VXLAN Interfaces

This chapter tells how to configure Virtual eXtensible LAN (VXLAN) interfaces. VXLANs act as Layer 2 virtual networks over Layer 3 physical networks to stretch Layer 2 networks.

About VXLAN Interfaces

VXLAN provides the same Ethernet Layer 2 network services as VLAN does, but with greater extensibility and flexibility. Compared to VLAN, VXLAN offers the following benefits:

- Flexible placement of multitenant segments throughout the data center.
- Higher scalability to address more Layer 2 segments: up to 16 million VXLAN segments.

This section describes how VXLAN works. For detailed information about VXLAN, see RFC 7348.

VXLAN Encapsulation

VXLAN is a Layer 2 overlay scheme on a Layer 3 network. It uses MAC Address-in-User Datagram Protocol (MAC-in-UDP) encapsulation. The original Layer 2 frame has a VXLAN header added and is then placed in a UDP-IP packet.

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces to which you apply your security policy, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

The following figure shows two ASAs and Virtual Server 2 acting as VTEPs across a Layer 3 network, extending the VNI 1, 2, and 3 networks between sites. The ASAs act as bridges or gateways between VXLAN and non-VXLAN networks.

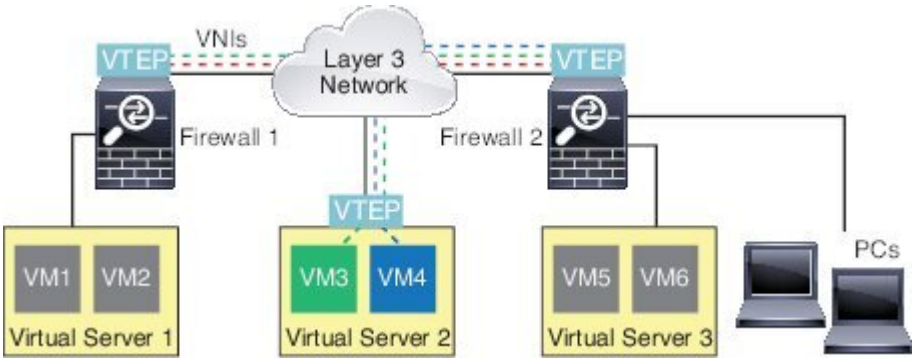

The underlying IP network between VTEPs is independent of the VXLAN overlay. Encapsulated packets are routed based on the outer IP address header, which has the initiating VTEP as the source IP address and the terminating VTEP as the destination IP address. The destination IP address can be a multicast group when the remote VTEP is not known. The destination port is UDP port 4789 by default (user configurable).

VTEP Source Interface

The VTEP source interface is a regular ASA interface (physical, redundant, EtherChannel, or even VLAN) with which you plan to associate all VNI interfaces. You can configure one VTEP source interface per ASA/security context.

The VTEP source interface can be devoted wholly to VXLAN traffic, although it is not restricted to that use. If desired, you can use the interface for regular traffic and apply a security policy to the interface for that traffic. For VXLAN traffic, however, all security policy must be applied to the VNI interfaces. The VTEP interface serves as a physical port only.

In transparent firewall mode, the VTEP source interface is not part of a BVI, and you do configure an IP address for it, similar to the way the management interface is treated.

VNI Interfaces

VNI interfaces are similar to VLAN interfaces: they are virtual interfaces that keep network traffic separated on a given physical interface by using tagging. You apply your security policy directly to each VNI interface.

You can only add one VTEP interface, and all VNI interfaces are associated with the same VTEP interface.

VXLAN Packet Processing

Traffic entering and exiting the VTEP source interface is subject to VXLAN processing, specifically encapsulation or decapsulation.

Encapsulation processing includes the following tasks:

- The VTEP source interface encapsulates the inner MAC frame with the VXLAN header.
- The UDP checksum field is set to zero.

- The Outer frame source IP is set to the VTEP interface IP.
- The Outer frame destination IP is decided by a remote VTEP IP lookup.

Decapsulation; the ASA only decapsulates a VXLAN packet if:

- It is a UDP packet with the destination port set to 4789 (this value is user configurable).
- The ingress interface is the VTEP source interface.
- The ingress interface IP address is the same as the destination IP address.
- The VXLAN packet format is compliant with the standard.

Peer VTEP

When the ASA sends a packet to a device behind a peer VTEP, the ASA needs two important pieces of information:

- The destination MAC address of the remote device
- The destination IP address of the peer VTEP

The ASA maintains a mapping of destination MAC addresses to remote VTEP IP addresses for the VNI interfaces.

There are two ways in which the ASA can find this information:

- A single peer VTEP IP address can be statically configured on the ASA.

  You cannot manually define multiple peers.

  The ASA then sends a VXLAN-encapsulated ARP broadcast to the VTEP to learn the end node MAC address.

- A multicast group can be configured on each VNI interface (or on the VTEP as a whole).

  The ASA sends a VXLAN-encapsulated ARP broadcast packet within an IP multicast packet through the VTEP source interface. The response to this ARP request enables the ASA to learn both the remote VTEP IP address along with the destination MAC address of the remote end node.

VXLAN Use Cases

This section describes the use cases for implementing VXLAN on the ASA.

VXLAN Bridge or Gateway Overview

Each ASA VTEP acts as a bridge or gateway between end nodes such as VMs, servers, and PCs and the VXLAN overlay network. For incoming frames received with VXLAN encapsulation over the VTEP source interface, the ASA strips out the VXLAN header and forwards it to a physical interface connected to a non-VXLAN network based on the destination MAC address of the inner Ethernet frame.

The ASA always processes VXLAN packets; it does not just forward VXLAN packets untouched between two other VTEPs.

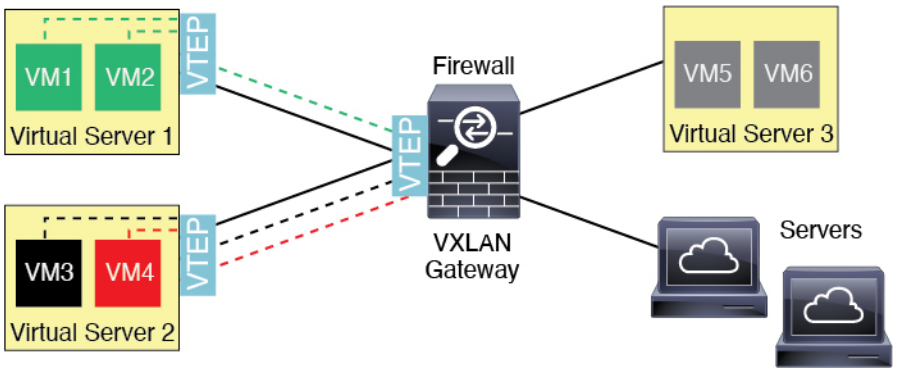

VXLAN Bridge

When you use a bridge group (transparent firewall mode, or optionally routed mode), the ASA can serve as a VXLAN bridge between a (remote) VXLAN segment and a local segment where both are in the same network. In this case, one member of the bridge group is a regular interface while the other member is a VNI interface.

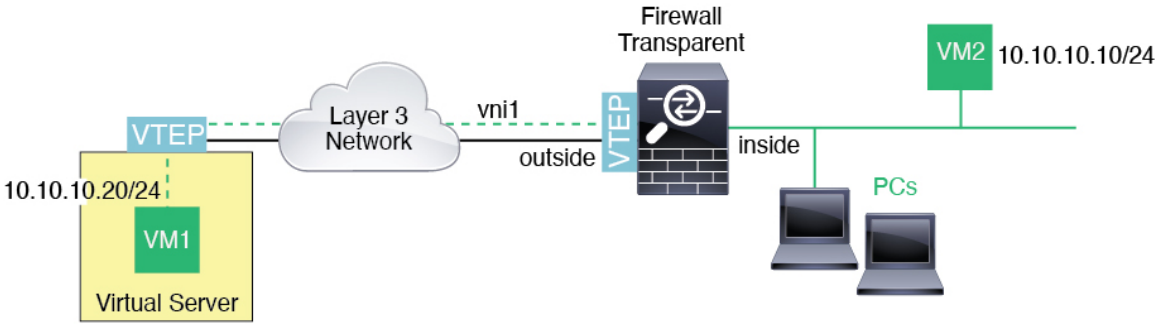

VXLAN Gateway (Routed Mode)

The ASA can serve as a router between VXLAN and non-VXLAN domains, connecting devices on different networks.

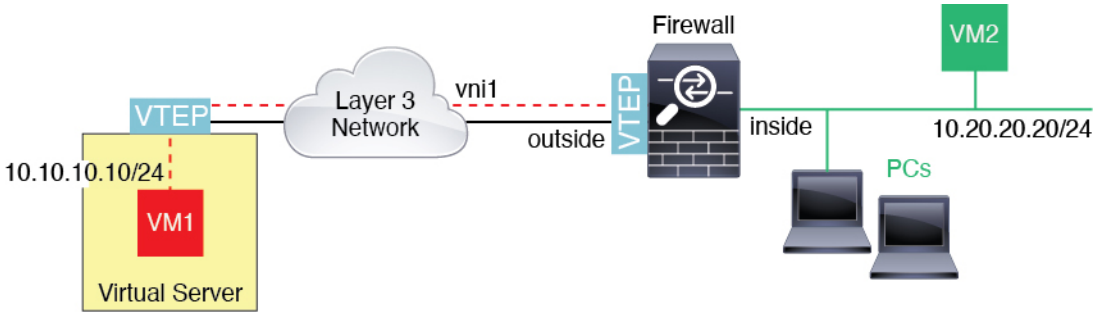

Router Between VXLAN Domains

With a VXLAN-stretched Layer 2 domain, a VM can point to an ASA as its gateway while the ASA is not on the same rack, or even when the ASA is far away over the Layer 3 network.

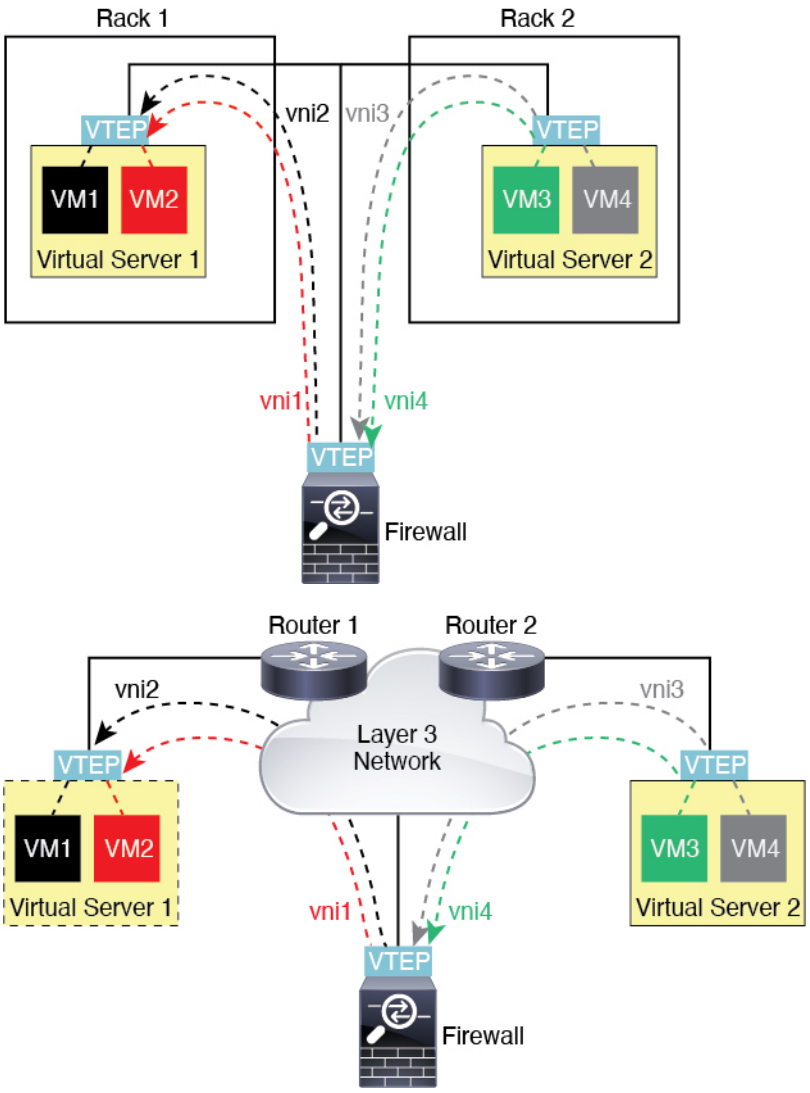

See the following notes about this scenario:

1. For packets from VM3 to VM1, the destination MAC address is the ASA MAC address, because the ASA is the default gateway.
2. The VTEP source interface on Virtual Server 2 receives packets from VM3, then encapsulates the packets with VNI 3's VXLAN tag and sends them to the ASA.
3. When the ASA receives the packets, it decapsulates the packets to get the inner frames.
4. The ASA uses the inner frames for route lookup, then finds that the destination is on VNI 2. If it does not already have a mapping for VM1, the ASA sends an encapsulated ARP broadcast on the multicast group IP on VNI 2.

Note The ASA must use dynamic VTEP peer discovery because it has multiple VTEP peers in this scenario.

5. The ASA encapsulates the packets again with the VXLAN tag for VNI 2 and sends the packets to Virtual Server 1. Before encapsulation, the ASA changes the inner frame destination MAC address to be the MAC of VM1 (multicast-encapsulated ARP might be needed for the ASA to learn the VM1 MAC address).
6. When Virtual Server 1 receives the VXLAN packets, it decapsulates the packets and delivers the inner frames to VM1.

Requirements and Prerequisites for VXLAN Interfaces

Model Requirements

- Firepower 1010 switch ports and VLAN interfaces are not supported as VTEP interfaces.

Guidelines for VXLAN Interfaces

IPv6

- The VNI interface supports both IPv4 and IPv6 traffic.
- The VTEP source interface IP address only supports IPv4.

Clustering

- ASA clustering does not support VXLAN in Individual Interface mode. Only Spanned EtherChannel mode supports VXLAN.

Routing

- Only static routing or Policy Based Routing is supported on the VNI interface; dynamic routing protocols are not supported.

MTU

- If the source interface MTU is less than 1554 bytes, then the ASA automatically raises the MTU. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. If the MTU used by other devices is larger, then you should set the source interface MTU to be the network MTU + 54 bytes. This MTU requires you to enable jumbo frame reservation on some models; see Enable Jumbo Frame Support (ASA Models, ASAv, ISA 3000).

Default Settings for VXLAN Interfaces

VNI interfaces are enabled by default.

Configure VXLAN Interfaces

To configure VXLAN, perform the following steps.

Procedure

Step 1 Configure the VTEP Source Interface, on page 7.

Step 2 Configure the VNI Interface, on page 8

Configure the VTEP Source Interface

You can configure one VTEP source interface per ASA or per security context. The VTEP is defined as a Network Virtualization Endpoint (NVE); VXLAN VTEP is the only supported NVE at this time.

Before you begin

For multiple context mode, complete the tasks in this section in the context execution space. In the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and edit the interface you want to use for the VTEP source interface.

Step 2 (Transparent Mode) Check the **VTEP Source Interface** check box.

This setting lets you configure an IP address for the interface. This command is optional for routed mode where this setting restricts traffic to VXLAN only on this interface.

Step 3 Configure the source interface name and IPv4 address, and click **OK**.

Step 4 Choose **Configuration** > **Device Setup** > **Interface Settings** > **VXLAN**.

Step 5 (Optional) Enter a **VXLAN Destination Port** value if you want to change from the default 4789.

In multiple context mode, configure this setting in the System execution space.

Step 6 Check the **Enable Network Virtualization Endpoint encapsulation using VXLAN** check box.

Step 7 Choose the **VTEP Tunnel Interface** from the drop-down list.

Note If the VTEP interface MTU is less than 1554 bytes, then the ASA automatically raises the MTU to 1554 bytes.

Step 8 (Optional) Check the **Configure Packet Recipient** check box.

- (Multiple context mode; Optional for single mode) Enter the **Specify Peer VTEP IP Address** to manually specify the peer VTEP IP address

If you specify the peer IP address, you cannot use multicast group discovery. Multicast is not supported in multiple context mode, so manual configuration is the only option. You can only specify one peer for the VTEP.

- (Single mode only) Enter the **Multicast traffic to default multicast address** to specify a default multicast group for all associated VNI interfaces.

  If you do not configure the multicast group per VNI interface, then this group is used. If you configure a group at the VNI interface level, then that group overrides this setting.

Step 9 Click **Apply**.

Configure the VNI Interface

Add a VNI interface, associate it with the VTEP source interface, and configure basic interface parameters.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and click **Add** > **VNI Interface**.

Step 2 Enter the **VNI ID**, between 1 and 10000.

This ID is just an internal interface identifier.

Step 3 Enter the **VNI Segment ID**, between 1 and 16777215.

The segment ID is used for VXLAN tagging.

Step 4 (Transparent Mode) Choose the **Bridge Group** to which you want to assign this interface.

See Configure Bridge Group Interfaces to configure the BVI interface and associate regular interfaces to this bridge group.

Step 5 Enter the **Interface Name**.

The name is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value.

Step 6 Enter the **Security Level**, between 0 (lowest) and 100 (highest). See Security Levels.

Step 7 (Single Mode) Enter the **Multicast Group IP Address**.

If you do not set the multicast group for the VNI interface, the default group from the VTEP source interface configuration is used, if available. If you manually set a VTEP peer IP for the VTEP source interface, you cannot specify a multicast group for the VNI interface. Multicast is not supported in multiple context mode.

Step 8 Check the **Map to VTEP Tunnel Interface** check box.

This setting associates the VNI interface with the VTEP source interface.

Step 9 Check the **Enable Interface** check box. This setting is enabled by default.

Step 10 (Routed Mode) In the **IP Address** area, configure an IPv4 address. To configure IPv6, click the **IPv6** tab.

Step 11 Click **OK**, and then **Apply**.

Allow Gateway Load Balancer Health Checks

The AWS Gateway Load Balancer requires appliances to answer a health check properly. The AWS Gateway Load Balancer will only send traffic to appliances that are considered healthy.

You must configure the ASAv to respond to an SSH, Telnet, HTTP, or HTTPS health check.

SSH Connection

For SSH, allow SSH from the Gateway Load Balancer. The Gateway Load Balancer will attempt to establish a connection to the ASAv, and the ASAv's prompt to log in is taken as proof of health.

Note An SSH login attempt will time out after 1 minute. You will need to configure a longer health check interval on the Gateway Load Balancer to accommodate this timeout.

Telnet Connection

For Telnet, allow Telnet from the Gateway Load Balancer. The Gateway Load Balancer will attempt to establish a connection to the ASAv, and the ASAv's prompt to log in is taken as proof of health.

Note You cannot Telnet to the lowest security level interface, so this method may not be practical.

HTTP(S) Cut-Through Proxy

You can configure the ASA to prompt the Gateway Load Balancer for an HTTP(S) login.

HTTP(S) Redirection Using Static Interface NAT with Port Translation

You can configure the ASAv to redirect health checks to a metadata HTTP(S) server. For HTTP(S) health checks, the HTTP(S) server must reply to the Gateway Load Balancer with a status code in the range 200 to 399. Because the ASAv has limits on the the number of simultaneous management connections, you may choose to offload the health check to an external server.

Static interface NAT with port translation lets you redirect a connection to a port (such as port 80) to a different IP address. For example, translate an HTTP packet from the Gateway Load Balancer with a destination of the ASAv outside interface so that it appears to be from the ASAv outside interface with a destination of the HTTP server. The ASAv then forwards the packet to the mapped destination address. The HTTP server responds to the ASAv outside interface, and then the ASAv forwards the response back to the Gateway Load Balancer. You need an access rule that allows traffic from the Gateway Load Balancer to the HTTP server.

Examples for VXLAN Interfaces

See the following configuration examples for VXLAN.

Transparent VXLAN Gateway Example

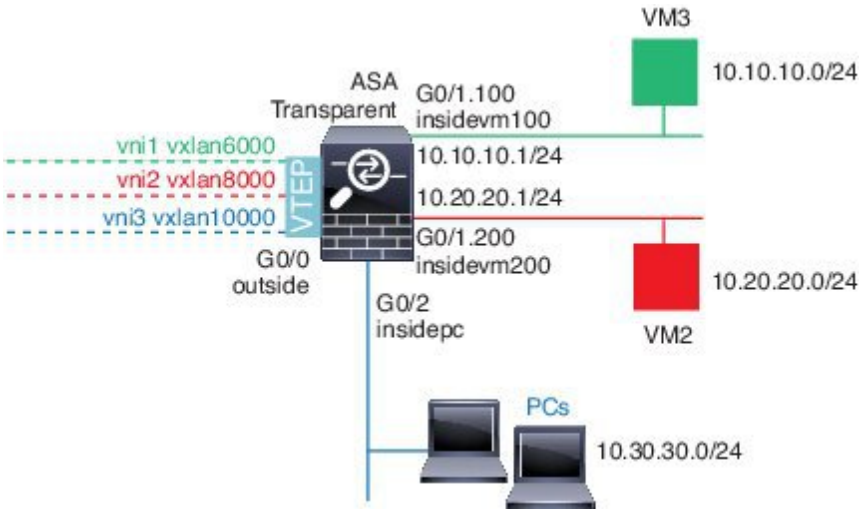

See the following description of this example:

- The outside interface on GigabitEthernet 0/0 is used as the VTEP source interface, and it is connected to the Layer 3 network.
- The insidevm100 VLAN subinterface on GigabitEthernet 0/1.100 is connected to the 10.10.10.0/24 network, on which VM3 resides. When VM3 communicates with VM1 (not shown; both have 10.10.10.0/24 IP addresses), the ASA uses VXLAN tag 6000.
- The insidevm200 VLAN subinterface on GigabitEthernet 0/1.200 is connected to the 10.20.20.0/24 network, on which VM2 resides. When VM2 communicates with VM4 (not shown; both have 10.20.20.0/24 IP addresses), the ASA uses VXLAN tag 8000.
- The insidepc interface on GigabitEthernet 0/2 is connected to the 10.30.30.0/24 network on which a few PCs reside. When those PCs communicate with VMs/PCs (not shown) behind a remote VTEP that belongs to same network (all have 10.30.30.0/24 IP addresses), the ASA uses VXLAN tag 10000.

ASA Configuration

```
firewall transparent
vxlan port 8427
!
interface gigabitethernet0/0
  nve-only
  nameif outside
  ip address 192.168.1.30 255.255.255.0
  no shutdown
!
nve 1
  encapsulation vxlan
  source-interface outside
!
interface vni1
  segment-id 6000
  nameif vxlan6000
  security-level 0
  bridge-group 1
```

```
  vtep-nve 1
  mcast-group 235.0.0.100
!
interface vni2
  segment-id 8000
  nameif vxlan8000
  security-level 0
  bridge-group 2
  vtep-nve 1
  mcast-group 236.0.0.100
!
interface vni3
  segment-id 10000
  nameif vxlan10000
  security-level 0
  bridge-group 3
  vtep-nve 1
  mcast-group 236.0.0.100
!
interface gigabitethernet0/1.100
  nameif insidevm100
  security-level 100
  bridge-group 1
!
interface gigabitethernet0/1.200
  nameif insidevm200
  security-level 100
  bridge-group 2
!
interface gigabitethernet0/2
  nameif insidepc
  security-level 100
  bridge-group 3
!
interface bvi 1
  ip address 10.10.10.1 255.255.255.0
!
interface bvi 2
  ip address 10.20.20.1 255.255.255.0
!
interface bvi 3
  ip address 10.30.30.1 255.255.255.0
```

Notes

- For VNI interfaces vni1 and vni2, the inner VLAN tag is removed during encapsulation.
- VNI interfaces vni2 and vni3 share the same multicast IP address for encapsulated ARP over multicast. This sharing is allowed.
- The ASA bridges the VXLAN traffic to non-VXLAN-supported interfaces based on the above BVIs and bridge group configurations. For each of the stretched Layer 2 network segments (10.10.10.0/24, 10.20.20.0/24 and 10.30.30.0/24), the ASA serves as a bridge.
- It is allowed to have more than one VNI or more than one regular interface (VLAN or just physical interface) in a bridge group. The forwarding or association between VXLAN segment ID to the VLAN ID (or a physical interface) is decided by the destination MAC address and which interface connects to the destination.

- The VTEP source-interface is a Layer 3 interface in transparent firewall mode indicated by **nve-only** in the interface configuration. The VTEP source interface is not a BVI interface or a management interface, but it has an IP address and uses the routing table.

VXLAN Routing Example

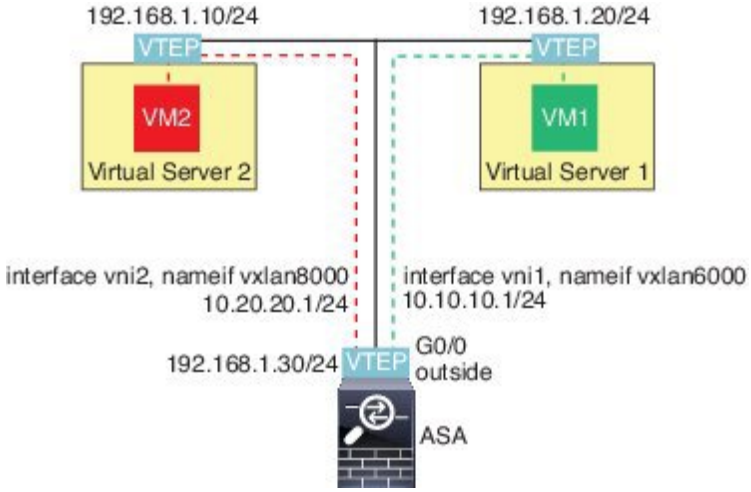

See the following description of this example:

- VM1 (10.10.10.10) is hosted on Virtual Server 1, and VM2 (10.20.20.20) is hosted on Virtual Server 2.
- The default gateway for VM1 is the ASA, which is not in the same pod as Virtual Server 1, but VM1 is not aware of it. VM1 only knows that its default gateway IP address is 10.10.10.1. Similarly, VM2 only knows that its default gateway IP address is 10.20.20.1.
- The VTEP-supported hypervisors on Virtual Server 1 and 2 are able to communicate with the ASA over the same subnet or through a Layer 3 network (not shown; in which case, the ASA and uplinks of virtual servers have different network addresses).
- VM1's packet will be encapsulated by its hypervisor's VTEP and sent to its default gateway over VXLAN tunneling.
- When VM1 sends a packet to VM2, the packet will be sent through default gateway 10.10.10.1 from its perspective. Virtual Server1 knows 10.10.10.1 is not local, so the VTEP encapsulates the packet over VXLAN and sends it to ASA's VTEP.
- On the ASA, the packet is decapsulated. The VXLAN segment ID is learned during decapsulation. The ASA then re-injects the inner frame to the corresponding VNI interface (vni1) based on the VXLAN segment ID. The ASA then conducts a route lookup and sends the inner packet through another VNI interface, vni2. All egressing packets through vni2 are encapsulated with VXLAN segment 8000 and sent through the VTEP to outside.
- Eventually the encapsulated packet is received by the VTEP of Virtual Server 2, which decapsulates it and forwards it to VM2.

ASA Configuration

```
interface gigabitethernet0/0
  nameif outside
  ip address 192.168.1.30 255.255.255.0
  no shutdown
!
nve 1
  encapsulation vxlan
  source-interface outside
  default-mcast-group 235.0.0.100
!
interface vni1
  segment-id 6000
  nameif vxlan6000
  security-level 0
  vtep-nve 1
  ip address 10.20.20.1 255.255.255.0
!
interface vni2
  segment-id 8000
  nameif vxlan8000
  security-level 0
  vtep-nve 1
  ip address 10.10.10.1 255.255.255.0
!
```

History for VXLAN Interfaces

Table 1: History for VXLAN Interfaces

Feature Name	Releases	Feature Information
VXLAN support	9.4(1)	VXLAN support was added, including VXLAN tunnel endpoint (VTEP) support. You can define one VTEP source interface per ASA or security context. We introduced the following screens: **Configuration > Device Setup > Interface Settings > Interfaces > Add > VNI Interface** **Configuration > Device Setup > Interface Settings > VXLAN**