

VXLAN

This chapter describes Arista's VXLAN implementation. Sections in this chapter include:

- [Section 22.1: VXLAN Introduction](#)
- [Section 22.2: VXLAN Description](#)
- [Section 22.3: VXLAN Configuration](#)
- [Section 22.4: VXLAN Command Descriptions](#)

22.1 VXLAN Introduction

Virtual Extensible LAN (VXLAN) is a networking technology that encapsulates MAC-based Layer 2 Ethernet frames within Layer 3 UDP packets to aggregate and tunnel multiple layer 2 networks across a Layer 3 infrastructure. VXLAN scales up to 16 million logical networks and supports layer 2 adjacency across IP networks. Multicast transmission architecture is used for broadcast, multicast, and unknown unicast traffic.

For a list of VXLAN feature support in a specific EOS release, consult the appropriate release notes here: <https://www.arista.com/en/support/software-download>.

For a list of VXLAN feature support by platform in the latest EOS release, see <https://www.arista.com/en/support/product-documentation/supported-features>

Note VXLAN and NAT cannot co-exist.

22.2 VXLAN Description

These sections describe VXLAN architecture, the data objects that comprise a VXLAN network, and process of bridging packets through a VXLAN network.

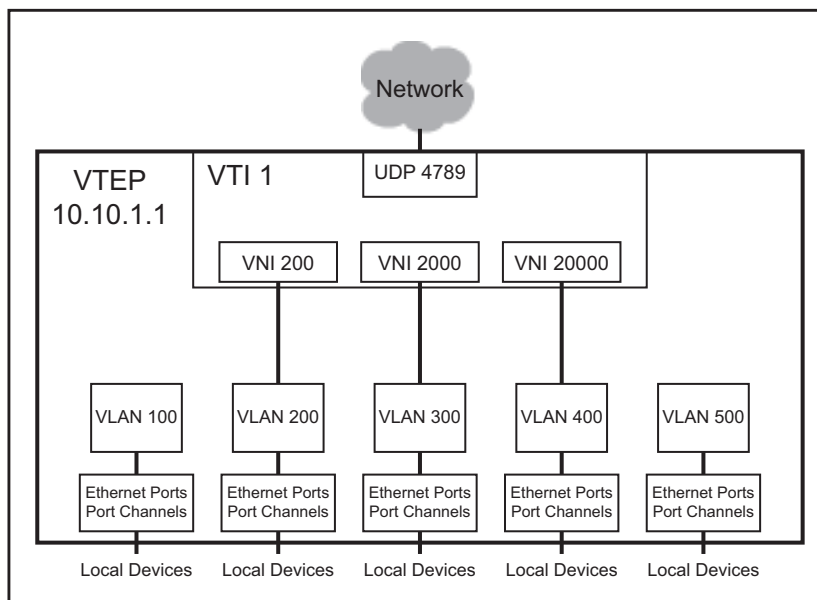
- [Section 22.2.1: VXLAN Architecture](#)
- [Section 22.2.2: VXLAN Processes](#)
- [Section 22.2.3: Multicast and Broadcast over VXLAN](#)
- [Section 22.2.4: VXLAN Gateway](#)
- [Section 22.2.5: VXLAN and MLAG](#)
- [Section 22.2.6: Data Structures](#)

22.2.1 VXLAN Architecture

The VXLAN architecture extends an L2 network by connecting VLANs from multiple hosts through UDP tunnels called VXLAN segments. VXLAN segments are identified by a 24-bit virtual network identifier (VNI). Within a host, each VLAN whose network is extended to other hosts is associated with a VNI. An extended L2 network comprises the devices attached to VLANs from all hosts that are on VLANs that are associated with the same VNI.

[Figure 22-1](#) displays the data objects that comprise a VXLAN implementation on a local host.

Figure 22-1: VXLAN Architecture



- **VXLAN Tunnel End Point (VTEP):** a host with at least one VXLAN Tunnel Interface (VTI).
- **VXLAN Tunnel Interface (VTI):** a switchport linked to a UDP socket that is shared with VLANs on various hosts. Packets bridged from a VLAN to the VTI are sent out the UDP socket with a VXLAN header. Packets arriving on the VTI through the UDP socket are demuxed to VLANs for bridging.
- **Virtual Network Identifier (VNI):** a 24-bit number that distinguishes between the VLANs carried on a VTI. It facilitates the multiplexing of several VLANs over a single VTI.

VNIs can be expressed in digital or dotted decimal formats. VNI values range from 1 to 16777215 or from 0.0.1 to 255.255.255.

The network in [Figure 22-1](#) has the following assignments:

- VTEP IP address of 10.10.1.1
- UDP port of 4789
- One VTI that supports three VXLAN segments (UDP tunnels): VNI 200, VNI 2000, and VNI 20000
- Five VLANs, of which three VLANs can communicate with remote devices over Layer 2.

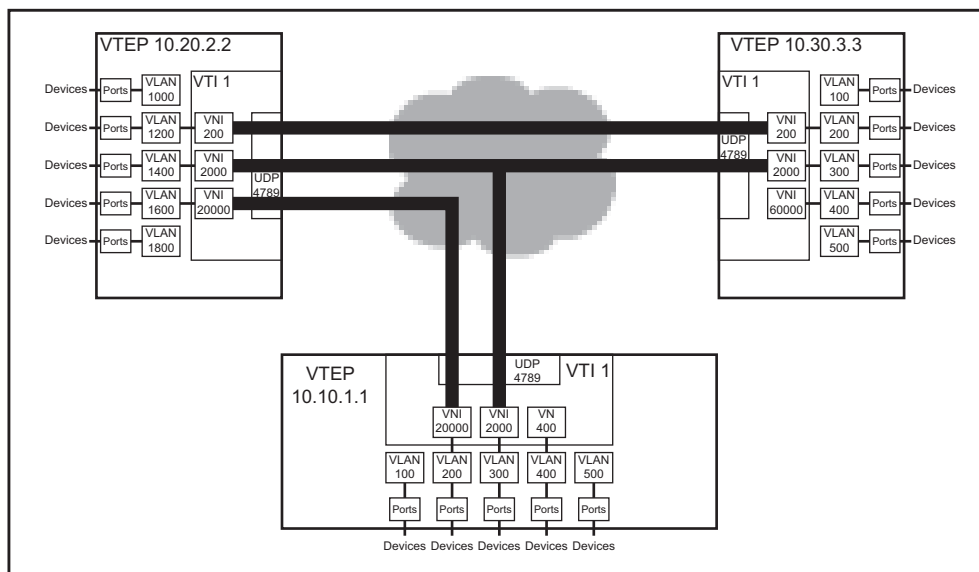
22.2.2 VXLAN Processes

When a packet enters a VLAN from a member (ingress) port, the VLAN learns the source address by adding an entry to the MAC address table that associates the source to the ingress-port. The VLAN then searches the table for destination address. If the MAC address table lists the address, the packet is sent out the corresponding port. If the MAC address table does not list the address, the packet is flooded to all ports except the ingress port.

VXLANs extend VLANs through the addition of a VXLAN address table that correlates remote MAC addresses to their port and resident host IP address. Packets that are destined to a remote device are sent to the VXLAN tunnel interface (VTI), which is the switchport that is linked to the UDP socket. The packet is encapsulated with a VXLAN header which includes the VNI associated with the VLAN and the IP mapping of the destination host. The packet is sent through a UDP socket to the destination VTEP IP. The VTI on the remote host extracts the original packet and bridges it to the VLAN associated with the VNI on the remote host.

UDP port 4789 is recognized as the VXLAN socket and listed as the destination port on the UDP packets. The UDP source port field is filled with a hash of the inner header to facilitate load balancing.

Figure 22-2: VXLAN Implementation



[Figure 22-2](#) displays a configuration that includes three VTEPs. The VXLAN defines three inter-host L2 networks. The VLANs that comprise the networks include:

- VNI 200: VTEP 10.20.2.2: VLAN 1200 and VTEP 10.30.3.3: VLAN 200
- VNI 2000: VTEP 10.10.1.1: VLAN 300, VTEP 10.20.2.2: VLAN 1400, and VTEP 10.30.3.3: VLAN 300
- VNI 20000: VTEP 10.10.1.1: VLAN 200, and VTEP 10.20.2.2: VLAN 1600

VXLAN Routing

VXLAN routing is enabled by creating a VLAN interface on the VXLAN-enabled VLAN and assigning an IP address to the VLAN interface. The IP address serves as VXLAN gateway for devices that are accessible from the VXLAN-enabled VLAN.

22.2.3 Multicast and Broadcast over VXLAN

These sections describe multicast and broadcast over VXLANs. Multicast packet flooding describes broadcast and multicast transmission by associating a multicast group to a VTI through a configuration command. Head-end Replication (HER) optimizes flooding of inter VTEP broadcast, unknown unicast and broadcast (BUM) traffic by using hardware and flood lists to perform replication on the supported platform.

22.2.3.1 Multicast Packet Flooding

Multicast packet flooding is supported with VXLAN bridging without MLAG. A VTI is associated with a multicast group through a configuration command.

VXLAN and Broadcast

When a VLAN receives or sends a broadcast packet the VTI is treated as a bridging domain L2 interface. The packet is sent from this interface on the multicast group associated with the VTI. The VTIs on remote VTEPs that receive this packet extract the original packet, which is then handled by the VLAN associated with the packet's VNI. The VLAN floods the packet, excluding the VTI. When the broadcast results in a response, the resulting packet can be unicast back to the originating VTEP because the VXLAN address table obtained the host MAC to VTEP association from the broadcast packet.

VXLAN and Multicast

A VTI is treated as an L2 interface in the VLAN for handling multicast traffic, which is mapped from the VLAN to the multicast group associated with the VTI. All VTEPs join the configured multicast group for inter-VTEP communication within a VXLAN segment; this multicast group is independent of any other multicast groups that the hosts in the VLAN join.

The IP address space for the inter-host VXLAN communication may be sourced from a different VRF than the address space of the hosts in the VLAN. The multicast group for inter-VTEP transmissions must not be used for other purposes by any device in the VXLAN segment space.

22.2.3.2 Head-end Replication

Head-end replication uses a flood list to support broadcast, unknown unicast, and multicast (BUM) traffic over VXLAN. The flood list specifies a list of remote VTEPs. The switch replicates BUM data locally for bridging across the remote VTEPs specified by the flood list. This data flooding facilitates remote MAC address learning by forwarding data with unknown MAC addresses.

Head-end replication is required for VXLAN routing and to support VXLANs over MLAG.

22.2.4 VXLAN Gateway

A VXLAN gateway is a service that exchanges VXLAN data and packets with devices connected to different network segments. VXLAN traffic must pass through a VXLAN gateway to access services on physical devices in a distant network.

A VXLAN gateway requires the following information:

- An IP address that is designated as the VXLAN interface source.

- VLAN to VNI mapping.
- VTEP list for each VNI.
- A method for handling broadcast, unknown unicast, and multicast (BUM) packets.

Arista switches manually perform VXLAN gateway services. The switch connects to VXLAN gateways that serve other network segments. MAC address learning is performed in hardware from inbound VXLAN packets. BUM packets are supported through one of the methods specified in [Section 22.2.3](#).

22.2.5 VXLAN and MLAG

VXLAN over MLAG provides redundancy in hardware VTEPs. VTI configuration must be identical on each MLAG peer for them to act as a single VTEP. This also prevents the remote MAC from flapping between the remote VTEPs by ensuring that the rest of the network sees a host that is connected to the MLAG interface as residing behind a single VTEP.

Differences between VXLAN bridging and routing implementations over MLAG are applicable for the DCS-7050X series platform.

- VXLAN routing recirculates a packet twice, with the first iteration performing the routing action involving an L2 header rewrite, and the second recirculation performing VXLAN encap and decap operations. Recirculation is achieved by MAC loopback on dedicated loopback interfaces.
- The configuration for VXLAN routing on an MLAG VTEP includes separate Recirc-Channel configuration on both peers. The virtual IP, virtual MAC, and virtual VARP VTEP IP addresses are identical on both peers.

The following VTI elements must be configured identically on both MLAG peers:

- VLAN-VNI mappings
- VTEP IP address of the source loopback interface
- Flood VTEP list used for head-end replication

If OSPF is also in use, configure the OSPF router ID manually to prevent the switch from using the common VTEP IP address as the router ID.

The following rules are observed by MLAG switches so that they behave as a single VXLAN VTEP:

- Only the MLAG peer that receives a packet performs VXLAN encapsulation on it.
- Packets are not VXLAN encapsulated if they are received from the peer link.
- If a packet is decapsulated and sent over the peer link, it should not be flooded to active MLAG interfaces.
- If a packet is sent over the peer link to the CPU, it is not head-end replicated to other remote VTEPs.
- If a packet's destination is the VTEP IP address, it is terminated by the MLAG peer that receives it.

Example

- These commands complete the configuration required for a VXLAN routing deployment.

```
switch(config)#interface Vxlan1
switch(config-if-Vx1)#vxlan source-interface Loopback0
switch(config-if-Vx1)#vxlan udp-port 4789
switch(config-if-Vx1)#vxlan vlan 2417 vni 8358534
switch(config-if-Vx1)#vxlan flood vtep 1.0.1.1 1.0.2.1
switch(config-if-Vx1)#interface Vlan2417
switch(config-if-Vl2417)#ip address 1.0.4.1/24
switch(config-if-Vl2417)#interface Loopback0
switch(config-if-Lo0)#ip address 1.0.1.1/32
switch(config-if-Lo0)#ip routing
switch(config)#interface Recirc-Channel627
switch(config-if-Re627)#switchport recirculation features vxlan
switch(config-if-Re627)#interface Ethernet 1
switch(config-if-Et1)#traffic-loopback source system device mac
switch(config-if-Et1)#channel-group recirculation 627
switch(config-if-Et1)#exit
switch(config)#interface Ethernet 2
switch(config-if-Et2)#traffic-loopback source system device mac
switch(config-if-Et2)#channel-group recirculation 627
switch(config-if-Et2)#
```

Configuring Unconnected Ethernet Interfaces for Recirculation

On systems where bandwidth is not fully used by the front panel ports, unused bandwidth is used for recirculation.

The following example is applicable to the DCS-7050X series platform.

Example

- These commands expose unconnected Ethernet interfaces which are used for recirculation, in order to use them to replace or use along with front panel Ethernet interfaces.

```
switch(config)#service interface unconnected expose
switch(config)#interface UnconnectedEthernet 2
switch(config-if-Ue2)#traffic-loopback source system device mac
switch(config-if-Ue2)#channel-group recirculation 627
```

22.2.6 Data Structures

VXLAN implementation requires two VXLAN tables and a MAC address table accommodation.

22.2.6.1 MAC Address Table VXLAN Support

MAC address table entries correlate MAC addresses with the port upon which packets arrive. In addition to Ethernet and port channels, the port column may specify a VTI for packets that arrive on a VLAN from a remote port through the VXLAN segment.

22.2.6.2 VTEP-MAC Address Table

VTEP-MAC address table entries correlate MAC address with the IP address of the VTEP from where packets bearing the MAC address arrive. The VTI uses this table to determine the destination address for packets that are sent to remote hosts.

22.2.6.3 VNI-VLAN Map

The VNI-VLAN map displays the one-to-one correspondence between the VNIs assigned on the switch and the VLANs to which they are assigned. Each VNI can be assigned to only one VLAN; each VLAN can be assigned a maximum of one VNI. Each VNI-VLAN assignment constitutes a VXLAN segment.

22.3 VXLAN Configuration

These sections describe VXLAN configuration tasks:

- [Section 22.3.1: Configuring the VTI](#)
- [Section 22.3.2: Head End Replication Configuration](#)
- [Section 22.3.3: VXLAN Routing Configuration](#)
- [Section 22.3.4: Configuring VXLAN Routing with Overlay VRFs](#)
- [Section 22.3.5: Configuring VXLAN over MLAG](#)
- [Section 22.3.6: Configuring VXLAN Control Service](#)
- [Section 22.3.7: Configuring VXLAN Multicast Decapsulation](#)
- [Section 22.3.8: VXLAN Rules Support for Mirror ACLs Configuration](#)
- [Section 22.3.9: Displaying VXLAN Configuration](#)

22.3.1 Configuring the VTI

Configuring the VTI enables VXLAN bridging and is a requirement for VXLAN Routing. The following sections describe the steps required to enabling VXLAN bridging by bringing up the VXLAN line protocol. [Section 22.3.3](#) describes the additional steps required to enable VXLAN routing.

Instantiating the VTI and VXLAN Configuration Mode

The **interface vxlan** command places the switch in VXLAN-interface configuration mode for modifying the specified VXLAN tunnel interface (VTI). The command also instantiates the interface if it was not previously created.

VXLAN interface configuration mode is not a group change mode; **running-config** is changed immediately after commands are executed. The **exit** command does not affect the configuration.

Example

- These commands create VXLAN tunnel interface 1, place the switch in VXLAN-interface configuration mode, and display parameters of the new VTI.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
switch(config-if-Vx1)#
```

Assigning an IP address to the VTEP

The **vxlan source-interface** command specifies the loopback interface from which the VTEP derives the source address (IP) that it uses when exchanging VXLAN frames. This address is used by UDP headers to specify source and destination addresses of hosts that send or receive VXLAN encapsulated packets.

There is no default source interface assignment. A valid VXLAN configuration requires the assignment of a loopback interface to the VTEP and the assignment of a valid IP address to the specified interface.

Example

- These commands configure VTI 1 to use IP address 10.25.25.3 (loopback interface 15) as the source interface in the encapsulation fields of outbound VXLAN frames.

```
switch(config)#interface loopback 15
switch(config-if-Lo15)#ip address 10.25.25.3/24
switch(config-if-Lo15)#exit
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan source-interface loopback 15
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan source-interface Loopback15
    vxlan udp-port 4789
switch(config-if-Vx1)#
```

Assigning a UDP Port to the VTEP

Packets bridged to the VTI from a VLAN are encapsulated with a VXLAN header, then sent through a pre-configured UDP port. Packets that arrive through this port are assumed to be VXLAN encapsulated and sent to the bridging domain of the recipient VLAN as determined by the VNI in the VXLAN header and the VNI-VLAN map.

The **vxlan udp-port** command associates a UDP port with the configuration mode VXLAN interface (VTI). By default, UDP port 4789 is associated with the VTI.

Important! UDP port 4789 is reserved by convention for VXLAN usage. Under most typical applications, this parameter should be set to the default value.

Example

- This command associates UDP port 5500 with VXLAN interface 1.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan udp-port 5500
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 5500
switch(config-if-Vx1)#
```

- This command resets the VXLAN interface 1 UDP port association of 4789.

```
switch(config-if-Vx1)#no vxlan udp-port
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
switch(config-if-Vx1)#
```

Assigning a VNI to a VLAN

When a VLAN bridges a packet to the VTI, the packet is encapsulated with a VXLAN header that includes the VNI associated with the VLAN. Packets that arrive on the VTI's UDP socket are bridged to the VLAN that is associated with the VNI specified by the VXLAN header that encapsulates the packet.

The VTI requires a one-to-one correspondence between specified VLANs and VNI values. Commands that assign a new VNI to a previously configured VLAN replace existing VLAN assignment statements in **running-config**. Commands that attempt to assign a VNI value to a second VLAN generate a CLI error.

The **vxlan vlan vni** command associates a VLAN ID with a virtual network identifier (VNI).

Example

- These commands associate VLAN 100 to VNI 100 and VLAN 200 to VNI 10.10.200.

```

switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan vlan 100 vni 100
switch(config-if-Vx1)#vxlan vlan 200 vni 10.10.200
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan udp-port 4789
  vxlan vlan 200 vni 658120
  vxlan vlan 100 vni 100
switch(config-if-Vx1)#vxlan vni notation dotted
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan udp-port 4789
  vxlan vlan 100 vni 0.0.100
  vxlan vlan 200 vni 10.10.200
switch(config-if-Vx1)#

```

Assigning a Multicast Group to the VTI

The VTI maps multicast traffic from its associated VLANs to a specified multicast group. Inter-VTEP multicast communications include all VTEPs that are associated with the specified multicast group, which is independent of any other multicast groups that VLAN hosts may join.

The **vxlan multicast-group** command associates a specified multicast group with the configuration mode VXLAN interface (VTI), which handles multicast and broadcast traffic as a layer 2 interface in a bridging domain.

Example

- This command associates the multicast address of 227.10.1.1 with VTI 1.

```

switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan multicast-group 227.10.1.1
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan multicast-group 227.10.1.1
  vxlan udp-port 4789
switch(config-if-Vx1)#

```

Verifying the VXLAN Configuration

The **show interface vxlan 1** displays the configuration and connection status of the VXLAN

Example

- This command indicates that the VXLAN line protocol status is **up**.

```

switch(config-if-Vx1)#show interface vxlan 1
Vxlan1 is up, line protocol is up (connected)
  Hardware is Vxlan
  Source interface is Loopback15 and is active with 10.25.25.3
  Static vlan to vni mapping is
    [100, 0.0.100]    [200, 10.10.200]
  Multicast group address is 227.1.1.1
switch(config-if-Vx1)#

```

22.3.2 Head End Replication Configuration

Head-end replication is a data distribution method that supports broadcast, unknown unicast, and multicast (BUM) traffic over VXLANs by replicating BUM data locally for transmission to the set of remote VTEPs specified by a flood list. This data flooding facilitates remote MAC address learning through the forwarding of data with unknown MAC addresses.

Each **vxlan flood vtep** statement in *running-config* associates a set of VTEP addresses to an access VNI. A default flood list is also configurable that applies to all VNIs for which a flood list is not configured.

The VTEP flood list is created and modified through the **vxlan flood vtep** command. When configuring VXLAN bridging, the flood list can replace **vxlan multicast-group**.

Example

- These commands create a default VXLAN head-end replication flood list.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan flood vtep 10.1.1.1 10.1.1.2
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan flood vtep 10.1.1.1 10.1.1.2
  vxlan udp-port 4789
switch(config-if-Vx1)#
```

- These commands create VXLAN head-end replication flood lists for the VNIs accessed through VLANs 101 and 102.

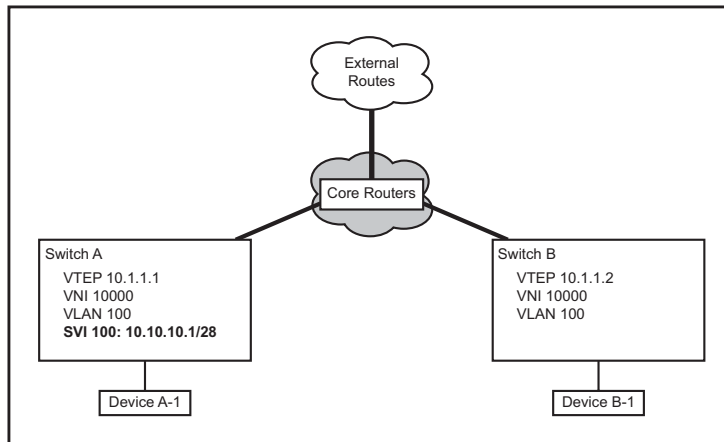
```
switch(config-if-Vx1)#vxlan vlan 101-102 flood vtep 11.1.1.1 11.1.1.2 11.1.1.3
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan flood vtep 10.1.1.1 10.1.1.2
  vxlan vlan 101 flood vtep 11.1.1.1 11.1.1.2 11.1.1.3
  vxlan vlan 102 flood vtep 11.1.1.1 11.1.1.2 11.1.1.3
  vxlan udp-port 4789
switch(config-if-Vx1)#
```

22.3.3 VXLAN Routing Configuration

22.3.3.1 Implementing VXLAN Routing

VXLAN routing is enabled by creating a VLAN interface (SVI) on a VLAN that is associated to a VNI. In [Figure 22-3](#), VXLAN routing is enabled on Switch A by configuring a VLAN interface with an IP address of 10.10.10.1. Packets from Devices A-1 and B-2 that have destinations other than 10.10.10.0/28 are VXLAN-bridged to the default gateway (10.10.10.1), then routed from Switch A.

Figure 22-3: Implementing VXLAN Routing



Example

- These commands configure Switch A to perform VXLAN routing. The example includes OSPF routing that is used for underlay routing.

```

switch-A(config)#route-map vxlanvlan permit 10
switch-A(config-route-map-vxlanvlan)#match interface loopb5
switch-A(config-route-map-vxlanvlan)#exit
switch-A(config)#route-map vxlanvlan permit 20
switch-A(config-route-map-vxlanvlan)#match interface vlan 100
switch-A(config-route-map-vxlanvlan)#exit
switch-A(config)#router ospf 1
switch-A(config-router-ospf)#redistribute connected route-map vxlanvlan
switch-A(config-router-ospf)#exit
switch-A(config)#interface loopback 5
switch-A(config-if-Lo5)#ip address 10.25.25.3/24
switch-A(config-if-Lo5)#exit
switch-A(config)#interface vxlan 1
switch-A(config-if-Vx1)#vxlan source-interface loopback 5
switch-A(config-if-Vx1)#vxlan vlan 100 vni 10000
switch-A(config)#interface vlan 100
switch-A(config-if-Vl100)#ip address 10.10.10.1/28
switch-A(config-if-Vl100)#exit
  
```

22.3.3.2 Configuring Direct VXLAN Routing

In [Figure 22-3](#), VXLAN routing is enabled on Switch A only; Switch B supports VXLAN bridging. Traffic from Switch B devices to the external routes must go through the core route twice: once as they are bridged to is VXLAN gateway and once when routed to its next hop device.

Direct VXLAN routing with VXLAN enabled addresses this issue by configuring each VTEP with all VLANs. This allows packets to be VXLAN-bridged to a local VTEP and routed to remote VTEPs. Indirect routing scales well but is complex to engineer efficiently, and naked routing provides the same scalability to indirect routing. Direct routing leads to the most efficient traffic flows, with the number of virtual subnets or virtual machines increasing at scale, and is thereby optimal from a data plane viewpoint.

The following sections describe conventions required to implement Direct VXLAN Routing, then presents a direct VXLAN routing implementation.

Configuring VARP addresses

For direct routing, an anycast IP address is used as the gateway address on the SVI for a VLAN on all hardware VTEPs associated with that VLAN.

Examples

- These commands configure an IP virtual-router and virtual MAC address.

```
switch(config)#interface Vlan2417
switch(config-if-Vl2417)#ip address 1.0.4.50/24
switch(config-if-Vl2417)#ip virtual-router address 1.0.4.1
switch(config-if-Vl2417)#ip virtual-router mac-address 00:00:11:11:22:22
switch(config)#
```

- These commands configure an IP virtual address (instead of IP virtual-router address) for the VLAN SVI, and a secondary address on the loopback interface for the virtual VTEP IP. The virtual VTEP IP is the logical VTEP hosting the virtual MAC address.

```
switch(config)#interface Vlan2417
switch(config-if-Vl2417)#ip address virtual 1.0.4.1/24
switch(config-if-Vl2417)#exit
switch(config)#interface Loopback0
switch(config-if-Lo0)#ip address 1.0.1.1/32
switch(config-if-Lo0)#ip address 1.0.1.2/32 secondary
switch(config-if-Lo0)#ip virtual-router mac-address 00:00:11:11:22:22
switch(config)#
```

Virtual IP and MAC Addresses

Virtual-router IP addresses can be configured on VLAN interfaces in addition to a primary address. All VTEPs in a direct VXLAN network can be configured with the same virtual router address. This allows devices to use a common IP address as their VXLAN gateway.

The **ip address virtual** command configures a specified address as the primary IPv4 address and as a virtual IP address for the configuration mode VLAN interface. This results in the virtual MAC address (**ip virtual-router mac-address**) assignment to the VLAN interface. In large VXLAN networks, using distinct primary IP addresses for each VTEP limits the number addresses on its subnet for connected hosts. Defining a common virtual IP address for all VTEPs and using that their primary addresses conserves subnet addresses

Example

- These commands specify a virtual router address of 00:00:00:00:00:48 for the switch and, for VLAN 100, a primary address of 10.10.10.10/28 and a virtual IP address of 10.10.10.10.

```
switch(config)#ip virtual-router mac-address 00:00:00:00:00:48
switch(config)#interface vlan 100
switch(config-if-Vl100)#ip address virtual 10.10.10.10/28
switch(config-if-Vl100)#show active
  interface Vlan100
    ip address virtual 10.10.10.10/28
switch(config-if-Vl100)#
```

Virtual VTEP Configuration

A virtual VTEP address is specified by configuring a secondary address on the loopback interface designated as the VXLAN's source interface. All VTEPs in the direct routing topology share the same virtual VTEP address.

You must also configure the secondary VTEP IP on the flood-list of the downstream VXLAN VTEPS as shown below.

Example

- These commands specify a primary (10.1.1.1) and virtual VTEP address (10.2.2.2).

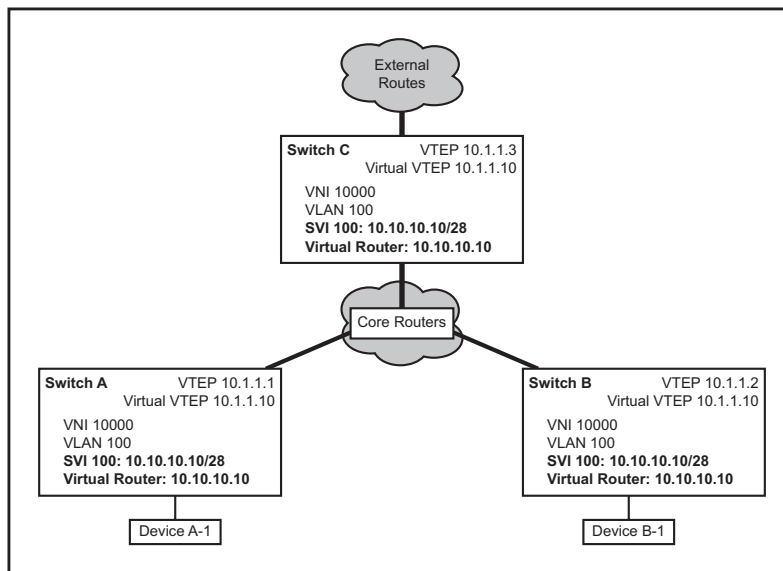
```
Switch1
switch(config)#interface loopback 5
switch(config-if-Lo5)#ip address 10.1.1.1/24
switch(config-if-Lo5)#ip address 10.2.2.2/24 secondary
switch(config-if-Lo5)#show active
  interface Loopback5
    ip address 10.1.1.1/24
    ip address 10.2.2.2/24 secondary
switch(config-if-Lo5)#exit
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan source-interface loopback 5
switch(config-if-Vx1)#show active
  interface Vxlan1
    vxlan source-interface Loopback5
    vxlan udp-port 4789
    vxlan vlan 100 vni 10000
switch(config-if-Vx1)#

Switch2
switch(config)#interface vxlan1
switch(config-if-Vx1)#vxlan flood vtep 10.1.1.1
switch(config-if-Vx1)#vxlan flood vtep 10.2.2.2
```

Direct VXLAN Topology

Figure 22-4 displays a direct VXLAN topology, where each VTEP is configured with the same set of VNIs, VLAN interfaces, and virtual VTEP address.

Figure 22-4: Direct VXLAN Routing



Example

- These commands configure VXLAN parameters for Switch-A in [Figure 22-4](#).

```

switch-A(config)#route-map vxlanvlan permit 10
switch-A(config-route-map-vxlanvlan)#match interface loopb5
switch-A(config-route-map-vxlanvlan)#exit
switch-A(config)#route-map vxlanvlan permit 20
switch-A(config-route-map-vxlanvlan)#match interface vlan 100
switch-A(config-route-map-vxlanvlan)#exit
switch-A(config)#router ospf 1
switch-A(config-router-ospf)#redistribute connected route-map vxlanvlan
switch-A(config-router-ospf)#exit
switch-A(config)#ip virtual-router mac-address 00:00:00:00:00:48
switch-A(config)#interface loopback 5
switch-A(config-if-Lo5)#ip address 10.1.1.3/24
switch-A(config-if-Lo5)#ip address 10.1.1.10/24 secondary
switch-A(config-if-Lo5)#exit
switch-A(config)#interface vxlan 1
switch-A(config-if-Vx1)#vxlan source-interface loopback 5
switch-A(config-if-Vx1)#vxlan vlan 100 vni 10000
switch-A(config)#interface vlan 100
switch-A(config-if-Vl100)#ip address virtual 10.10.10.10/28
switch-A(config-if-Vl100)#exit
  
```

22.3.3.3 Configuring VXLAN VTEP counters

The switch platforms which use this feature are:

- DCS-7050X
- DCS-7250X
- DCS-7300X

The VXLAN VTEP counters feature enables a device to count VXLAN packets received and sent by the device on a per VTEP basis. Specifically, it enables the device to count bytes and packets that are getting encapsulated and decapsulated as they are passing through.

The counters are logically split up in the two VXLAN directions. Encapsulated on the device and directed to the core, “encap” counters count packets coming from the edge. Decapsulated on the device and heading towards the edge, “decap” counters count packets coming from the core.

To be able to count VXLAN packets the device has to support VXLAN and have a VXLAN interface correctly configured.

Examples

- This command configures the enabling of VXLAN VTEP counters for encap.

```
switch(config)#hardware counter feature vtep encap
switch(config)#
```

- This command configures the disabling of VXLAN VTEP counters for encap.

```
switch(config)#no hardware counter feature vtep encap
switch(config)#
```

- This commands configures the enabling of VXLAN VTEP counters for decap.

```
switch(config)#hardware counter feature vtep decap
switch(config)#
```

- This commands configures the disabling of VXLAN VTEP counters for decap.

```
switch(config)#no hardware counter feature vtep decap
switch(config)#
```

22.3.4 Configuring VXLAN Routing with Overlay VRFs

The switch platforms which use this feature are:

- DCS-7050X
- DCS-7250X
- DCS-7300X

VXLAN SVIs configured in non-default VRFs are supported with VXLAN routing using overlay VRFs. Overlay SVIs are configured in non-default VRFs but underlay SVIs, which provide IP connectivity between VTEPs, must remain in the default VRF. VXLAN routing is deployable by allowing users to configure separate overlay routing domains using VRFs per tenant, thereby allowing support for overlapping IP addresses in the overlay. This provides separation between overlay and underlay traffic, including simpler and cleaner protocol configuration, without using complicated route-maps to control distribution of prefixes to peers in the overlay VRFs and underlay SVIs. IPv4 based VXLAN routing is currently supported.

22.3.5 Configuring VXLAN over MLAG

VTI configuration must be identical on each MLAG peer for them to act as a single VTEP.

The following VTI elements must be configured identically on both MLAG peers:

VLAN-VNI Mappings

Configure identical VLAN to VNI mappings on both MLAG peers using the **vxlan vlan vni** command.

Example

- These commands associate VLAN 100 to VNI 100 and VLAN 200 to VNI 10.10.200.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan vlan 100 vni 100
switch(config-if-Vx1)#vxlan vlan 200 vni 10.10.200
switch(config-if-Vx1)#
```


VTEP IP Address of the Source Loopback Interface

Configure the same VTEP IP address for the source loopback interface on both MLAG peers using the **vxlan source-interface** command.

Example

- These commands configure a primary VTEP address.


```
switch(config)#interface loopback 5
switch(config-if-Lo5)#ip address 10.1.1.1/24
switch(config-if-Lo5)#exit
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan source-interface loopback 5
switch(config-if-Vx1)#
```

Flood VTEP List

Configure the same VTEP flood list on both MLAG peers using the **vxlan flood vtep** command.

Example

- These commands create a default VXLAN head-end replication flood list.


```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan flood vtep 10.1.1.1 10.1.1.2
switch(config-if-Vx1)#
```

OSPF Configuration

If OSPF is in use, configure the OSPF router ID using the **router-id (OSPFv2)** command to prevent the switch from using the common VTEP IP address as the router ID.

Example

- These commands assign 10.0.0.1 as the OSPFv2 router ID.


```
switch(config)#router ospf 100
switch(config-router-ospf)#router-id 10.0.0.1
switch(config-router-ospf)#
```

22.3.6 Configuring VXLAN Control Service

The VXLAN Control Service (VCS) provides a mechanism by which hardware VTEPs share states between each other in order to establish VXLAN tunnels, without the need for a multicast control plane. This feature enables the use of a VCS client.

Examples

- These commands connect a switch to the VCS running on CVX. The server host IP address is the management IP address of the CVX controller or the IP address that CVX is listening on for client connections.

```
switch(config)#management cvx
switch(config-mgmt-cvx)#server host 172.27.6.248
switch(config-mgmt-cvx)#no shutdown
switch(config-mgmt-cvx)#
```

- These commands configure the VXLAN interface, except for the multicast group configuration, in order to learn from the controller.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan controller-client
switch(config-if-Vx1)#
```

22.3.7 Configuring VXLAN Multicast Decapsulation

The switch platforms which use this feature are:

- DCS-7050X
- DCS-7250X
- DCS-7300X

VXLAN multicast decapsulation enables VTEPs that support Head End Replication (HER). Multicast encapsulated Broadcast/Unknown/Multicast (BUM) packets terminate VTEPs from remote VTEPs that do not support HER.

Examples

- These commands enable VXLAN multicast decapsulation.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan multicast-group decap 230.1.1.1
switch(config-if-Vx1)#
```

- These commands disable VXLAN multicast decapsulation.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#no vxlan multicast-group decap 230.1.1.1
switch(config-if-Vx1)#
```

22.3.8 VXLAN Rules Support for Mirror ACLs Configuration

The switch platforms which use this feature are:

- DCS-7150S

VXLAN rules support for mirror ACLs configuration permit VXLAN deep inspection rules to be specified in the mirroring ACLs when the switch is operating in normal mode.

Examples

The following are examples of VXLAN rules specified in mirroring ACLs.

- These commands permit all VXLAN traffic (udp protocol and destination port 4789).

```
switch(config)#ip access-list miracl
switch(config-acl-miracl)#permit vxlan any any
switch(config-acl-miracl)#
```

- These commands permit VXLAN traffic with vni 1001 only.

```
switch(config)#ip access-list miracl
switch(config-acl-miracl)#permit vxlan any any vni 1001 0x000000
switch(config-acl-miracl)#
```

- These commands deny VXLAN traffic with vni 0x1000 through 0x100f.

```
switch(config)#ip access-list miracl
switch(config-acl-miracl)#permit vxlan any any vni 0x1000 0x100f
switch(config-acl-miracl)#
```

22.3.9 Displaying VXLAN Configuration

The following section describes the commands that control the display format of VNIs and the commands that list VXLAN configuration and transmission information.

Configuring VNI Display Format

The **vxlan vni notation dotted** command configures the switch to display VNIs in dotted decimal notation. VNI values range from 1 to 16777215 in decimal notation and from 0.0.1 to 255.255.255 in dotted decimal notation.

The command affects the VNI number display in all **show** commands, including **show running-config**. Commands that include VNI as a parameter may use decimal or dotted decimal notation regardless of the setting of this command. By default, show commands display VNI number in decimal notation.

Examples

- These commands configure the switch to display vni numbers in dotted decimal notation, then displays a configuration that includes a VNI setting.

```
switch(config)#vxlan vni notation dotted
switch(config)#interface vxlan 1
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
    vxlan vlan 333 vni 3.4.5
switch(config-if-Vx1)#
```

- These commands configure the switch to display vni numbers in decimal notation, then displays a configuration that includes a VNI setting.

```
switch(config)#no vxlan vni notation dotted
switch(config)#interface vxlan 1
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
    vxlan vlan 333 vni 197637
switch(config-if-Vx1)#
```

MAC Address Table

The MAC address table indicates a MAC address from a device on a remote host by indicating Vx interface as the port that corresponds to the address.

Example

- This command displays a MAC address table that includes entries of devices from remote hosts by specifying Vx1 as the corresponding port.

```
switch>show mac address-table
          Mac Address Table
-----
Vlan      Mac Address      Type      Ports      Moves      Last Move
----      -
1         0050.5682.6725   DYNAMIC   Et16       1          0:02:01 ago
1         0050.568e.58e9   DYNAMIC   Et23       2          0:08:53 ago
1         0050.56a0.474a   DYNAMIC   Et16       1          0:18:04 ago
51        0000.0051.0004   DYNAMIC   Et5        1          12 days, 1:02:44 ago
51        0000.0051.0005   DYNAMIC   Et5        1          12 days, 1:02:44 ago
51        0000.0051.0101   DYNAMIC   Vx1        1          12 days, 0:17:30 ago
51        0000.0051.0102   DYNAMIC   Vx1        1          12 days, 0:17:30 ago
61        0000.0061.0005   DYNAMIC   Et5        1          12 days, 1:02:44 ago
Total Mac Addresses for this criterion: 8

          Multicast Mac Address Table
-----
Vlan      Mac Address      Type      Ports
----      -
Total Mac Addresses for this criterion: 0
switch>
```

VXLAN MAC Address Table

VXLAN MAC address table entries correlate MAC addresses accessible through remote VTEPs with the local VLAN and the IP address of the VTEP through which the addressed device is accessed. The VTI uses this table when constructing the VXLAN encapsulation to specify the destination IP address of the recipient VTEP and the VNI segment through which the device's remote VLAN is accessed.

The **show vxlan address-table** command displays the VXLAN MAC address table.

Example

- This command displays the VXLAN address table.

```
switch>show vxlan address-table
          Vxlan Mac Address Table
-----
Vlan      Mac Address      Type      Prt  Vtep      Moves      Last Move
----      -
51        0000.0051.0101   DYNAMIC   Vx1  10.25.2.12  1          4 days, 0:37:14 ago
51        0000.0051.0102   DYNAMIC   Vx1  10.25.2.12  1          4 days, 0:37:14 ago
51        0000.0051.0103   DYNAMIC   Vx1  10.25.2.12  1          4 days, 0:37:14 ago
51        0000.0051.0104   DYNAMIC   Vx1  10.25.2.12  1          4 days, 0:37:14 ago
51        0000.0051.0105   DYNAMIC   Vx1  10.25.2.12  1          4 days, 0:37:14 ago
61        0000.0061.0103   DYNAMIC   Vx1  10.25.2.12  1          4 days, 0:37:14 ago
61        0000.0061.0104   DYNAMIC   Vx1  10.25.2.12  1          4 days, 0:37:14 ago
61        0000.0061.0105   DYNAMIC   Vx1  10.25.2.12  1          4 days, 0:37:14 ago
switch>
```

VXLAN MAC Address Table

The **show vxlan vtep** command displays information about remote VTEPs that the configured VTI has discovered and with whom it has exchanged packets.

Example

- These commands display the VTEPs that have exchanged data with the configured VTI.

```
switch>show vxlan vtep
Remote vteps for Vxlan1:
10.52.2.12
Total number of remote vteps: 1
switch>
```

VXLAN Counters

The **clear vxlan counters** command resets the VXLAN counters. The **show vxlan counters** command displays the VXLAN counters.

Example

- This command displays the VXLAN counters

```
switch>show vxlan counters software
encap_bytes:3452284
encap_pkts:27841
encap_read_err:1
encap_discard_runt:0
encap_discard_vlan_range:0
encap_discard_vlan_map:0
encap_send_err:0
encap_timeout:1427
decap_bytes_total:382412426
decap_pkts_total:2259858
decap_bytes:0
decap_pkts:0
decap_runt:0
decap_pkt_filter:45128
decap_bytes_filter:5908326
decap_discard_vxhdr:0
decap_discard_vlan_map:2214730
decap_timeout:0
decap_sock_err:1
switch>
```

22.4 VXLAN Command Descriptions

VXLAN Global Configuration Commands

- interface vxlan
- ip address virtual
- vxlan vni notation dotted

VXLAN Interface Configuration Commands

- vxlan flood vtep
- vxlan multicast-group
- vxlan multicast-group decap
- vxlan source-interface
- vxlan udp-port
- vxlan vlan vni

VXLAN Display and Clear Commands

- clear vxlan counters
- show service vxlan
- show vxlan address-table
- show vxlan counters
- show vxlan flood vtep
- show vxlan vtep

clear vxlan counters

The **clear vxlan counters** command resets the VXLAN counters.

Command Mode

Privileged EXEC

Command Syntax

```
clear vxlan counters ROUTE_TYPE
```

Parameters

- ***ROUTE_TYPE*** Specifies the type of VXLAN counter reset by the command.
 - **software** Command resets software counters.
 - **varp** Command resets virtual-ARP counters.

Related Commands

- **show vxlan counters** displays the VXLAN counters.

Example

- This command resets the VXLAN counters

```
switch#clear vxlan counters software
switch#show vxlan counters software
encap_bytes:0
encap_pkts:0
encap_read_err:0
encap_discard_runt:0
encap_discard_vlan_range:0
encap_discard_vlan_map:0
encap_send_err:0
encap_timeout:0
decap_bytes_total:0
decap_pkts_total:0
decap_bytes:0
decap_pkts:0
decap_runt:0
decap_pkt_filter:0
decap_bytes_filter:0
decap_discard_vxhdr:0
decap_discard_vlan_map:0
decap_timeout:0
decap_sock_err:0
switch#
```

interface vxlan

The **interface vxlan** command places the switch in VXLAN-interface configuration mode for modifying the specified VXLAN tunnel interface (VTI). The command also instantiates the interface if it was not previously created.

VXLAN interface configuration mode is not a group change mode; **running-config** is changed immediately after commands are executed. The **exit** command does not affect the configuration.

The **no interface vxlan** deletes the specified VTI interface, including its configuration statements, from **running-config**. The **default interface vxlan** command removes all configuration statements for the specified VTI from **running-config** without deleting the interfaces.

Command Mode

Global Configuration

Command Syntax

```
interface vxlan vx_range
no interface vxlan vx_range
default interface vxlan vx_range
```

Parameter

- **vx_range** VXLAN interface number. The only permitted value is 1.

Commands Available in link-flap Configuration Mode

- **vxlan multicast-group**
- **vxlan source-interface**
- **vxlan udp-port**
- **vxlan vlan vni**

Example

- These commands create VXLAN tunnel interface 1, place the switch in VXLAN-interface configuration mode, then display parameters of the new VTI.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan udp-port 4789
switch(config-if-Vx1)#
```

- This command exits VXLAN-interface configuration mode, placing the switch in global configuration mode.

```
switch(config-if-Vx1)#exit
switch(config)#
```


ip address virtual

The **ip address virtual** command configures a specified address as the primary IPv4 address and as a virtual IP address for the configuration mode VLAN interface. The address resolves to the virtual MAC address configured through the **ip virtual-router mac-address** command. The command includes a subnet designation that is required in primary IP address assignments.

This command is typically used in VXLAN routing configurations as an alternative to assigning a unique IP address to each VTEP. All existing IPv4 addresses must be removed from the interface before executing this command.

The **no ip address virtual** and **default ip address virtual** commands remove the IPv4 address and virtual IP assignment from the configuration mode interface by deleting the **ip address virtual** command from *running-config*.

Removing the IPv4 address assignments from an interface disables IPv4 processing on that port.

Command Mode

Interface-VLAN Configuration

Command Syntax

```
ip address virtual ipv4_subnet
no ip address virtual
default ip address virtual
```

Parameters

- *ipv4_subnet* IPv4 and subnet address (CIDR or address-mask notation).

Related Commands

- **ip address**
- **ip address virtual**
- **ip virtual-router mac-address**

Example

- This command configures 10.10.10.1 as the IPv4 address and virtual address for VLAN 100.

```
switch(config-if-Vl100)#show active
interface Vlan100
  ip address virtual 10.10.10.1/28
switch(config-if-Vl100)#
```

show service vxlan

The **show service vxlan** command displays the status of the Vxlan Control Service (VCS) and the received (from all connected VTEPs) and advertised (to all connected VTEPs) MAC address reachability information.

Command Mode

EXEC

Command Syntax

```
show service vxlan [status | switch [SWITCH_TYPE] | vni [VNI_INFO]]
```

Parameters

- **SWITCH_TYPE** displayed by switch type. Options include:
 - *word* hostname, IP address, or ID of the switch.
 - **all** all switches.
- **VNI_INFO** displayed with VNI information. Options include:
 - **advertised** advertised MAC addresses.
 - **received** received MAC addresses.

Example

- This command displays the status of the VCS.

```
switch(config)#show service vxlan status
Vxlan Controller Service is      : stopped
Mac learning                     : Control plane
Resync period                    : 300 seconds
Resync in progress               : No
Capability                       : VXLAN v4 overlay routing
                                VXLAN v4 overlay indirect routing

fm319(config-if-Vx1)#show service vxlan status
Vxlan Controller Service is      : stopped
Mac learning                     : Control plane
Resync period                    : 300 seconds
Resync in progress               : No
Capability                       : VXLAN v4 overlay routing
                                VXLAN v4 overlay indirect routing

switch(config)#
```

show vxlan address-table

The **show vxlan address-table** command displays the VXLAN address table. Entries are created by extracting information from packets received from remote VTEPs.

The VXLAN address table correlates MAC addresses that are accessible through remote VTEPs with the local VLAN and the IP address of the VTEP through which the addressed device is accessible. The VTI uses this table when constructing the VXLAN encapsulation fields to specify the destination IP address of the recipient VTEP and the VNI segment through which the device's remote VLAN is accessed.

Command Mode

EXEC

Command Syntax

```
show vxlan address-table [ENTRY_TYPE] [MAC_ADDR] [VLANS] [REMOTE_VTEP]
```

Parameters

- **ENTRY_TYPE** command filters display by entry type. Options include:
 - <no parameter> all table entries.
 - **configured** static entries; includes unconfigured VLAN entries.
 - **dynamic** entries learned through packet receipts.
 - **static** entries entered by CLI commands.
 - **unicast** entries with unicast MAC address.
- **MAC_ADDR** command uses MAC address to filter displayed entries.
 - <no parameter> all MAC addresses table entries.
 - **address mac_address** displays entries with specified address (dotted hex notation – H.H.H).
- **VLANS** command filters display by VLAN.
 - <no parameter> all VLANs.
 - **vlan v_num** VLAN specified by *v_num*.
- **REMOTE_VTEP** Filters entries by IP address of the remote VTEPs. Options include:
 - <no parameter> all items.
 - **vtep ipaddr_1 [ipaddr_2...ipaddr_n]** Identifies VTEPs by their IP address.

Example

- This command displays the VXLAN address table.

```
switch>show vxlan address-table
```

```
      Vxlan Mac Address Table
```

```
-----
```

Vlan	Mac Address	Type	Prt	Vtep	Moves	Last Move
----	-----	----	---	----	-----	-----
51	0000.0051.0101	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago
51	0000.0051.0102	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago
51	0000.0051.0103	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago
51	0000.0051.0104	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago
51	0000.0051.0105	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago
61	0000.0061.0102	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago
61	0000.0061.0103	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago
61	0000.0061.0104	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago
61	0000.0061.0105	DYNAMIC	Vx1	10.25.2.12	1	4 days, 0:37:14 ago

```
switch>
```

show vxlan counters

The **show vxlan counters** command displays the VXLAN counters.

Command Mode

EXEC

Command Syntax

```
show vxlan counters ROUTE_TYPE
```

Parameters

- **ROUTE_TYPE** Specifies the type of VXLAN counter displayed by the command.
 - **software** Command displays software routers.
 - **varp** Command displays virtual-ARP counters.
 - **vtep** Command displays counters for VTEPs which are identified by their IP address. An optional keyword allows the user to view a single direction of the counters:
 - **encap** “encap” counters count packets coming from the edge, encapsulated on the device and directed to the core.
 - **decap** “decap” counters count packets coming from the core, decapsulated on the device and heading towards the edge.

Related Commands

- **clear vxlan counters** resets the VXLAN counters.

Example

- This command displays the VXLAN counters for software routers.

```
switch>show vxlan counters software
encap_bytes:3452284
encap_pkts:27841
encap_read_err:1
encap_discard_runt:0
encap_discard_vlan_range:0
encap_discard_vlan_map:0
encap_send_err:0
encap_timeout:1427
decap_bytes_total:382412426
decap_pkts_total:2259858
decap_bytes:0
decap_pkts:0
decap_runt:0
decap_pkt_filter:45128
decap_bytes_filter:5908326
decap_discard_vxhdr:0
decap_discard_vlan_map:2214730
decap_timeout:0
decap_sock_err:1
switch>
```

- This command displays the VXLAN counters for VTEPs.

```
switch>show vxlan counters vtep
```

VTEP	Decap Bytes	Decap Known Unicast Packets	Decap BUM Packets	Decap Drop or Exception Packets
1.0.14.1	62526968000	312632701	312636979	2
1.0.16.1	800	2	6	312279633
1.0.23.1	800	2	6	2
unlearned	0	0	0	0

VTEP	Encap Bytes	Encap Packets	Encap Drop or Exception Packets
1.0.14.1	30579308814	268239551	2
1.0.16.1	1140	10	2
1.0.23.1	0	0	0

```
switch>
```

show vxlan flood vtep

The **show vxlan flood vtep** command displays the flood list that the switch is using to perform head-end replication. Head-end replication is a data distribution method that supports broadcast, unknown unicast, and multicast (BUM) traffic over VXLANs by replicating BUM data locally for transmission to the set of remote VTEPs that a flood list specifies. The command displays the VLAN ID that references the configured VNIs (**vxlan vlan vni**).

The flood list is determined by the **vxlan flood vtep** command.

Command Mode

EXEC

Command Syntax

```
show vxlan flood vtep [VLANS]
```

Parameters

- **VLANS** command filters display by the reference VLAN.
 - <no parameter> all VLANs.
 - **vlan v_range** VLANs specified by *v_range*.

Valid *v_range* formats include number, range, or comma-delimited list of numbers and ranges.

Guidelines

The command displays flood list contents only when the VLAN line protocol status is **up**.

Related Commands

- **vxlan flood vtep** configures the flood list.

Example

- These commands display the VTEPs that have exchanged data with the configured VTI.

```
switch>show vxlan flood vtep vlan 100-102
      Vxlan Flood Vtep Table
```

```
-----
Vlan   Ip Address
-----
100    3.3.3.3
101    11.1.1.1      11.1.1.2      11.1.1.3
102    11.1.1.1      11.1.1.2      11.1.1.3
      12.1.1.1
```

```
switch>
```

show vxlan vtep

The **show vxlan vtep** command displays information about remote VTEPs that the configured VTI has discovered and with whom it has exchanged packets.

Command Mode

EXEC

Command Syntax

```
show vxlan vtep
```

Example

- These commands display the VTEPs that have exchanged data with the configured VTI.

```
switch>show vxlan vtep
Remote vteps for Vxlan1:
10.52.2.12
Total number of remote vteps: 1
switch>
```


vxlan flood vtep

The **vxlan flood vtep** command supports VXLAN head-end replication by creating or modifying a list that specifies remote VTEPs to which the switch bridges replicated traffic. Head-end replication is a data distribution method that supports broadcast, unknown unicast, and multicast (BUM) traffic over VXLANs by replicating BUM data locally for transmission to the set of remote VTEPs that a flood list specifies. This data flooding facilitates remote MAC address learning through the forwarding of data with unknown MACs.

Each **vxlan flood vtep** statement in *running-config* associates a set of VTEP addresses to an access VNI. A default flood list is also configurable that applies to all VNIs for which a flood list is not configured. The **vxlan flood vtep** command is available in the following formats to create or modify corresponding *running-config* statements:

- **vxlan flood vtep** creates a statement for a specified VNI and replaces existing statements for that VNI.
- **vxlan flood vtep add** modifies an existing flood statement by adding the specified VTEPs. This statement creates a list if it references a VNI that has no flood statement.
- **vxlan flood vtep remove** modifies an existing flood statement by deleting the specified VTEPs. This statement has no effect if it references a VNI that has no flood statement.

The **vxlan flood vtep** command specifies a VNI by referencing its associated VLAN ID (**vxlan vlan vni**). The command provides these options for specifying the reference VLANs:

- **a single VLAN**: creates or modifies a single statement referenced by the command.
- **a range of VLANs**: creates or modifies all statements referenced by the VLAN range.
- **no VLAN**: creates or modifies the default list

The **no vxlan flood vtep** and **default vxlan flood vtep** commands remove the specified flood list by deleting the corresponding **vxlan flood vtep** statements from *running-config*. Commands that specify a VLAN range remove all corresponding statements.

Command Mode

Interface-VXLAN Configuration

Command Syntax

```
vxlan [ACCESS_VNI] flood vtep [MODIFY] VTEP_1 [VTEP_2] ... [VTEP_N]
no vxlan [ACCESS_VNI] flood vtep
default vxlan [ACCESS_VNI] flood vtep
```

Parameters

- **ACCESS_VNI** VLAN ID associated to the flood list's target VNI. Value ranges from 1 to 4094.
 - <no parameter > default list.
 - **vlan vlan_range** List of VLANs. (Number, range, comma-delimited list of numbers and ranges). Numbers range from 1 to 4094.
- **MODIFY** Statement modification method. Options include:
 - <no parameter > creates new list for specified VLANs. Current list is overwritten.
 - **add** specified VTEPs are added to existing list.
 - **remove** specified VTEPs are deleted from existing list.
- **VTEP_X** IPv4 address of VTEPs that are added or removed from the list.

Example

- These commands create a default VXLAN head-end replication flood list.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan flood vtep 10.1.1.1 10.1.1.2
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan flood vtep 10.1.1.1 10.1.1.2
  vxlan udp-port 4789
switch(config-if-Vx1)#
```

- These commands create VXLAN head-end replication flood lists for the VNIs accessed through VLANs 101 and 102.

```
switch(config-if-Vx1)#vxlan vlan 101-102 flood vtep 11.1.1.1 11.1.1.2 11.1.1.3
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan flood vtep 10.1.1.1 10.1.1.2
  vxlan vlan 101 flood vtep 11.1.1.1 11.1.1.2 11.1.1.3
  vxlan vlan 102 flood vtep 11.1.1.1 11.1.1.2 11.1.1.3
  vxlan udp-port 4789
switch(config-if-Vx1)#
```

- These commands add two VTEPs for the VNI access through VLAN 102.

```
switch(config-if-Vx1)#vxlan vlan 102 flood vtep add 12.1.1.1
switch(config-if-Vx1)#show active
interface Vxlan1
  vxlan flood vtep 10.1.1.1 10.1.1.2
  vxlan vlan 101 flood vtep 11.1.1.1 11.1.1.2 11.1.1.3
  vxlan vlan 102 flood vtep 11.1.1.1 11.1.1.2 11.1.1.3 12.1.1.1
  vxlan udp-port 4789
switch(config-if-Vx1)#
```

vxlan multicast-group

The **vxlan multicast-group** command associates a specified multicast group with the configuration mode VXLAN interface (VTI), which handles multicast and broadcast traffic as a layer 2 interface in a bridging domain.

The VTI maps multicast traffic from its associated VLANs to the specified multicast group. Inter-VTEP multicast communications include all VTEPs that are associated with the specified multicast group, which is independent of any other multicast groups that VLAN hosts may join.

A VTI can be associated with one multicast group. By default, a VTI is not associated with any multicast group.

The **no vxlan multicast-group** and **default vxlan multicast-group** commands removes the multicast group – VTI association by removing the **vxlan multicast-group** command from *running-config*.

Command Mode

Interface-VXLAN Configuration

Command Syntax

```
vxlan multicast-group group_addr
no vxlan multicast-group
default vxlan multicast-group
```

Parameters

- **group_addr** IPv4 address of multicast group. Dotted decimal notation of a valid multicast address.

Related Commands

- **interface vxlan** places the switch in VXLAN interface configuration mode.

Examples

- This command associates the multicast address of 227.10.1.1 with VTI 1.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan multicast-group 227.10.1.1
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan multicast-group 227.10.1.1
    vxlan udp-port 4789
switch(config-if-Vx1)#
```

- This command changes VTI 1's multicast group association.

```
switch(config-if-Vx1)#vxlan multicast-group 227.10.5.5
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan multicast-group 227.10.5.5
    vxlan udp-port 4789
switch(config-if-Vx1)#
```

- This command removes the multicast group association from VTI 1.

```
switch(config-if-Vx1)#no vxlan multicast-group
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
switch(config-if-Vx1)#
```

vxlan multicast-group decap

The **vxlan multicast-group decap** command enables VXLAN multicast decapsulation.

VTEPs are enabled by VXLAN multicast decapsulation, supporting Head End Replication (HER). Multicast encapsulated Broadcast/Unknown/Multicast (BUM) packets terminate VTEPs from remote VTEPs that do not support HER.

The **no vxlan multicast-group decap** and **default vxlan multicast-group decap** commands disable VXLAN multicast decapsulation.

Command Mode

Interface-VXLAN Configuration

Command Syntax

```
vxlan multicast-group decap group_addr
no vxlan multicast-group decap
default vxlan multicast-group decap
```

Parameters

- *group_addr* IPv4 address of multicast group. Dotted decimal notation of a valid multicast address.

Examples

- This command enables VXLAN multicast decapsulation.

```
switch(config)#interface vxlan 1
switch(config-config-if-Vx1)#vxlan multicast-group decap 230.1.1.1
switch(config-config-if-Vx1)#
```

- This command disables VXLAN multicast decapsulation.

```
switch(config)#interface vxlan 1
switch(config-config-if-Vx1)#no vxlan multicast-group decap 230.1.1.1
switch(config-config-if-Vx1)#
```

vxlan source-interface

The **vxlan source-interface** command specifies the interface from which the configuration mode VXLAN interface (VTI) derives the source address (IP) that it uses when exchanging VXLAN frames. There is no default source interface assignment.

The **no vxlan source-interface** and **default vxlan source-interface** commands remove the source interface assignment from the configuration mode VXLAN interface by deleting the corresponding **ip vxlan source-interface** command from *running-config*.

Command Mode

Interface-VXLAN Configuration

Command Syntax

```
vxlan source-interface INT_NAME
no vxlan source-interface
default vxlan source-interface
```

Parameters

- **INT_NAME** Interface type and number. Options include:
 - **loopback I_num** Loopback interface specified by *I_num*.

Guidelines

A VXLAN interface is inoperable without the source-interface assignment.

Related Commands

- **interface vxlan** places the switch in VXLAN interface configuration mode.

Example

- These commands configure VTI 1 to use the IP address 10.25.25.3 as the source address of outbound VXLAN frames.

```
switch(config)#interface loopback 15
switch(config-if-Lo15)#ip address 10.25.25.3/24
switch(config-if-Lo15)#exit
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan source-interface loopback 15
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan source-interface Loopback15
    vxlan udp-port 4789
switch(config-if-Vx1)#
```

vxlan udp-port

The **vxlan udp-port** command associates a UDP port with the configuration mode VXLAN interface (VTI). By default, UDP port 4789 is associated with the VTI.

Packets bridged to the VTI from a VLAN are encapsulated with a VXLAN header that includes the VNI associated with the VLAN and the IP address of the VTEP that connects to the recipient, then sent through the UDP port. Packets that arrive through the UDP port are sent to the bridging domain of the recipient VLAN as determined by the VNI number in the VXLAN header and the interface's VNI-VLAN map.

The **no vxlan udp-port** and **default vxlan udp-port** command restores the default UDP port association (4789) on the configuration mode interface by deleting the corresponding **vxlan udp-port** command from *running-config*.

Command Mode

Interface-VXLAN Configuration

Command Syntax

```
vxlan udp-port port_id
no vxlan udp-port
default vxlan udp-port
```

Parameters

- *port_id* UDP port number. Value ranges from 1024 to 65535.

Guidelines

UDP port 4789 is reserved by convention for VXLAN usage. Under most typical applications, this parameter should be set to the default value.

Related Commands

- **interface vxlan** places the switch in VXLAN interface configuration mode.

Example

- This command associates UDP port 5500 with VXLAN interface 1.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan udp-port 5500
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 5500
switch(config-if-Vx1)#
```

- This command resets the VXLAN interface 1 UDP port association of 4789.

```
switch(config-if-Vx1)#no vxlan udp-port
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
switch(config-if-Vx1)#
```

vxlan vlan vni

The **vxlan vlan vni** command associates a VLAN ID with a virtual network identifier (VNI). A VNI is a 24-bit number that is assigned to a VLAN to distinguish it from other VLANs that are on a VXLAN tunnel interface (VTI). VNI values range from 1 to 16777215 in decimal notation and from 0.0.1 to 255.255.255 in dotted decimal notation.

When a VLAN bridges a packet to the VTI, the packet is encapsulated with a VXLAN header that includes the VNI that is associated with the VLAN. Packets that arrive on the VTI's UDP socket are bridged to the VLAN that is associated with the VNI specified by the VXLAN header that encapsulates the packet.

The VTI requires a one-to-one correspondence between specified VLANs and VNI values. Commands that assign a new VNI to a previously configured VLAN replace the existing VLAN assignment statement in **running-config**. Commands that attempt to assign a VNI value to a second VLAN generate a CLI error.

The **no vxlan vlan vni** and **default vxlan vlan vni** commands remove the specified VLAN-VNI association from the configuration mode interface by deleting the corresponding **vxlan vlan** command from **running-config**.

Command Mode

Interface-VXLAN Configuration

Command Syntax

```
vxlan vlan vlan_id vni vni_id
no vxlan vlan vlan_id vni [vni_id]
default vxlan vlan vlan_id vni [vni_id]
```

Parameters

- **vlan_id** number of access VLAN. Value ranges from 1 to 4094.
- **vni_id** VNI number. Valid formats: decimal <1 to **16777215**> or dotted decimal <**0.0.1** to **255.255.255**>.

Example

- These commands associate VLAN 100 to VNI 100 and VLAN 200 to VNI 10.10.200.

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan vlan 100 vni 100
switch(config-if-Vx1)#vxlan vlan 200 vni 10.10.200
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
    vxlan vlan 200 vni 658120
    vxlan vlan 100 vni 100
switch(config-if-Vx1)#vxlan vni notation dotted
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
    vxlan vlan 200 vni 10.10.200
    vxlan vlan 100 vni 0.0.100
switch(config-if-Vx1)#
```

vxlan vni notation dotted

The **vxlan vni notation dotted** command configures the switch to display VNIs in dotted decimal notation. A virtual network identifier (VNI) is a 24-bit number that is assigned to a VLAN to distinguish it from other VLANs that are on a VXLAN tunnel interface. VNI values range from 1 to 16777215 in decimal notation and from 0.0.1 to 255.255.255 in dotted decimal notation.

The command affects the VNI number display in all **show** commands, including **show running-config**. Commands that include VNI as a parameter may use decimal or dotted decimal notation regardless of the setting of this command. By default, show commands display VNI number in decimal notation.

The **no vxlan vni notation dotted** and **default vxlan vni notation dotted** commands restore the default setting of displaying vni numbers in decimal notation by deleting the **vxlan vni notation dotted** command from *running-config*.

Command Mode

Global Configuration

Command Syntax

```
vxlan vni notation dotted
no vxlan vni notation dotted
default vxlan vni notation dotted
```

Examples

- These commands configure the switch to display vni numbers in dotted decimal notation, then displays a configuration that includes a vni setting.

```
switch(config)#vxlan vni notation dotted
switch(config)#interface vxlan 1
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
    vxlan vlan 333 vni 3.4.5
switch(config-if-Vx1)#
```

- These commands configure the switch to display vni numbers in decimal notation, then displays a configuration that includes a vni setting.

```
switch(config)#no vxlan vni notation dotted
switch(config)#interface vxlan 1
switch(config-if-Vx1)#show active
interface Vxlan1
    vxlan udp-port 4789
    vxlan vlan 333 vni 197637
switch(config-if-Vx1)#
```