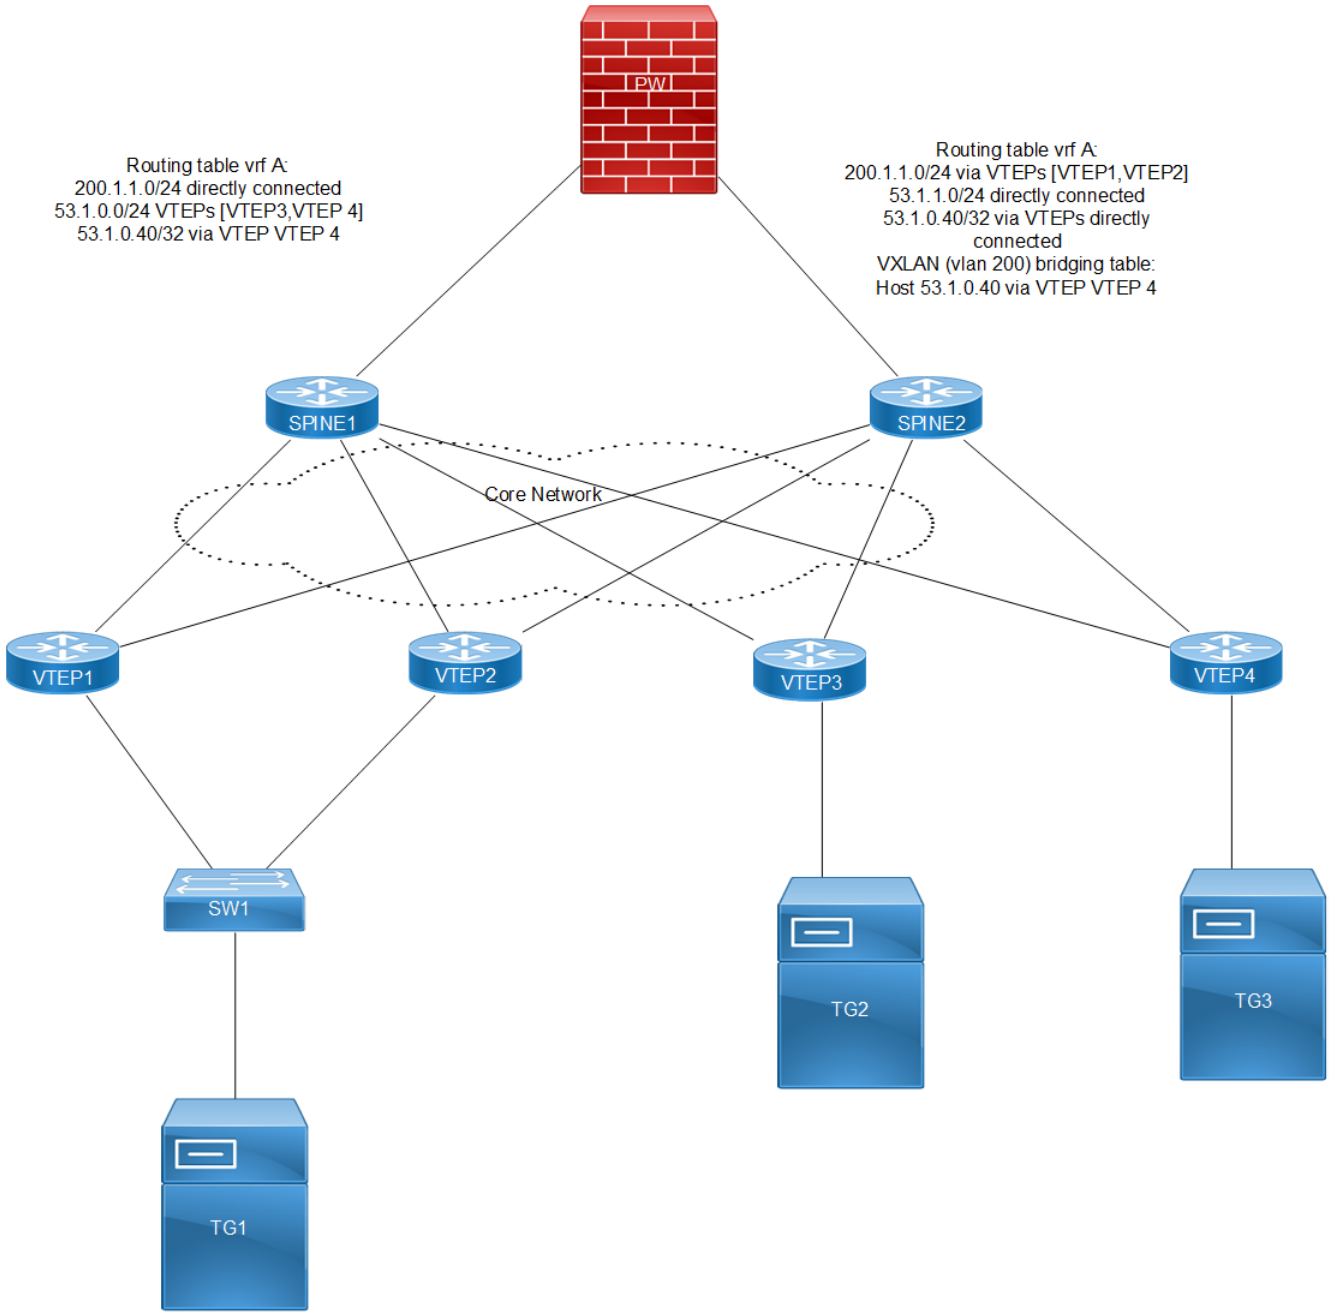

Figure 8-1: VxLAN EVPN IRB Connected host

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 VxLAN-EVPN 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 9 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 use QoS 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.

Two commands support L2 VxLAN QoS:

- `qos profile dscp-encap`: Use this command to create new profiles or to update "default" profiles for dscp-dscpEncap and Queue to DSCP value. This profile will be mapped to nvo vxlan tunnel at the egress direction of the VTEP. Default Qos profile dscp-encap would take preference than default qos profile queue-color-to dscp when no user-defined qos profile is configured on the nvo vxlan tunnel mode.
- `l2 queue dscp`: Use this command to configure or update user defined mapping for queue to dscp for egress L2 Traffic over VxLAN tunnel.

Topology

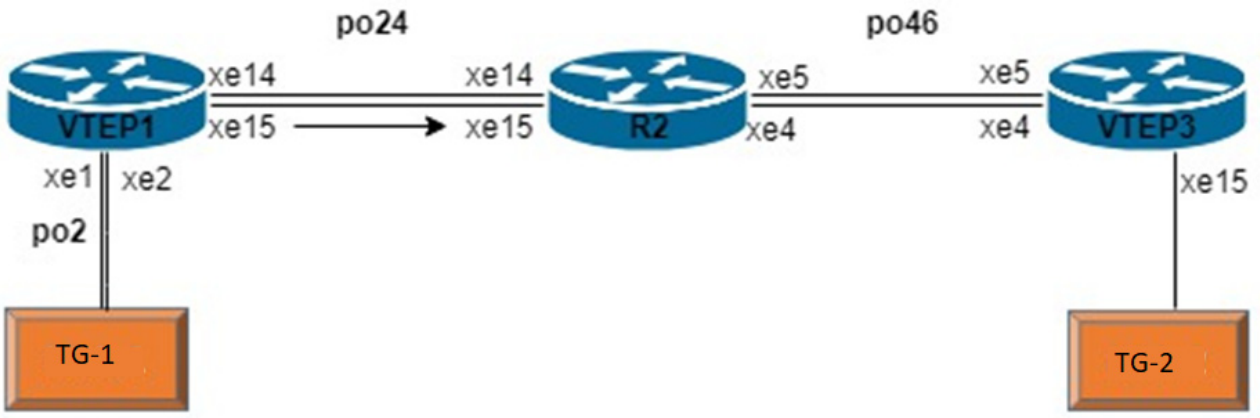

Figure 9-2: VxLAN QoS

COS-DSCP

RTR1/VTEP1

`VTEP1#configure terminal`	Enter Configure mode.
`VTEP1(config)#mac vrf vrf1`	Create mac routing/forwarding instance with vrf1 name and enter into vrf mode
`VTEP1(config-vrf)#rd 1.1.1.1:11`	Assign RD value
`VTEP1(config-vrf)#route-target both 10.10.10.10:100`	Assign route-target value for import/export
`VTEP1(config-vrf)#exit`	Exit form vrf mode
`VTEP1(config)#hardware-profile filter vxlan enable`	Enable hardware profile for vxlan
`VTEP1(config)#qos enable`	Enable qos
`VTEP1(config)#qos profile cos-to-queue COS-QUE`	Create qos profile for mapping traffic towards tunnel from access-if.
`VTEP1(config-ingress-cos-map)#cos 2 queue 3`	Configure particular COS value to the queue value for configured profile.
`VTEP1(config-ingress-cos-map)#exit`	Exit from qos profile config mode

`VTEP1(config)#qos profile queue-color-to-dscp QUE-DSCP`
`OR`
`VTEP1(config)#qos profile dscp-encap DSCP-ENCAP`
	Create qos profile for attaching in vxlan tunnel egress.
	
	Either one of the qos profile CLI can be configured. Functionality of queue-color-to-dscp and dscp-encap qos profiles remains the same.

VTEP1(config-egress-dscp-map)#queue 3 dscp 16 OR VTEP1(config-egress-dscp-encap-map)#l2 queue 3 dscp 16	Configure particular queue value to the dscp value for configured profile. Either one of the qos profile CLI can be configured. Functionality of queue-color-to-dscp and dscp-encap qos profiles remains the same.
VTEP1(config-egress-dscp-encap-map)#exit	Exit from qos profile config mode
VTEP1(config)#interface po2	Create a port channel po2
VTEP1(config-if)#switchport	Configure port as switchport
VTEP1(config-if)#load-interval 30	Set load-interval
VTEP1(config-if)#interface po24	Create a port channel po24
VTEP1(config-if)#load-interval 30	Configure port as switchport
VTEP1(config-if)#ip address 24.1.1.1/30	Set load-interval
VTEP1(config-if)#interface lo	Enter in to loopback interface
VTEP1(config-if)#ip address 1.1.1.1/32 secondary	Configure ip address
VTEP1(config-if)#interface xe1	Enter in to interface mode
VTEP1(config-if)#channel-group 2 mode active	Map to channel-group
VTEP1(config-if)#interface xe2	Enter in to interface mode
VTEP1(config-if)#channel-group 2 mode active	Map to channel-group
VTEP1(config-if)#interface xe14	Enter in to interface mode
VTEP1(config-if)#channel-group 24 mode active	Map to channel-group
VTEP1(config-if)#interface xe15	Enter in to interface mode
VTEP1(config-if)#channel-group 24 mode active	Map to channel-group
VTEP1(config-if)#router ospf 1	Create ospf instance
VTEP1(config-router)#ospf router-id 1.1.1.1	Configure ospf router-id
VTEP1(config-router)#network 1.1.1.1/32 area 0.0.0.0	Configure loopback network address in to ospf
VTEP1(config-router)#network 24.1.1.0/30 area 0.0.0.0	Configure network address in to ospf
VTEP1(config-router)#router bgp 100	Enter into Router BGP mode
VTEP1(config-router)#neighbor 6.6.6.6 remote-as 100	Specify a neighbor router with peer ip address and remote-as defined
VTEP1(config-router)#neighbor 6.6.6.6 update-source lo	Specify the neighbor to use loopback address as source
VTEP1(config-router)#address-family l2vpn evpn	Enter into l2vpn evpn address-family
VTEP1(config-router-af)#neighbor 6.6.6.6 activate	Activate the neighbor to address-family
VTEP1(config-router)#nvo vxlan vtep-ip-global 1.1.1.1	Configure Source vtep-ip-global configuration

VTEP1(config)#nvo vxlan tunnel qos-map-mode cos-dscp egress QUE-DSCP OR VTEP1(config)#nvo vxlan tunnel qos-map-mode cos-dscp egress DSCP-ENCAP	Configure the mapping qos profile in to vxlan tunnel egress Configure the mapping qos profile in to vxlan tunnel egress
VTEP1(config)#nvo vxlan id 1 in-gress-replication inner-vid-disabled	Create vnid 1 and disable inner-vid
VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1	Assign vrf for evpn-bgp to carry EVPN route
VTEP1(config-nvo)#nvo vxlan access-if port-vlan po2 1001	Create vxlan access-if with vlan 1001
VTEP1(config-nvo-acc-if)#no shutdown	No shut the vxlan access-if
VTEP1(config-nvo-acc-if)#map vnid 1	Map vnid to the vxlan access-if
VTEP1(config-nvo-acc-if)#map qos-profile cos-to-queue COS-QUE	Map qos profile for vxlan access-if ingress traffic from CE
VTEP1(config-nvo-acc-if)#exit	Exit from VxLAN access-interface mode and enter into configuration mode.
VTEP1(config)#commit	Commit the candidate configuration to the running configuration

RTR2

R2#configure terminal	Enter Configure mode.
R2(config)#interface po24	Create port channel
R2(config-if)#load-interval 30	Set load-interval
R2(config-if)#ip address 24.1.1.2/30	Assign ip address
R2(config-if)#interface po46	Create port channel
R2(config-if)#load-interval 30	Set load-interval
R2(config-if)#ip address 46.1.1.1/30	Assign ip address
R2(config-if)#interface lo	Enter in to loopback interface
R2(config-if)#ip address 4.4.4.4/32 sec-ondary	Assign secondary ip address
R2(config-if)#interface xe4	Enter into interface mode
R2(config-if)#channel-group 46 mode active	Map port channel to the interface
R2(config-if)#interface xe5	Enter into interface mode
R2(config-if)#channel-group 46 mode active	Map port channel to the interface
R2(config-if)#interface xe14	Enter into interface mode
R2(config-if)#channel-group 24 mode active	Map port channel to the interface
R2(config-if)#interface xe15	Enter into interface mode
R2(config-if)#channel-group 24 mode active	Map port channel to the interface
R2(config-if)#router ospf 1	Create ospf instance
R2(config-router)#ospf router-id 4.4.4.4	Configure ospf router-id
R2(config-router)#network 4.4.4.4/32 area 0.0.0.0	Configure ospf network address with respective area

R2(config-router)#network 24.1.1.0/30 area 0.0.0.0	Configure ospf network address with respective area
R2(config-router)#network 46.1.1.0/30 area 0.0.0.0	Configure ospf network address with respective area
R2(config-router)#exit	Exit from router mode.
R2(config)#commit	Commit the candidate configuration to the running configuration

RTR3/VTEP2

VTEP3#configure terminal	Enter Configuration mode
VTEP3(config)#mac vrf vrf1	Create mac routing/forwarding instance with vrf1 name and enter into vrf mode
VTEP3(config-vrf)#rd 6.6.6.6:11	Assign RD value
VTEP3(config-vrf)#route-target both 10.10.10.10:100	Assign route-target value for import/export
VTEP1(config-vrf)#exit	Exit from vrf mode
VTEP1(config)#hardware-profile filter vxlan enable	Enable hardware profile for vxlan
VTEP3(config)#qos enable	Enable QOS
VTEP3(config)#qos statistics	Enable QOS statistics
VTEP3(config)#qos profile queue-color-to-cos QUE-COS	Create qos profile for mapping incoming traffic from tunnel to access-if.
VTEP3(config-egress-cos-map)#queue 4 cos 5	Configure particular queue value to the cos value for configured profile.
VTEP3(config-egress-cos-map)#qos profile dscp-to-queue DSCP-QUE	Create qos profile for attaching in vxlan tunnel ingress.
VTEP3(config-ingress-dscp-map)#dscp 16 queue 4	Configure particular dscp value to the queue value for configured profile.
VTEP3(config-egress-dscp-map)#interface po46	Create port channel
VTEP3(config-if)#load-interval 30	Set load interval
VTEP3(config-if)#ip address 46.1.1.2/30	Assign ip address
VTEP3(config-if)#interface lo	Enter into loopback interface
VTEP3(config-if)#ip address 6.6.6.6/32 secondary	Assign secondary ip address
VTEP3(config-if)#interface xe4	Enter into interface mode
VTEP3(config-if)#channel-group 46 mode active	Map channel group into the interface
VTEP3(config-if)#interface xe5	Enter into interface mode
VTEP3(config-if)#channel-group 46 mode active	Map channel group into the interface
VTEP3(config-if)#interface xe15	Enter into interface mode
VTEP3(config-if)#switchport	Make interface as L2 port
VTEP3(config-if)#load-interval 30	Set load interval
VTEP3(config-if)#router ospf 1	Create ospf instance

`VTEP3(config-router)#ospf router-id 6.6.6.6`	Configure ospf router-id
`VTEP3(config-router)#network 6.6.6.6/32 area 0.0.0.0`	Configure ospf network address with respective area
`VTEP3(config-router)#network 46.1.1.0/30 area 0.0.0.0`	Configure ospf network address with respective area
`VTEP3(config-router)#router bgp 100`	Enter into Router BGP mode
`VTEP3(config-router)#neighbor 1.1.1.1 remote-as 100`	Specify a neighbor router with peer ip address and remote-as defined
`VTEP3(config-router)#neighbor 1.1.1.1 update-source lo`	Specify the neighbor to use loopback address as source
`VTEP3(config-router)#address-family l2vpn evpn`	Enter into l2vpn evpn address-family
`VTEP3(config-router-af)#neighbor 1.1.1.1 activate`	Activate the neighbor to address-family
`VTEP3(config)#nvo vxlan vtep-ip-global 6.6.6.6`	Configure Source vtep-ip-global configuration
`VTEP3(config)#nvo vxlan tunnel qos-map-mode cos-dscp ingress DSCP-QUE`	Configure the mapping qos profile in to vxlan tunnel ingress
`VTEP3(config)#nvo vxlan id 1 in-gress-replication inner-vid-disabled`	Create vnid 1 and disable inner-vid
`VTEP3(config-nvo)#vxlan host-reachability-protocol evpn-bgp vrf1`	Assign vrf for evpn-bgp to carry EVPN route
`VTEP3(config-nvo)#nvo vxlan access-if port-vlan xe15 1000`	Create vxlan access-if with vlan 1000
`VTEP3(config-nvo-acc-if)#no shutdown`	No shut the vxlan access-if
`VTEP3(config-nvo-acc-if)#map vnid 1`	Map vnid to the vxlan access-if
`VTEP3(config-nvo-acc-if)#map qos-profile queue-color-to-cos QUE-COS`	Map qos profile for vxlan access-if egress traffic to CE
`VTEP3(config-nvo-acc-if)#exit`	Exit from VxLAN access-interface mode and enter into configuration mode
`VTEP3(config)#commit`	Commit the candidate configuration to the running configuration

Validation

As per the QoS configuration, when L2 traffic with cos value 2 is sent to VTEP1 access-if, the packets are forwarded to queue 3 and packets in queue 3 are mapped with dscp value 16 while egress out of tunnel. At VTEP2, when packets with dscp value 16 ingress 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 (OR DSCP-ENCAP)
```

```
!
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 queue 3
!
qos profile queue-color-to-dscp QUE-DSCP
 queue 3 color all dscp 16
!
               OR
VTEP1#show run qos
qos enable
!
qos profile cos-to-queue COS-QUE
 cos 2 queue 3
!
qos profile dscp-encap DSCP-ENCAP
 l2 queue 3 dscp 16
!

VTEP1#sh int 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#sh int 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#
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#
```

RTR3/VTEP3

```
VTEP3#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
!
!
VTEP3#sh 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
!
VTEP3#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

VTEP3#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
```

```
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      ----        AC   xe15     --- Single Homed port ---    1000 ----     ----
----

Total number of entries are 3
VTEP3#

VTEP3#show int 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

VTEP3#
```

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 when 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 1-3: VxLAN EVPN EVC

Configuration

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)# hardware-profile filter vxlan enable	Enable hardware profile for vxlan
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for VxLAN access-if port counters
(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)#interface xe15.1000 switchport	Create L2 subinterface xe15.1000
(config-if)#encapsulation dot1q 1000	Configure encapsulation dot1q with vlan id 1
(config-if)#access-if-evpn	Enter access-if-evpn mode

(config-acc-if-evpn)#map vpn-id 1	Map vnid 1 to access interface xe15.1000
(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)# hardware-profile filter vxlan enable	Enable hardware profile for vxlan
(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for VxLAN access-if port counters
(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)#interface xe6.1000 switchport	Create L2 subinterface xe6.1000
(config-if)#encapsulation dot1q 1000	Configure encapsulation dot1q with vlan id 1
(config-if)#access-if-evpn	Enter access-if-evpn mode
(config-acc-if-evpn)#map vpn-id 1	Map vnid 1 to access interface xe6.1000
(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]:[32,33.33.33.0]
                     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 configure on VTEP2 as follows:

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 configure on VTEP2 as follows:

`(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 and MPLS 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).

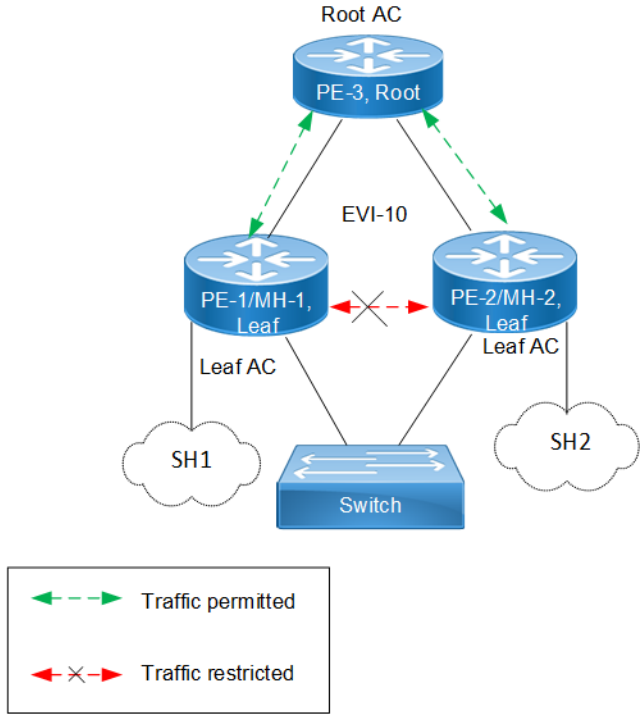

EVPN E-Tree

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.

   ```
   !
   hardware-profile filter vxlan enable
   hardware-profile filter vxlan-mh enable
   !
   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.

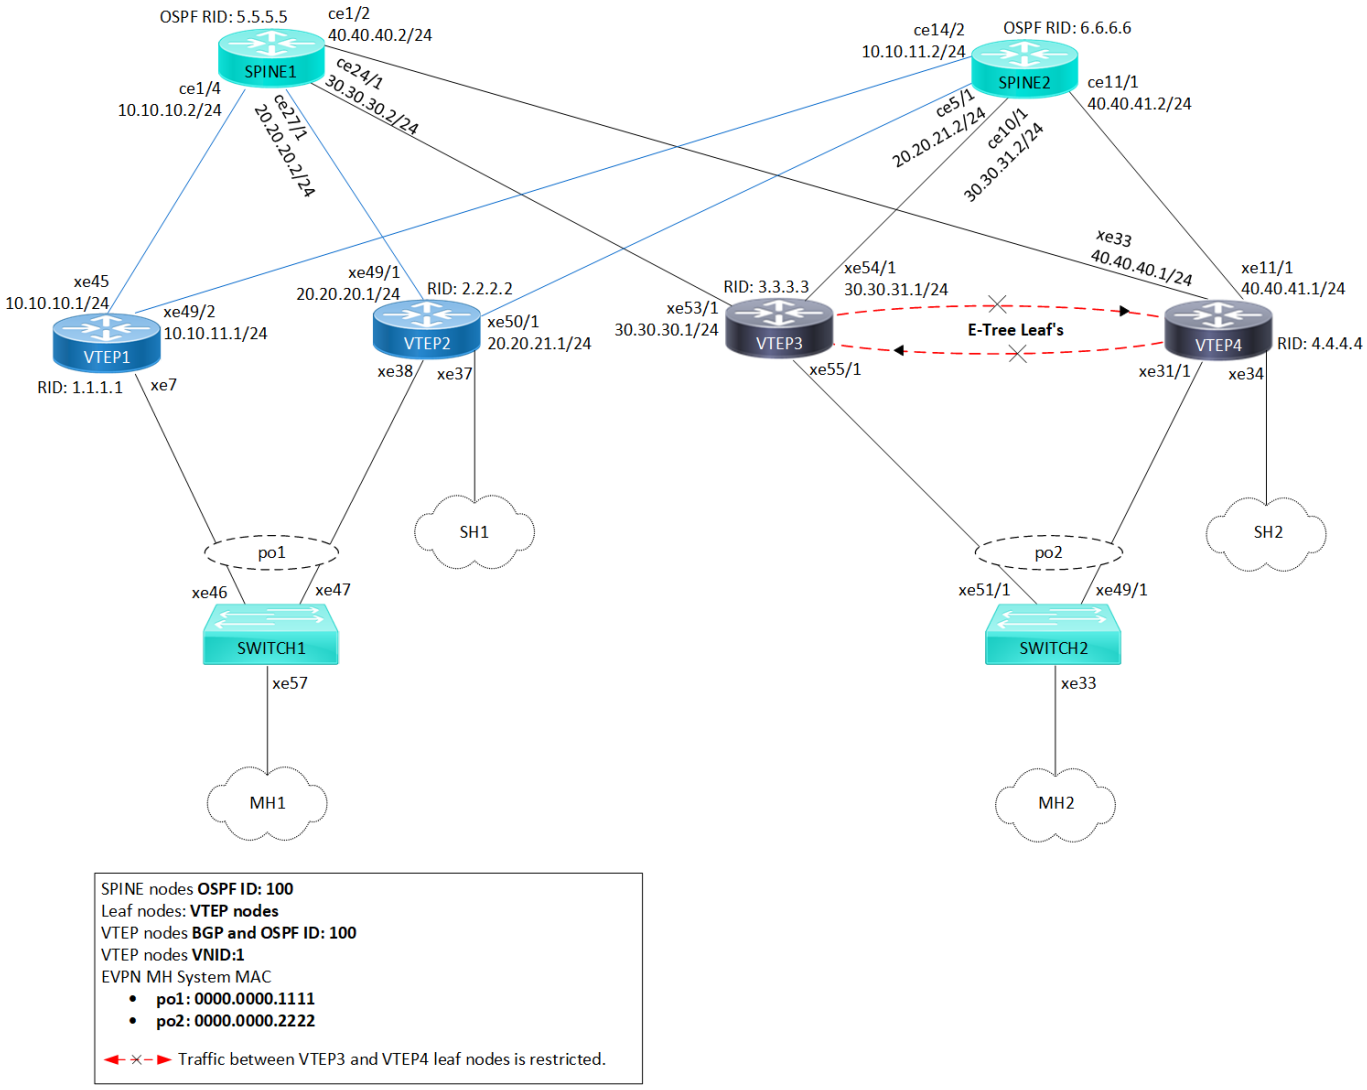

VXLAN EVPN E-Tree Topology

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

```
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
```

```
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

```
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
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

```
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
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

```
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
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. Verify if BUM traffic is transmitted between Leaf sites inter-PE:
   - Ensure that a BUM tunnels are not established between inter-PE devices.
   - Validate this by examining the Multicast ingress group, using the `show evpn mpls tunnel` command. For EVPN MPLS, confirm that BUM tunnels are not created.
4. 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.
5. 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.
6. 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.
7. 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 is 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.

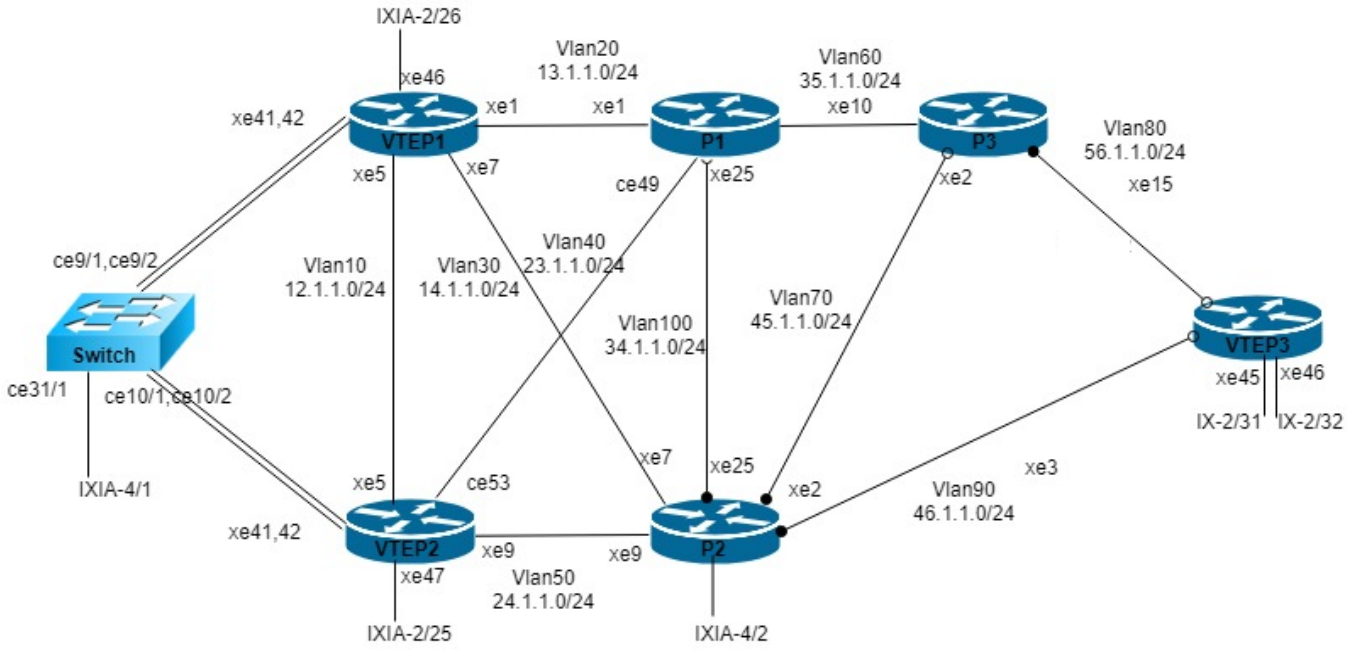

Figure 3-4: VxLAN over SVI

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)#hardware-profile filter egress-ipv4 enable`	Enable hardware filter for egress ipv4
`VTEP1(config)#hardware-profile filter vxlan enable`	Enable vxlan in hardware
`VTEP1(config)# hardware-profile filter vxlan-mh enable`	Enable vxlan-mh in hardware
`VTEP1(config)#hardware-profile statistics ac-lif enable`	Enable ac-lif for enabling vxlan counters
`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
`(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
`(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)#interface po1.2000 switchport	Create L2 subinterface po1.2000
VTEP1(config-if)#encapsulation dot1q 2000	Configure encapsulation dot1q with vlan id 1000
VTEP1(config-if)#access-if-evpn	Enter access-if-evpn mode
VTEP1(config-acc-if-evpn)#map vpn-id 1000	Map vnid 1000 to access interface po1.2000
VTEP1(config)#interface xe46 switchport	Create L2 subinterface xe46
VTEP1(config-if)#encapsulation dot1q	Configure encapsulation dot1q with vlan id 1
VTEP1(config-if)#access-if-evpn	Enter access-if-evpn mode
VTEP1(config-acc-if-evpn)#map vpn-id 1	Map vnid 1 to access interface xe46
VTEP1(config)#interface po1.1000 switchport	Create L2 subinterface po1.1000
VTEP1(config-if)#encapsulation dot1q 1000	Configure encapsulation dot1q with vlan id 1
VTEP1(config-if)#access-if-evpn	Enter access-if-evpn mode
VTEP1(config-acc-if-evpn)#map vpn-id 1	Map vnid 1 to access interface po1.1000
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)#hardware-profile filter egress-ipv4 enable	Enable hardware filter for egress ipv4
VTEP2(config)#hardware-profile filter vxlan enable	Enable vxlan in hardware
VTEP2(config)#nvo vxlan enable	Enable vxlan
VTEP2(config)# hardware-profile filter vxlan-mh enable	Enable vxlan-mh in hardware
VTEP2(config)#evpn vxlan multihoming enable	Enable evpn vxlan multihoming

VTEP2(config)#hardware-profile statistics ac-lif enable	Enable ac-lif for enabling vxlan counters
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 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 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
`VTEP2(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
`VTEP2(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)#interface xe47 switchport`	Create L2 subinterface xe47
`VTEP2(config-if)#encapsulation dot1q`	Configure encapsulation dot1q with vlan id 1
`VTEP2(config-if)#access-if-evpn`	Enter access-if-evpn mode
`VTEP2(config-acc-if-evpn)#map vpn-id 1`	Map vnid 1 to access interface xe47
`VTEP2(config)#interface po1.2001 switchport`	Create L2 subinterface po1.2001
`VTEP2(config-if)#encapsulation dot1q 2001`	Configure encapsulation dot1q with vlan id 1000
`VTEP2(config-if)#access-if-evpn`	Enter access-if-evpn mode
`VTEP2(config-acc-if-evpn)#map vpn-id 1`	Map vnid 1000 to access interface po1.2001
`VTEP2(config)#interface po1.2000 switchport`	Create L2 subinterface po1.2000
`VTEP2(config-if)#encapsulation dot1q 2001`	Configure encapsulation dot1q with vlan id 1000
`VTEP2(config-if)#access-if-evpn`	Enter access-if-evpn mode
`VTEP2(config-acc-if-evpn)#map vpn-id 1000`	Map vnid 1000 to access interface po1.2000
`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	Configure bfd interval globally
VTEP3(config)#load-balance enable	Enable load balancing
VTEP3(config)#load-balance ipv4 dest-ipv4 src-ipv4	Enable load balnce based on souce and destination ip address
VTEP3(config)#hardware-profile filter egress-ipv4 enable	Enable hardware filter for egress ipv4
VTEP3(config)#hardware-profile filter vxlan enable	Enable vxlan-mh in hardware
VTEP3(config)# hardware-profile filter vxlan-mh enable	Enable evpn vxlan multihoming
VTEP3(config)#evpn vxlan multihoming enable	Enable vxlan in hardware

VTEP3(config)#hardware-profile statistics ac-lif enable	Enable statistics on vxlan access 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 neighbors
`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)#interface xe45.3001 switchport	Create L2 subinterface xe45.3001
VTEP3(config-if)#encapsulation dot1q 3001	Configure encapsulation dot1q with vlan id 1
VTEP3(config-if)#access-if-evpn	Enter access-if-evpn mode
VTEP3(config-acc-if-evpn)#map vpn-id 1	Map vnid 1 to access interface xe45.3001
VTEP3(config)#interface xe46 switchport	Create L2 subinterface xe46
VTEP3(config-if)#encapsulation dot1q	Configure encapsulation dot1q with vlan id 1000
VTEP3(config-if)#access-if-evpn	Enter access-if-evpn mode
VTEP3(config-acc-if-evpn)#map vpn-id 1000	Map vnid 1000 to access interface xe46
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)#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 - 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 1-5: VxLAN OAM Packet Flow

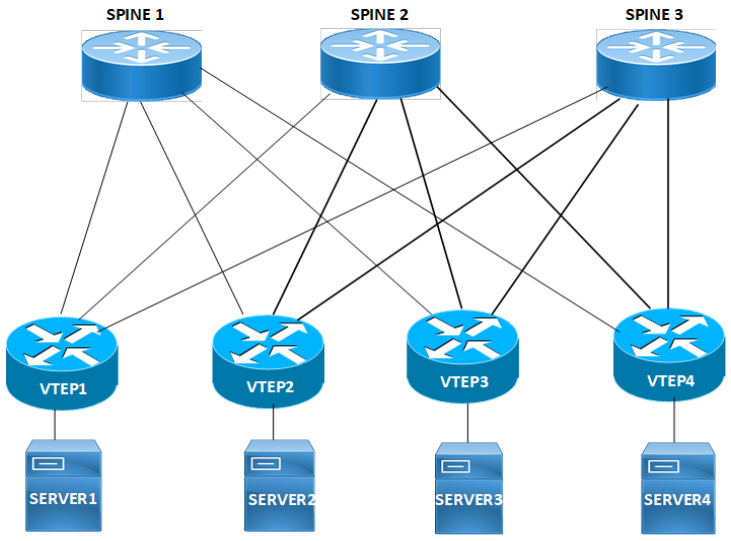

Following is the flow of VxLAN OAM Packet communication:

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 1-6 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

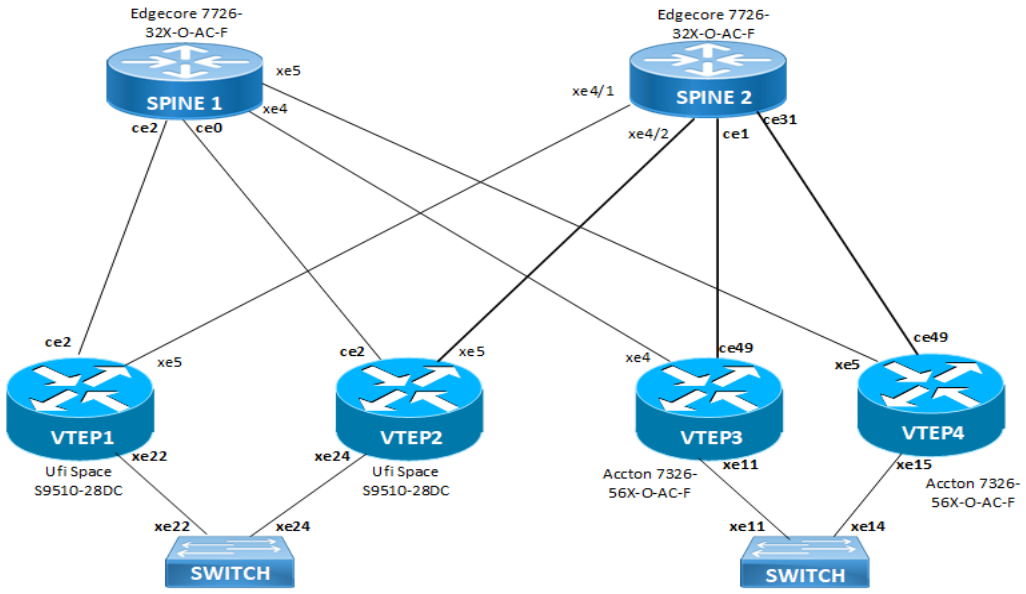

Figure 1-6: VxLAN OAM Topology

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:

access-if-evpn

Use this command to create the evpn access-port.

Use the no form of this command to delete the evpn access-port.

Command Syntax

```
access-if-evpn
no access-if-evpn
```

Parameters

None

Command Mode

L2 Sub-interface mode

Applicability

This command was introduced in OcNOS version 6.0.0.

Examples

```
#configure terminal
(config)#interface xe1.1 switchport
(config-if)#access-if-evpn
(config-access-if)#end
```

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

Access interface EVPN mode

Applicability

This command was introduced before OcNOS version 1.3.

The Access interface EVPN mode is supported only from OcNOS version 6.0.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#arp-cache disable
(config-nvo-acc-if)#exit

(config)#interface xe7.100 switchport
(config-if)#access-if-evpn
(config-acc-if-evpn)#arp-cache disable
```

arp-nd cos

Use this command to explicitly set or override the COS value `(0-7)` in the ARP/ND reply sent (also used in ARP/ND proxy sub-feature). By default ARP/ND reply will use CoS value of received ARP/ND request.

Use the `no` form of this command to use the default COS value.

Command Syntax

```
arp-nd cos <0-7>
```

```
no arp-nd cos
```

Parameters

`<0-7>`

Default

`CoS 0` or previous ARP/ND reply CoS being used.

Command Mode

Access if evpn mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Examples

```
#configure terminal
(config)#interface xe1.1 switchport
(config-if)#access-if-evpn
(config-access-if)#arp-nd cos 5
(config-access-if)#end
```

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

Access interface EVPN mode

Applicability

This command was introduced before OcNOS version 1.3.

The Access interface EVPN mode is supported only from OcNOS version 6.0.

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

(config)#interface xe7.100 switchport
(config-if)#access-if-evpn
(config-acc-if-evpn)#arp-nd flood-supress disable
```

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 seconds.

Use the `no` form of this command to remove the configuration.

Note: After this timer interval, it sends out ARP to revalidate and 3 times of this would lead to removal of the dynamic entry.

Command Syntax

```
nvo vxlan arp-nd refresh-timer <3-190>
no nvo vxlan arp-nd refresh-timer
```

Parameters

`<3-190>`	Refresh timer value in seconds (age-out is refresh time * 3)

Command Mode

Configuration mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#config mode
(config)#nvo vxlan arp-nd refresh-timer 100
(config)#no nvo vxlan arp-nd refresh-timer
```

clear mac address-table dynamic vxlan

Use this command to clear dynamically learned MACs.

Command Syntax

```
clear mac address-table dynamic vxlan
clear mac address-table dynamic vxlan vnid <1-16777215>
clear mac address-table dynamic vxlan vnid <1-16777215> (address MACADDR|)
```

Parameters

`address`	Clear the specified MAC Address
`VNID`	VxLAN network identifier
`<1-16777215>`	Range supported for VNID

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear mac address table dynamic vxlan
#clear mac address-table dynamic vxlan vnid 100
#clear mac address-table dynamic vxlan vnid 100 address 0000.0005.0505
```

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 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

#show running-config  nvo vxlan
!
nvo vxlan enable
!
nvo vxlan access-if port-vlan xe1 2
 description member-port xe1 with vlan 2
 no shutdown
```

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

Access interface EVPN mode

Applicability

This command was introduced before OcNOS version 1.3.

The Access interface EVPN mode is supported only from OcNOS version 6.0.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#dynamic-learning disable
(config-nvo-acc-if)#exit

(config)#interface xe7.100 switchport
(config-if)#access-if-evpn
(config-acc-if-evpn)#dynamic-learning disable
```

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.

Note: Before configuring the TPID in the NVO access interface mode, first configure it at port level.

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)#no encapsulation
(config-nvo-acc-if)#exit
```

evpn esi hold-time

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 multihoming.

Use the `no` form of this command to make the ESI up immediately when configuring the access-if cli.

Command Syntax

```
evpn esi hold-time <10-300>
no evpn esi hold-time <10-300>
```

Parameters

`<10-300>`	Hold time in seconds

Default

The default value is 0.

Command Mode

Configuration Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)# evpn esi hold-time 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

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 before OcNOS version 4.1.

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 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 before OcNOS version 4.1.

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 before OcNOS version 4.1.

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
```

evpn multi-homed

Use this command to configure single-active or port-active load-balancing Ethernet Segment Identifier (ESI) configuration on a link with multihomed Customer Edge (CE).

Use the no parameter of this command to unconfigure Ethernet Segment Identifier (ESI) configuration on a link with multihomed Customer Edge (CE).

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)) load-balancing single-active
evpn multi-homed system-mac (XX-XX-XX-XX-XX-XX|XX:XX:XX:XX:XX:XX|XXXX.XXXX.XXXX)
  load-balancing port-active
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

The default value is 0.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3 and underwent modifications in the OcNOS version 6.4.1.

Example

```
#configure terminal
(config)#interface xe1
(config-if)#evpn multi-homed esi 00:11:22:33:44:55:66:77:88 load-balancing
single-active
(config)#exit

#configure terminal
(config)#interface po1
(config-if)#evpn multi-homed system-mac 0000.0000.1111 load-balancing port-
active
(config)#exit
```

evpn vxlan multihoming enable

Use this command to enable evpn vxlan multi-homing

Use the no form of this command to disable evpn vxlan multi-homing.

Note: You must restart the device after giving this command. If there are devices in the topology which have multi-homed CEs, then devices which do not have multi-homed CEs should also enable multihoming so that they can load share traffic to the multi-homed CEs.

Note: Before enabling multi-homing, configure the hardware-profiles:

- hardware-profile filter vxlan-mh
- hardware-profile filter (Qumran1) with the `egress-ipv4` parameter

Command Syntax

```
evpn vxlan multihoming enable
no evpn vxlan multihoming enable
```

Parameters

None

Default

By default, multi-homing is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and changed in OcNOS version 4.0.

Example

```
#configure terminal
#(config)#evpn vxlan multihoming enable
#(config)#exit
```

evpn-vlan-service

Use this command to configure VLAN-based 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-based
no evpn-vlan-service
```

Parameters

`vlan-based`	VLAN-based EVPN-Service type

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
```

hardware-profile filter vxlan

Use this command to configure hardware profile for nvo vxlan. This profile should be enabled before enabling VxLAN.

Note: You need to save the configuration and do a reboot after giving this command.

Command Syntax

```
hardware-profile filter vxlan enable
hardware-profile filter vxlan disable
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
#config mode
(config)# hardware-profile filter vxlan enable
(config)# hardware-profile filter vxlan disable
```

hardware-profile filter vxlan-mh

Use this command to enable the hardware-profile for VxLAN multi-homing to successfully activate multi-homing in the hardware.

Before enabling EVPN multi-homing (evpn multi-homed command), give this command.

Before disabling the hardware-profile, disable EVPN multi-homing.

Note: You need to save the configuration and do a reboot after giving this command.

Use the `disable` form of this command to disable the configured hardware-profile.

Command Syntax

```
hardware-profile filter vxlan-mh enable
hardware-profile filter vxlan-mh disable
```

Parameters

None

Default

By default, the VxLAN multi-homing hardware-profile is disabled.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
#configure terminal
#(config)#hardware-profile filter vxlan-mh enable
#(config)#hardware-profile filter vxlan-mh disable
#(config)#exit
```

hardware-profile filter vxlan-short-pkt

Use this command to enable/disable the VxLAN short packet forwarding filter group.

When this filter group is enabled, any short packet received less than 36B is sent to the CPU to add extra bytes to make the packet size up to 64B and sent back to access-port.

Use the `disable` form of this command to disable the configured hardware-profile.

Command Syntax

```
hardware-profile filter vxlan-short-pkt enable
hardware-profile filter vxlan-short-pkt disable
```

Parameters

None

Default

By default, the VxLAN short packet hardware-profile is disabled.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.3.1.

Example

```
#configure terminal
#(config)#hardware-profile filter vxlan-short-pkt enable
#(config)#exit
```

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

`<1-4094>`	IRB interface number

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.1.

Example

```
#configure teminal
(config)#interface irb1
(config)#no interface irb1
```

ip address

Use this command to set anycast flag for primary and secondary subnets under IRB interface.

With this anycast gateway can be supported for multiple subnets.

Command Syntax

```
ip address [ <A.B.C.D/M>  | anycast]
ip address [ <A.B.C.D>  | <A.B.C.D> | anycast]
ip address [ <A.B.C.D/M>  | secondary | anycast]
ip address [ <A.B.C.D>  | <A.B.C.D> | secondary | anycast]
```

Default

The default value is router mac

Parameters

`anycast`	Anycast flag
`secondary`	Used for secondary address option

Command Mode

IRB_IF Mode

Applicability

The anycast flag was introduced in OcNOS version 6.3.0.

Example

```
(config)#interface irb2
(config-irb-if)#ip address 40.1.1.1/24 anycast
(config-irb-if)#ip address 41.1.1.1/24 secondary anycast
(config-irb-if)#
(config)#interface irb1
(config-irb-if)#ip address 42.1.1.1 255.255.255.0 anycast
(config-irb-if)#ip address 43.1.1.1 255.255.255.0 secondary anycast
(config-irb-if)#
```

ipv6 address

Use this command to set anycast flag for any configured subnets under IRB interface.

With this anycast gateway can be supported for multiple subnets.

Command Syntax

```
ipv6 address [ < X:X::X:X/M > | anycast]
```

Default

The default value is router mac

Parameters

`anycast`	Anycast flag

Command Mode

IRB_IF Mode

Applicability

The anycast flag was introduced in OcNOS version 6.3.0.

Example

```
(config)#interface irb1
(config-irb-if)# ipv6 address 1100::1/64 anycast
```

ip dhcp relay uplink

Use this command to configure uplink interface towards server per vrf.

Use the `no` form of this command to un-configure uplink interface.

Command Syntax

```
ip dhcp relay uplink evpn
no ip dhcp relay uplink evpn
```

Parameters

`evpn`	IP VRF IRB interface

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 5.0.

Example

```
#configure teminal
(config)#ip vrf vrf1
(config-vrf)#ip dhcp relay uplink evpn
(config-vrf)#no ip dhcp relay uplink evpn
```

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 4.1.

Example

```
#configure teminal
(config)#ip vrf vrfip
(config-vrf)#l3vni 10002
(config-vrf)#no l3vni 10002
```

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

Access interface EVPN mode

Applicability

This command was introduced before OcNOS version 1.3.

The Access interface EVPN mode is supported only from OcNOS version 6.0.

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
(config)#interface xe7.100 switchport
(config-if)#access-if-evpn
(config-acc-if-evpn)#arp-nd flood-supress disable
```

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` form 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.

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 vpn-id 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 access interface mode

Access interface EVPN mode

Note: When configured in both modes, then the NVO_ACC_IF_MODE or ACC-IF-EVPN 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.

The Access interface EVPN mode is supported only from OcNOS version 6.0.

Example

```
#configure terminal
(config)#nvo vxlan id 3 ingress-replication inner-vid-disabled
(config-nvo)#mac-holdtime -1
(config-nvo)#exit

(config)#interface xe7.100 switchport
(config-if)#access-if-evpn
(config-acc-if-evpn)#mac-holdtime -1
```

map vpn-id

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 vpn-id <1-16777215>
no map vpn-id <1-16777215>
```

Parameters

`<1-16777215>`	VxLAN network identifier.

Default

No default value is specified for map vnid command.

Command Mode

NVO access interface mode

Access interface EVPN mode

Applicability

This command was introduced before OcNOS version 1.3.

The Access interface EVPN mode is supported only from OcNOS version 6.0.

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#map vpn-id 100
(config-nvo-acc-if)#exit

(config)#interface xe7.100 switchport
(config-if)#access-if-evpn
(config-acc-if-evpn)#map vpn-id 100
```

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

Access interface EVPN mode

Applicability

This command was introduced before OcNOS version 1.3.

The Access interface EVPN mode is supported only from OcNOS version 6.0

Example

```
#configure terminal
(config)#nvo vxlan access-if port-vlan xe1 2
(config-nvo-acc-if)#nd-cache disable
(config-nvo-acc-if)#exit

(config)#interface xe7.100 switchport
(config-if)#access-if-evpn
(config-acc-if-evpn)#nd-cache disable
```

nvo vxlan

Use this command to enable or disable VxLAN.

You must enable the VxLAN hardware profile with the hardware-profile filter vxlan command before enabling VxLAN.

Note: To make nvo vxlan disable and enable effective, system reboot is required.

Command Syntax

```
nvo vxlan (enable | disable)
```

Parameters

None

Default

By default, the VxLAN 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 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.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and VLAN_RANGE option is introduced in OcNOS version 5.0. This command is not available on Qumran2 devices.

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 id

Use this command to add a tenant and the type of VPN. This command changes the mode to NVO mode.

Use `nvo vxlan id` for creating ELAN and use `nvo vxlan id xconnect` with source and target identifier for ELINE/XConnect.

Use `no` form of this command to unconfigure the VxLAN ID.

Note:

- To use this command, user must first enable VXLAN with the nvo vxlan command.
- IRB: VXLAN IRB is supported with an `inner-vid-disabled` configuration. For sub-interfaces configured as access interfaces, use the `rewrite POP` command to remove VLAN tags with the POP on the sub-interface.

Command Syntax

```
nvo vxlan id <1-16777215> (multicast |default) | ((ingress-replication | xconnect
  target-vxlan-id <1-16777215>) (| inner-vid-disabled)) (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 group (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.
`xconnect target-vxlan-id <1-16777215>`	Specifies the cross-connect Identifier.
`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.

Command Mode

Configure mode

Applicability

Introduced before OcNOS version 1.3. Introduced the `etree-leaf and xconnect target-vxlan-id <1-16777215>` parameters 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

(config)#nvo vxlan id 30 xconnect target-vxlan-id 20
(config-nvo)#exit

(config)#nvo vxlan id 40 xconnect target-vxlan-id 50 ingress-replication
(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
```

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 before OcNOS version 4.1.

Example

```
#configure terminal
(config)#nvo vxlan irb
(config)#no nvo vxlan irb
```

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>`	Aging 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
(config)#no nvo vxlan mac-ageing-time 10
```

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.

You must enable VxLAN with the nvo vxlan command before you give this command.

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 bgp l2vpn evpn

Use this command to display details about Layer 2 Virtual Private Network (L2VPN) Ethernet Virtual Private Network (EVPN) routes.

Note: A BGP EVPN route update received for an unreachable IP address is also listed by this command and as a best route. This is because the next hop tracking feature is not supported for the EVPN address family. However, the tunnel to this IP address is shown in unresolved state by the show nvo vxlan tunnel output.

Command Syntax

```
show bgp l2vpn evpn (((vrf WORD)|(rd WORD)|)((prefix-route)|)(detail| time|))
show bgp l2vpn evpn mac-ip ((vrf WORD)|(rd WORD)|)
show bgp l2vpn evpn mcast
show bgp l2vpn evpn multihoming es-route <(rd WORD)|(vrf WORD)>
show bgp l2vpn evpn multihoming ethernet-ad-per-evi <(rd WORD)|(vrf WORD)>
show bgp l2vpn evpn multihoming ethernet-ad-per-es <(rd WORD)|(vrf WORD)>
```

Parameters

`vrf`	Virtual Routing and Forwarding instance
`WORD`	VRF name
`rd`	Route distinguisher
`WORD`	Route distinguisher: ASN:nn or IP:nn
`prefix-route`	Shows detail of the Prefix-Route (Type:5))
`detail`	Detailed output of the route-path
`time`	Display learnt time for details for evpn routes.
`mac-ip`	Show detail of the MAC-IP route (Type:2)
`mcast`	Show detail of the Inclusive MULTICAST route (Type:3)
`multihoming`	Show multihoming information
`peer-group`	Dynamic peer-group

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
show bgp l2vpn evpn detail
BGP route entry for prefix : [1]:[00:00:00:00:11:12:12:22:11:11]:[100]:[100]
 Route-Distinguisher: 1.1.1.1:1
 Flags : Valid, Selected, IBGP, Labelled
 Nexthop : 1.1.1.1  MED value : 0
 Community:
 Extended Community: RT:100:1  Encapsulation:VxLAN  ESI-Label:0
 Weight :0,  Local Preference :100
 AS Path : Local
```

```
 Origin : IGP
 Last Update : Thu Apr 13 12:05:23 2023
 Peer : 1.1.1.1

BGP route entry for prefix : [2]:[0]:[100]:[48,0000:1111:0000]:[32,10.12.11.12]:[100]
 Route-Distinguisher: 1.1.1.1:1
 Flags : Valid, Selected, IBGP, Labelled
 Nexthop : 1.1.1.1  MED value : 0
 Community:
 Extended Community: RT:100:1  Encapsulation:VxLAN  MAC_mob_seq:Static
 Weight :0,  Local Preference :100
 AS Path : Local
 Origin : IGP
 Last Update : Thu Apr 13 12:05:23 2023
 Peer : 1.1.1.1
Total number of prefixes 2
```

Table 1-3 shows the status codes displayed at the start of a route entry.

Table 1-3: status codes

Status code	Description	Comments
s	suppressed	Whether the route is suppressed and is not advertised to neighbors.
d	damped	When the penalty of a flapping route exceeds the suppress limit, the route is damped and remains in a withdrawn state until its penalty decreases below the reuse limit.
h	history	When the penalty of a flapping route does not exceed the suppress limit, the route is not damped and BGP maintains a history of the flapping route.
*	valid	Whether the route is valid. When a route is not suppressed, damped, or present in the history, it is valid.
>	best	The selected route to be installed in the kernel routing table.
i	internal	The prefix was learned from an iBGP peer.
l	labeled	BGP Labeled Unicast advertises route information between inter region routers.

Table 1-4 shows the codes at the end of each route entry that indicate where the route originated.

Table 1-4: origin codes

Origin Code	Description	Comments
i	IGP	The route is from an Interior Gateway Protocol.
e	EGP	The route is from an Exterior Gateway Protocol.
?	incomplete	Origin not known. Typically, these are routes redistributed from an IGP.

Table 1-5 explains the fields for each route.

Table 1-5: route entry fields

Field	Description
RD	Route distinguisher: AS number or IP address.
VRF	Name of the VRF.
Network	EVPN route information. The route type indicates the type of routing information advertised by the EVPN control plane: 2 MAC/IP Route: Endpoint reachability information, including MAC and IP addresses of the endpoints. 3 Inclusive Multicast Route: Information about how to forward Broadcast, Unknown Unicast and Multicast (BUM) traffic. The other fields included depend on the route type: Type 2: [ESI]:[E-Tag]:[Length, Host MAC address]:[Length, Host IP address]:[Label/VNID] Type 3: [E-Tag]:[Length, PE IP address] ESI (Ethernet Segment Identifier): a unique non-zero identifier that identifies an Ethernet segment, which is a set of links that connects a network or device to one or more PEs. ESI 0 denotes a single-homed site. E-Tag (Ethernet tag): identifies a particular broadcast domain such as a VLAN or VNID in the VxLAN case. An EVPN instance consists of one or more broadcast domains. VNID (VxLAN network identifier): identifies Layer 2 segments and maintains Layer 2 isolation between the segments, allowing the addressing of up to 16 million logical networks in the same administrative domain. The status codes are explained in Table 1-3.
Next Hop	IP address of the nexthop for this route.
Metric	Multiple-Exit Discriminator (MED). If there are multiple paths to the same destination from a single routing protocol, then the multiple paths have the same administrative distance and the best path is selected based on this metric. The path with the lowest metric is selected as the optimal path and installed in the routing table.
LocPrf	This value is used only with iBGP sessions within the local autonomous system to determine if a route towards a destination is the “best” one. The path with the highest local preference is preferred.
Weight	This field applies only to routes within an individual router. If a route was learned from a peer, it has a default weight of 0. All routes generated by the local router have a weight of 32,768.
Path	The autonomous systems through which the prefix advertisement passed. The origin codes are explained in Table 1-4.
Peer	Neighbor address.
Total number of prefixes	The total number of prefixes listed.

show bgp l2vpn evpn prefix-route

Use this command to display the Type-5 prefix remote routes.

Command Syntax

```
show bgp l2vpn evpn prefix-route <vrf WORD | rd Word>
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 4.1.

Example

```
rtr1#show bgp l2vpn evpn prefix-route

RD[300:1]
ESI                 Eth-Tag Prefix-Length    IP-Address
GW-IP Address                              L3VNID    Nexthop          Encap
0                    45001   24              80.80.80.0
0.0.0.0                                    45001     2.2.2.2          VxLAN
0                    45001   24              90.90.90.0
0.0.0.0                                    45001     2.2.2.2          VxLAN
0                    45001   64              8001::
::                                         45001     2.2.2.2          VxLAN
0                    45001   64              9001::
::                                         45001     2.2.2.2          VxLAN
rtr1#
```

show bgp l2vpn evpn summary

Use this command to display a summary of BGP EVPN neighbor status.

Command Syntax

```
show bgp l2vpn evpn summary
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 17
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
8.8.8.8                 4   100  111        112      17      0       0  00:53:03              3      0      0      3      0
9.9.9.9                 4   100  110        110      17      0       0  00:52:10             15      0     13      2      0
13.13.13.13             4   100  132        109      17      0       0  00:51:57              4      0      2      2      0

Total number of neighbors 3

Total number of Established sessions 3
```

The start of the output shows:

- The BGP router identifier and the local router AS number.
- The BGP table version tracks the local BGP table version. Any time the BGP best path algorithm executes, the table version increments.
- BGP AS-PATH entry and community entries.

Table 1-6 explains the fields for each neighbor entry.

Table 1-6: neighbor fields

Field	Description
Neighbor	IP address of peer.
V	BGP version of peer.
AS	Autonomous system number of peer.
MsgRcvd	Messages received since the BGP connection was established.
MsgSent	Messages sent since the BGP connection was established.

Table 1-6: neighbor fields (Continued)

Field	Description
TblVer	Last version of the local router's BGP database advertised to the peer.
InQ	Received messages waiting in the input queue for further processing.
OutQ	Messages waiting in the output queue to be sent.
Up/Down	Connection up time in the interface.
State/PfxRcd	If the TCP session is up and the BGP peers have formed an adjacency, this field shows how many prefixes have been received from the remote neighbor. Other states: Idle: The local router has not allocated resources for the peer connection, so incoming connection requests are refused Idle (Admin): The peer has shut down Idle (PfxCt): Prefix overflow Idle (G-shut): Graceful shutdown Connect: BGP is waiting for the TCP connection to complete Active: the local router is trying to establish a TCP connection to the remote peer. You might see this if the local peer has been configured, but the remote peer is unreachable or has not been configured. OpenSent: BGP is waiting for an open message from its peer OpenConfirm: BGP received an open message from the peer and is now waiting for a keepalive or notification message. If BGP receives a keep alive message from the peer, the state changes to established. If the message is a notification, the state changes to idle. Established: BGP is ready to exchange update, notification, and keep alive messages with its peer Invalid: The session state is invalid.
AD	Number of EVPN type 1 Ethernet Auto-discovery routes: Only originated for multi-homed sites. Type 1 routes allow fast convergence where PE devices can change the next-hop adjacencies for all MAC addresses associated with a particular Ethernet Segment and aliasing where traffic can be balanced across multiple egress points
MACIP	Number of EVPN type 2 MAC/IP routes: Endpoint reachability information, including MAC and IP addresses of the endpoints.
MCAST	Number of EVPN type 3 Inclusive Multicast routes: Broadcast, Unknown Unicast and Multicast (BUM) traffic.
ESI	Number of EVPN type 4 Ethernet Segment Routes: Used in multi-homing for Designated Forwarder Election. The Designated Forwarder sends BUM traffic to the CE on a particular Ethernet Segment.

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 1-7 explains the output fields.

Table 1-7: 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 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 before OcNOS version 4.1.

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 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    ----             ----
45001    ----         L3   NW        ----                            ---- ----      1.1.1.1          3.3.3.3

Total number of entries are 7
```

Table 1-8 explains the fields in the output.

Table 1-8: VxLAN fields

Field	Description
VNID	VxLAN network identifier.
Type	NW - Network Port: VxLAN tunnel AC - Access Port: Host connection
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.

Example

```
#show nvo vxlan access-if-config
     nvo vxlan access-if port-vlan xe1 2
      map vnid 100
access-if-description member-port with xe1 as vlan 2
shutdown
mac 0000.0000.1111
mac 0000.0000.aaaa  ip 12.12.12.1
map qos-profile cos-to-queue ac_port_ingress
map qos-profile queue-color-to-cos ac_port_egress
     !
     nvo vxlan access-if port-vlan po1 6
      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>|)
```

Parameters

`<1-16777215>`	VxLAN network identifier

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show nvo vxlan  arp-cache
VxLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
10       12.12.12.2       0000.0000.2222 Static Local   ----
Total number of entries are 1
```

Table 1-9 explains the fields in the output.

Table 1-9: 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 Local: Learned by data plane source learning Dynamic Remote: 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
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.

Note: Due to a limitation in the hardware, the transmit packet counters includes the BUM traffic received on that port.

Note: To see the statistics, you must enable the hardware profile for the access-port with the hardware-profile filter (Qumran1) command with the `ac-lif` parameter.

Note: All the expected packets might not be seen in the output of this command.

Command Syntax

```
show nvo vxlan counters access-port (port IFNAME | port-vlan IFNAME VLAN_ID_RANGE
  (| inner-vlan INNER_VLAN_ID) | all)
```

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 VLAN_RANGE option is introduced in OcNOS version 5.0.

Example

```
#show nvo vxlan  counters access-port port-vlan xe13 10

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            : 0
  Tx Vxlan Arp discard count           : 0
  Tx Vxlan Nd discard count            : 0
  Rx Vxlan Arp Request count           : 0
```

```
Tx Vxlan Arp Request count           : 0
Rx Vxlan Arp Reply count             : 5
Tx Vxlan Arp Reply count             : 1
Rx Vxlan Neighbor Solicitation count : 0
Tx Vxlan Neighbor Solicitation count : 0
Rx Vxlan Neighbor Advertisement count: 3
Tx Vxlan Neighbor Advertisement count: 0
```

Table 1-10 explains the fields in the output.

Table 1-10: access port counters

Field	Description
RX: packets	Number of packets received on an access-interface.
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.
Tx 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 received from neighbor.
Rx Vxlan Neighbor Advertisement count	Number of Neighbor Advertisement that is received from neighbor.
Tx Vxlan Neighbor Advertisement count	Number of 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.

Command Syntax

```
show nvo vxlan counters network-port (dst A.B.C.D | ALL)
```

Parameters

A.B.C.D	Tunnel IPv4 address
ALL	All addresses

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show nvo vxlan counters network-port dst 2.2.2.2

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  : 0
       bytes    : 0
  TX:  packets  : 3570006
       bytes    : 406980684

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
```

Table 1-11 explains the each network entry fields.

Table 1-11: 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.

Table 1-11: show nvo vxlan counters network-port output fields (Continued)

Field	Description
TX: packets	Number of hello packets transmitted to neighbor.
TX: bytes	Number of hello packets transmitted to neighbor in bytes transmitted.
Tx VxLAN Arp discard count	Number of discarded Arp that is transmitted to peer.
Tx VxLAN Nd discard count	Number of discarded ND that is transmitted to peer.
Tx Vxlan Arp Request count	Number of request ARP that is transmitted to peer.
Tx Vxlan Arp Reply count	Number of replied ARP which is transmitted to peer.
Tx Vxlan Neighbor Solicitation count	Number of replied ND that is received from neighbor.
Tx Vxlan Neighbor Advertisement count	Number of Neighbor Advertisement which is transmitted to peer.

show evpn

Use this command to display the L3 VNI, L2 VNI and IRB interface mapping.

Command Syntax

```
show evpn l3vni-map
show evpn irb-status
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 4.1 and enhanced in OcNOS version 7.0.0 to show the output in ascending order of L3VNI.

Example

```
OcNOS#show evpn l3vni-map
 L3VNI        L2VNI        IRB-interface
==================================
 1603         1601         irb1601
 1604         1602         irb1602
 1605         ---          ---
 1606         ---          ---
 1607         ---          ---
 1608         ---          ---
 1609         ---          ---
 1610         ---          ---


OcNOS#

rtr3#show evpn irb-status
IRB is ACTIVE in Hardware
```

show nvo vxlan mac-table

Use this command to display the host MAC address table. Use the hardware option to see the ageout time for the dynamically learn macs.

Command Syntax

```
show nvo vxlan mac-table (vnid <1-16777215>|) (summary | hardware |)
```

Parameters

`<1-16777215>`	VxLAN network identifier
`summary`	Display a count of MAC addresses
`hardware`	Display hardware information

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3 and was adjusted in OcNOS version 7.0.0 to show the output in ascending order of VNID.

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
Dyanamic Remote      -----             ---
10         ---       ---    ---          0000.0000.aa11 3.3.3.3
Dyanamic 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
Dyanamic Remote      -----             ---

Total number of entries are 7
```

Table 1-12 explains the fields in the output.

Table 1-12: MAC table fields

Field	Description
VNID	VxLAN network identifier
Interface	Interface name
VlanId	VLAN identifier
Mac-Addr	MAC address
VTEP-Ip	VTEP identifier
Type	How a host learns a MAC/IP pair: Dynamic Remote: 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 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 1-12: MAC table fields

Field	Description
AccessPortDesc	Access port description.
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

```
#show nvo vxlan  static host state
VxLAN Static Host Information
============================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     Ifname       Outer-Vlan Inner-vlan Ip-Addr                                   Mac-
Addr       Status
_____________________________________________________________________________________
_____________________
10       xe13         10         ---        12.12.12.2
0000.0000.2222 Active
10       xe17         10         ---        11.11.11.1
0000.0000.bbbb Inactive
10       xe17         30         40         0.0.0.0
0000.1111.2222 Inactive

Total number of entries are 5
```

Table 1-13 explains the output fields.

Table 1-13: Static host fields

Field	Description
VNID	VxLAN network identifier
Ifname	Interface name
Vlan	VLAN name
Ip-Addr	IP address

Table 1-13: 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 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 1-14 explains the output fields.

Table 1-14: 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 nvo vxlan route-count

Use this command to display the VxLAN active route (MAC-IP 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 the GW-IPv4, and Prefix IPv4 in OcNOS version 1.3.5. This command was adjusted in OcNOS version 7.0.0 to show the output in ascending order of VNID.

Example

```
VTEP1#show nvo vxlan route-count
VXLAN Active route count information
====================================
Max supported route count   : 131072


-------------------------------------------
 VNID     Total    MACONLY  MACIPv4  MACIPv6
-------------------------------------------
 10       1        1        0        0
 99       0        0        0        0
 100      1        1        0        0


Active route count: 2
Total number of entries are 3
VTEP1#
```

Table 1-15 explains the output fields.

Table 1-15: 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.
VNID	VNID is used to identify Layer 2 segments and to maintain Layer 2 isolation between the segments.

Table 1-15: show nvo vxlan route-count output details

Field	Description
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 based on vni-name.

Command Syntax

```
show nvo vxlan vni-name (WORD)
```

Parameters

WORD	VxLAN id name

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 nvo vxlan xconnect

Use this command to display the VPWS xconnect details of the MTU, AC-NW connections, and network status.

Command Syntax

```
show nvo vxlan xconnect
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
MH-VTEP3#sh nvo vxlan xconnect
EVPN Xconnect Info
======================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                           Remote       Connection-Details
=============================== ============
================================================================================
VPN-ID       EVI-Name      MTU   VPN-ID     Source       Destination                  PE-IP         MTU   Type   NW-Status
================================= =============
================================================================================
92           ----          1500  91         xe2 92       --- Single Homed Port ---    102.1.1.1     1500  AC-NW  NW-SET
94           ----          1500  93         xe2 94       --- Single Homed Port ---    102.1.1.1     1500  AC-NW  NW-SET
38052        ----          1500  38051      xe2 380 716  00:11:22:33:00:00:00:55:66:77 101.1.1.1     1500  AC-NW  NW-SET
                                                                                     102.1.1.1     1500  ----   ----
39012        ----          1500  39011      xe2 390 715  00:00:00:12:34:90:90:00:00:00 101.1.1.1     1500  AC-NW  NW-SET
                                                                                     102.1.1.1     1500  ----   ----
Total number of entries are 4
```

Table 1-16 explains the output fields.

Table 1-16: show nvo vxlan xconnect fields

Field	Description
Local VPN-ID	Source VPWS ID
Local EVI-Name	EVI name of the local VPWS ID
Local MTU	Local MTU
Remote VPN-ID	Remote (target) VPWS ID
Source	Source AC port

Table 1-16: show nvo vxlan xconnect fields (Continued)

Field	Description
Destination	Destination AC port: “Single Homed Port” or ESI value if it is multi-homed port
PE-IP	IP Address of the provider edge.
MTU	Remote MTU
Type	Connection details of the AC port
NW-Status	Connection details of the NW port
Total number of entries	Total number of entries listed.

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 a version before OcNOS version 4.1.

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
      shutdown
     !
```

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; `nvo vxlan id xconnect` was introduced in OcNOS version 5.1.

Example

```
#sh running-config  nvo vxlan
!
evpn vxlan multihoming enable
!
evpn esi hold-time 100
!
nvo vxlan enable
!
nvo vxlan irb
!
evpn irb-forwarding anycast-gateway-mac 0000.0000.ff10
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan mac-ageing-time 10
!
nvo vxlan id 10 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp vrfred
 evpn irb1
!
nvo vxlan id 100 xconnect target-vxlan-id 200
 vxlan host-reachability-protocol evpn-bgp vrfblue
!
nvo vxlan access-if port-vlan xe13 10
 no shutdown
 map vnid 10
 mac 0000.0000.2222  ip 12.12.12.2
!
nvo vxlan access-if port-vlan xe17 10
 no shutdown
 map vnid 10
 mac 0000.0000.bbbb  ip 11.11.11.1
!
nvo vxlan access-if port-vlan xe17 30
```

```
 no shutdown
 map vnid 10
 mac 0000.1111.2222
!
nvo vxlan access-if port-vlan xe1 11
 description member-port xe1 with vlan 11
 no shutdown
 map vnid 10
!
nvo vxlan access-if port-vlan xe1 10
 no shutdown
 map vnid 10
!
nvo vxlan access-if port-vlan xe1 12
 no shutdown
 map vnid 10
!
nvo vxlan access-if port-vlan xe9 100
 no shutdown
 map vnid 100
!
```

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

Command Mode

NVO mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#nvo vxlan id 3 ingress-replication inner-vid-disabled
(config-nvo)#vxlan host-reachability-protocol evpn-bgp Blue
```

CHAPTER 2 VXLAN Quality of Service Commands

This chapter describes the VXLAN commands for Quality of Service (QoS):

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
```

cos queue

Use this command to configure user defined mapping for Class of Service (CoS) values and output queues.

Use the `no` form of this command to remove the mapping.

Command Syntax

```
cos <0-7> queue <0-7>
no cos <0-7>
```

Parameters

`cos <0-7>`	Sets the CoS value, ranging from 0 to 7, which prioritizes traffic based on its importance.
`queue <0-7>`	Determines the output queue, ranging from 0 to 7, where matched traffic is directed.

Default

If the mapping between CoS and queue values is not configured, it follows a one-to-one mapping.

Command Mode

Ingress CoS Map mode

Applicability

Introduced in OcNOS version 3.0.

Example

```
OcNOS(config)#qos profile cos-to-queue ac_port_ingress
OcNOS(config-ingress-cos-map)#cos 1 queue 7
OcNOS(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 3.0.

Example

```
(config)#qos profile dscp-to-queue nw_profile
(config-ingress-dscp-map)#dscp 50 queue 1
(config-ingress-dscp-map)#no dscp 50
```

l2 queue dscp

Use this command to configure or update user defined mapping for queue to dscp for egress L2 Traffic over VxLAN tunnel.

Use the `no` form of this command to delete the mapping.

Command Syntax

```
l2 queue <0-7> (color (green|yellow|red|all)|) dscp <0-63>
no l2 queue <0-7> (color (green|yellow|red|all|)|)
```

Parameters

`<0-7>`	Queue values
`<0-63>`	Select DSCP value
`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 dscp value is one-one default mapping if it is not configured.

Command Mode

egress-dscp-encap-map mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
#configure terminal
(config)#qos profile dscp-encap DSCP_ENCP
(config-egress-dscp-encap-map)#l2 queue 1 dscp 32
(config-egress-dscp-encap-map)#l2 queue 5 color green dscp 16
(config-egress-dscp-encap-map)#no l2 queue 1
```

l3 dscp dscpEncap

Use this command to configure or update user defined mapping for dscp to dscp for egress L3 Traffic over VxLAN tunnel.

Use the `no` form of this command to delete the mapping.

Command Syntax

```
l3 dscp <0-63> dscpEncap <0-63>
no l3 dscp <0-63>
```

Parameters

`<0-63>`	DSCP values

Default

Default dscp and dscp encap value is one-one default mapping if it is not configured.

Command Mode

egress-dscp-encap-map mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
#configure terminal
(config)#qos profile dscp-encap DSCP_ENCP
(config-egress-dscp-encap-map)#l3 dscp 10 dscpEncap 48
(config-egress-dscp-encap-map)#no l3 dscp 10
```

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.

We can map the following qos profile type on IRB interface.

- dscp-to-queue for ingress traffic and
- dscp-to-dscp profile for egress traffic

from/to the IRB interface.

On doing no map qos-profile, default profile is applied to the IRB interface.

In order for dscp-to-dscp - qos profile type to be effective on L3 IRB interface , "qos remark dscp" is enabled globally.

Note: Default profile is applied to all the IRB interface only when qos is enabled.

Command Syntax

```
map qos-profile (dscp-to-dscp | dscp-to-queue) <NAME>
no map qos-profile (dscp-to-dscp | dscp-to-queue) <NAME>
```

Parameters

`NAME`	Profile name

Default

By default, the default dscp-to-queue and dscp-to-dscp profile is attached to all IRB interface and if user changes the value in the default profile, it will affect to all the IRB interfaces and L3 interfaces.

Command Mode

IRB_IF_Mode

Default

Default queue and cos value is one-one default mapping if it is not configured.

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
#configure terminal
(config)#interface irb 1
(config-irb-if)# map qos-profile dscp-to-dscp DSCP_DSCP
(config-irb-if)# no map qos-profile dscp-to-dscp DSCP_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 before OcNOS version 3.0.

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 before OcNOS version 3.0.

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 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 3.0.

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 3.0.

Example

```
(config)#qos profile cos-to-queue ac_port_ingress
(config)#no qos profile cos-to-queue ac_port_ingress
```

qos profile dscp-encap

Use this command to create new profiles or to update "default" profiles for dscp-dscpEncap and Queue to DSCP value.

Use the `no` form of this command to remove the profiles.

Note: "default" profiles can only be updated and not be deleted. User can create/delete user-defined profiles.

Note: Use this command to configure QoS profile for both VxLAN L2 and VxLAN L3 traffic.

Note: This profile will be mapped to nvo vxlan tunnel at the egress direction of the VTEP.

Note: Default Qos profile dscp-encap would take preference than default qos profile queue-color-to dscp when no user-defined qos profile is configured on the nvo vxlan tunnel mode.

Command Syntax

```
qos profile dscp-encap (NAME | default)
no qos profile dscp-encap NAME
```

Parameters

`NAME`	Profile name
`default`	update global mapping profile

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
#configure terminal
(config)#qos profile dscp-encap DSCP_ENCP
(config)#no qos profile dscp-encap DSCP_ENCP
```

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 3.0.

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 3.0.

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.

Note: The profile name of "default" is not supported for VxLAN QoS.

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 3.0.

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 3.0.

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 3.0.

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
```

show qos-profile type dscp-encap

Use this command to show all the default and configured QoS profiles configurations.

Command Syntax

```
show qos-profile type dscp-encap (| NAME)
```

Parameters

NAME	Profile name

Command Mode

Exec Mode.

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
VTEP1#show qos-profile type dscp-encap
profile name: default
profile type: dscp-encap
configured mapping:
Detailed mapping:
 L3 DSCP to DSCP-ENCAP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  INPUT | OUTPUT   |  INPUT | OUTPUT  |  INPUT | OUTPUT  |  INPUT | OUTPUT
 -------|--------  | -------|-------- | -------|-------- | -------|--------
  DSCP  | DSCP     |  DSCP  | DSCP    |  DSCP  | DSCP    |  DSCP  | DSCP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  0        0       |  16       16       |  32       32       |  48       48
  1        1       |  17       17       |  33       33       |  49       49
  2        2       |  18       18       |  34       34       |  50       50
  3        3       |  19       19       |  35       35       |  51       51
  4        4       |  20       20       |  36       36       |  52       52
  5        5       |  21       21       |  37       37       |  53       53
  6        6       |  22       22       |  38       38       |  54       54
  7        7       |  23       23       |  39       39       |  55       55
  8        8       |  24       24       |  40       40       |  56       56
  9        9       |  25       25       |  41       41       |  57       57
  10       10       |  26       26       |  42       42       |  58       58
  11       11       |  27       27       |  43       43       |  59       59
  12       12       |  28       28       |  44       44       |  60       60
  13       13       |  29       29       |  45       45       |  61       61
  14       14       |  30       30       |  46       46       |  62       62
  15       15       |  31       31       |  47       47       |  63       63

 L2 Queue + Color to DSCP-ENCAP
 ---------------+--------| ----------------+--------| -------------+--------
      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     8      |  1       yellow    8      |  1       red       8
  2       green     16      |  2       yellow    16      |  2       red       16
  3       green     24      |  3       yellow    24      |  3       red       24
  4       green     32      |  4       yellow    32      |  4       red       32
  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

profile name: DSCP_ENCAP
profile type: dscp-encap
profile attached to 0 instances
configured mapping:
 l3 dscp 16 dscpEncap 24
 l2 queue 3 color all dscp 40
Detailed mapping:
 L3 DSCP to DSCP-ENCAP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  INPUT | OUTPUT   |  INPUT | OUTPUT  |  INPUT | OUTPUT  |  INPUT | OUTPUT
 -------|--------  | -------|-------- | -------|-------- | -------|--------
  DSCP  | DSCP     |  DSCP  | DSCP    |  DSCP  | DSCP    |  DSCP  | DSCP
 -------+--------  | -------+-------- | -------+-------- | -------+--------
  0        0       |  16       24       |  32       32       |  48       48
  1        1       |  17       17       |  33       33       |  49       49
  2        2       |  18       18       |  34       34       |  50       50
  3        3       |  19       19       |  35       35       |  51       51
  4        4       |  20       20       |  36       36       |  52       52
  5        5       |  21       21       |  37       37       |  53       53
  6        6       |  22       22       |  38       38       |  54       54
  7        7       |  23       23       |  39       39       |  55       55
  8        8       |  24       24       |  40       40       |  56       56
  9        9       |  25       25       |  41       41       |  57       57
  10       10       |  26       26       |  42       42       |  58       58
  11       11       |  27       27       |  43       43       |  59       59
  12       12       |  28       28       |  44       44       |  60       60
  13       13       |  29       29       |  45       45       |  61       61
  14       14       |  30       30       |  46       46       |  62       62
  15       15       |  31       31       |  47       47       |  63       63

 L2 Queue + Color to DSCP-ENCAP
 ---------------+--------| ----------------+--------| --------------+--------
      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     8      |  1       yellow    8      |  1       red       8
  2       green     16      |  2       yellow    16      |  2       red       16
  3       green     40      |  3       yellow    40      |  3       red       40
  4       green     32      |  4       yellow    32      |  4       red       32
  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

VTEP1#
```

show running-config interface irb

Use this command to display the current QoS 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 4.2.

Example

```
VTEP1#sh run int irb1010
!
interface irb1010
 ip vrf forwarding FAX
 ip address 10.210.32.1/24
 ipv6 address 3001:ab8:1::1/64
 qos map-profile dscp-to-queue DSCP_QUE
 qos map-profile dscp-to-dscp DSCP_DSCP
!
```

Index

P

Q

R

S

T

U

V

W