Multi-tenancy

• Multi-tenancy , on page 1

Multi-tenancy

Multi-tenancy Overview

Multi-tenancy is a mode of operation where multiple independent instances (Layer-3 VRFs, Layer-2 VLANs) of a tenant (business entity, user group, applications, or security) operate in a shared environment (VXLAN BGP EVPN fabric), while ensuring logical segmentation between the instances. The tenant instances such as VRF and VLANs are logically isolated but physically operate on the same fabric.

Layer-3 and Layer-2 VNIs

In a VXLAN BGP EVPN fabric, a Layer-3 Virtual Network Identifier (VNI) identifies each tenant at the Layer-3 level and is associated with a unique tenant VRF.

As a pendant to Layer-3, Layer-2 virtual networks (VLANs) can carry a unique Layer-2 Virtual Network Identifier (VNI) in the fabric. Separate Layer-2 and Layer-3 networks can be created to achieve Layer-2 and Layer-3 segmentation, like for business units, user groups, applications, etc. Typically, a Layer-2 virtual network is associated with a single IP subnet while a VRF can contain multiple Layer-2 networks.

Note **Important**—From a global, VXLAN BGP EVPN fabric perspective, the VNI is the important identifier that is used across the fabric.

Servers belonging to a Layer-2 virtual network can be spread across the fabric, and might be associated with different Top of Rack (ToR)/leaf switches. Communication between servers or end hosts of the same Layer-2 virtual network is typically bridged.

Communication between end hosts belonging to different Layer-2 virtual networks represents Layer-3 communication, and is achieved through routing. Routed traffic traversing through the fabric logically traverses through the Layer-3 VNI or VRF VNI. The Layer-3 virtual network is similarly spread across different TOR/leaf switches to match the respective Layer-2 virtual networks that require routing.

Across the VXLAN BGP EVPN fabric, the L2 and L3 VNIs are unique and have a global significance. All the Layer-2 virtual networks of a tenant or VRF are associated with a common, unique, Layer-3 VRF VNI.

Note The L2 and L3 VNI use the same VNI field in the VXLAN encapsulation and hence can't overlap.

Figure 1: Logical representation of server traffic and segregation across the VXLAN BGP EVPN fabric

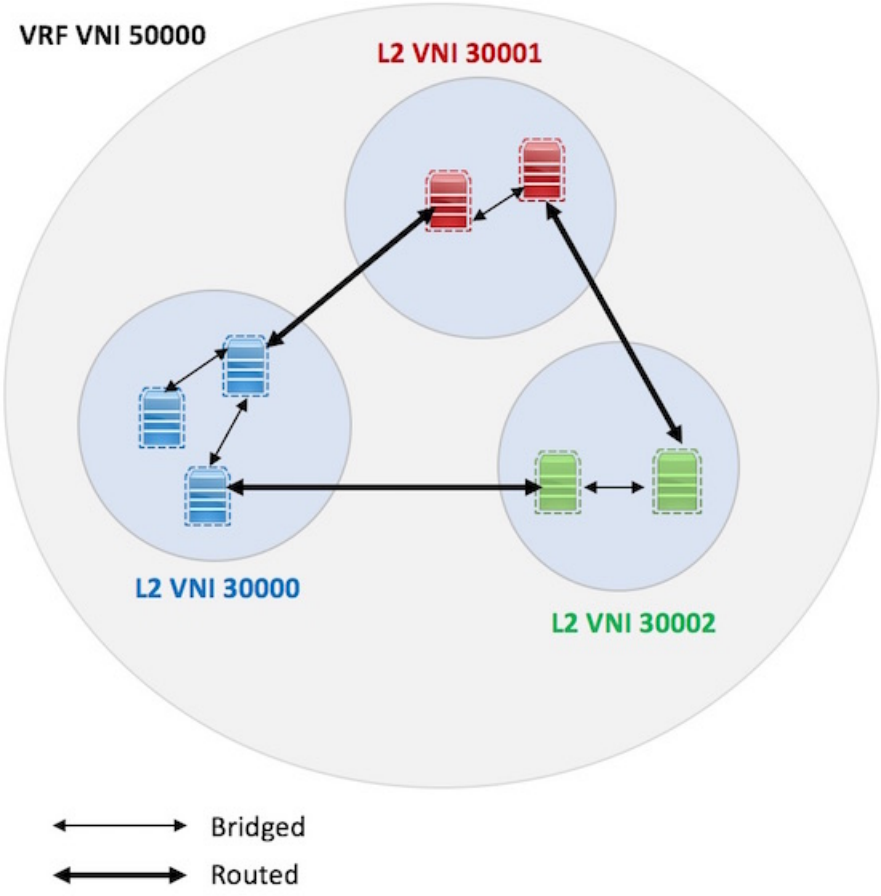

Figure 2: Physical representation of server traffic and segregation across the VXLAN BGP EVPN fabric

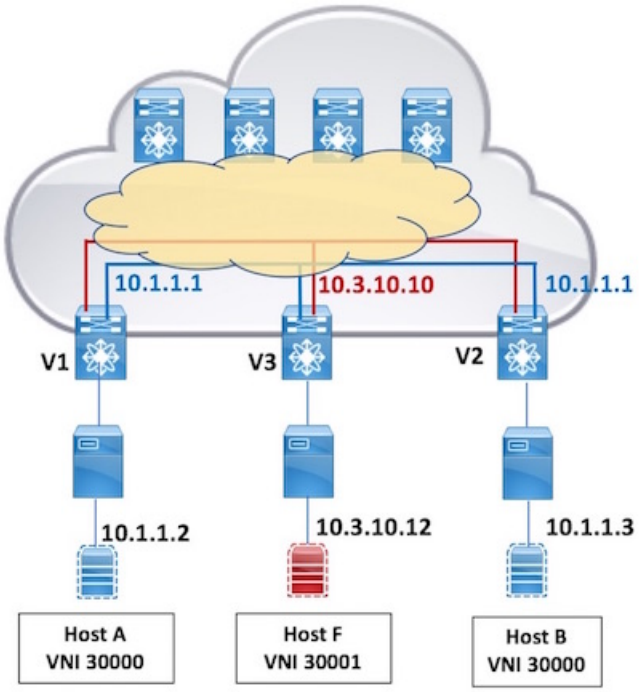

In the above sample topology, Host A and Host B belong to the same Layer-2 virtual network (VNI 30000, in blue color), so the traffic between them is bridged. Traffic from Host A to Host F (VNI 30001, in red color) is routed through the VRF VNI (say VNI 50000).

Routing between Layer-2 virtual networks

In the VXLAN BGP EVPN fabric, each Layer-2 virtual network needs to be configured with a first hop gateway approach. This first hop gateway will allow to traverse through the Layer-3 boundary, and send traffic to an end host in another Layer-2 virtual network. Since a Layer-2 virtual network might have presence across the fabric with its end hosts attached to multiple ToRs, the same first hop gateway IP address should be configured on those ToR switches where it has presence (Distributed Anycast Gateway).

In the above example, to route traffic from Host A (Layer-2 VNI 30000) to Host F (Layer-2 VNI 30001), you should configure a first hop gateway IP address (say 10.1.1.1/24) on the attached ToR switch V1. To route Host B traffic to Host F, you should enable the same first hop gateway (10.1.1.1/24) on the attached ToR switch V2. This is because Host A and B belong to the same Layer-2 network (VNI 30000). When using the Distributed Anycast Gateway as a first hop gateway, the IP address as well as the MAC address for the gateway itself will be the same across all ToR switches.

Layer-2 and Layer-3 Multi-tenancy

Let us consider bridging (Layer-2 multi-tenancy operation) and routing (Layer-3 multi-tenancy operation) between end hosts across the VXLAN BGP EVPN fabric, from a multi-tenancy perspective. The following sections explain how tenant instances (VLANs and VRFs) are connected across the fabric to send Layer-2 bridged and Layer-3 routed traffic from an end host to another.

For convenience, Cisco Nexus 9000 Series and 7000 Series switch concepts are explained separately.

Layer-2 Multi-tenancy on Cisco Nexus 9000 Series (and Cisco Nexus 5600 Series) Switches

VLAN based multi-tenancy can be implemented on the 9000 and 5600 Series switches. Some pointers are given below:

Figure 3: Layer-2 multi-tenancy—Same parent VLAN and VLAN on the wire

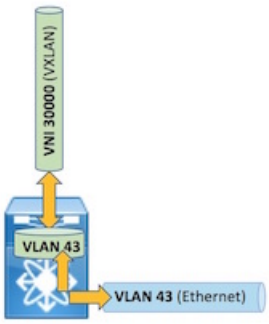

- From a ToR switch's perspective, a Layer-2 virtual network is represented by a VNI on the VXLAN BGP EVPN fabric side (VNI 30000 in the image) and a unique VLAN (43) on the tenant side. The 1:1 mapping between the parent VLAN and the VNI should be configured on the ToR switch.

Note A ToR switch can traditionally only accommodate the 12-bit VLAN namespace. However, the VLAN limitation at the network level is removed due to the introduction of VNIs or segments in the fabric.

- The use case below shows an end-to-end example.

Layer-2 multi-tenancy use case—Same parent VLAN and VLAN on the wire

Figure 4: Same parent VLAN and VLAN on the wire—VXLAN BGP EVPN fabric

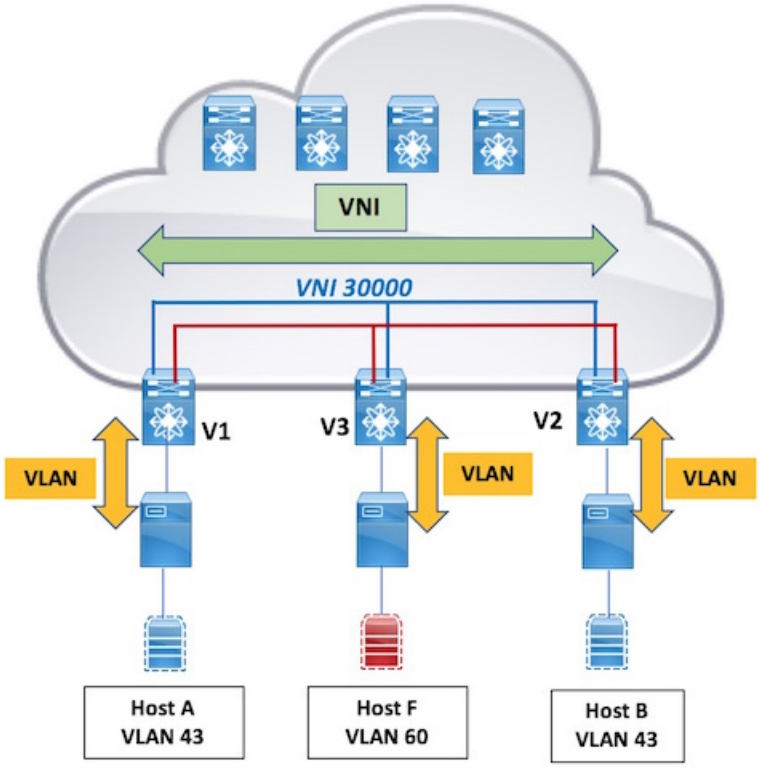

- In the above example, on ToR switch VTEP V1, VLAN 43 is mapped to L2 VNI 30000. All end host ports that have servers of this network should be associated with VLAN 43.
- Host B on V2 belongs to the same Layer-2 virtual network (30000). On ToR switch VTEP V2, Host B is mapped to VLAN 43 and VLAN 43 to VNI 30000.

Note Since VLANs only have ToR/leaf switch significance, different switches can have different VLAN IDs to represent a L2 network. However, it is convenient if you use the same VLAN ID across switches. The L2 VNI binds the Layer-2 virtual network and extends Layer-2 reachability across the fabric.

A sample configuration is given below:

Note **Important**—Configurations in this chapter/white paper are partial configurations in the overall VXLAN BGP EVPN fabric configuration, and should not be done in isolation. For complete configuration, refer the *Forwarding Configurations* chapter, Cisco Nexus 9000, Cisco Nexus 5600 Series switch configuration and Cisco Nexus 7000 Series switch configuration sections.

Figure 5: Same parent VLAN and VLAN on the wire – ToR switch VTEP

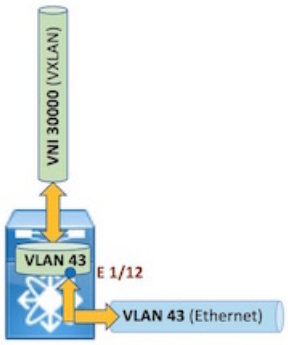

Create a VLAN and map it to a L2 VNI

(config) #

```
vlan 43
  vn-segment 30000
```

VLAN 43 and VNI 30000 are mapped with each other.

Configure an access port to enable Layer-2 traffic on the interface

(config) #

```
interface Ethernet 1/12
   switchport mode access
   switchport access vlan 43
```

Layer-2 multi-tenancy use case – VLANs on the wire are mapped to a parent VLAN

Note The manual configuration of VLAN translation is only supported on the Cisco Nexus 9000 Series switches and not on Cisco Nexus 5600 Series switches.

Figure 6: VLAN translation—Multiple VLANs mapped to a parent VLAN

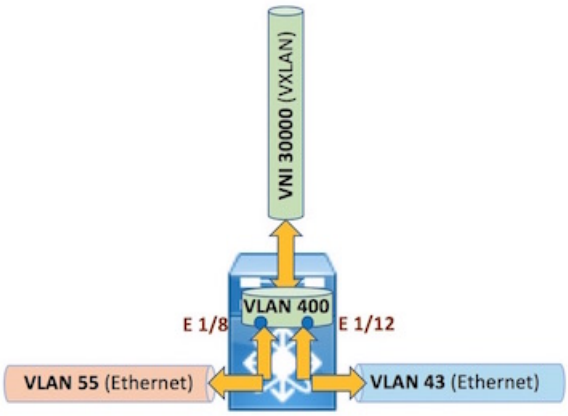

Parent VLAN

In this use case, VNI 30000 represents a Layer-2 virtual network on the fabric side. VLAN 400 is the parent VLAN on the ToR switch that represents the Layer-2 virtual network on the tenant side. VLAN 400 needs to be mapped to VNI 30000 for Layer-2 stitching or extension.

VLANs on the wire

VLANs 55 and 43 are on the wire, and end hosts are mapped to these VLANs. They also represent the Layer-2 virtual network 30000, but only through VLAN 400, the parent VLAN.

A sample configuration is given below:

Create a Parent VLAN and map to a L2 VNI

(config) #

```
vlan 400
  vn-segment 30000
```

Configure a trunk port to enable Layer-2 traffic across VLANs—Ethernet interface 1/8

(config) #

```
interface Ethernet 1/8
  switchport mode trunk
  switchport vlan mapping enable
  switchport vlan mapping  55 400
  switchport trunk allowed vlan 400
```

VLAN 55 is local and significant only to the port on which it is configured (Ethernet 1/8 in this case). VLAN 55 is mapped to parent VLAN 400.

Configure a trunk port to enable Layer-2 traffic across VLANs—Ethernet interface 1/12

(config) #

```
interface Ethernet 1/12
  switchport mode trunk
  switchport vlan mapping enable
  switchport vlan mapping 43 400
  switchport trunk allowed vlan 400
```

VLAN 43 is local and significant only to the port on which it is configured (Ethernet 1/12 in this case). VLAN 43 is also mapped to parent VLAN 400

Layer-3 multi-tenancy on Cisco Nexus 9000 Series (and Cisco Nexus 5600 Series) Switches

To route traffic between L2 virtual networks, an L3 virtual interface (called switch virtual interface [SVI]) should be created for each L2 virtual network, on a ToR/leaf switch. Each SVI should be associated with the tenant VRF, thereby adding it to the VRF table. A sample scenario is given below:

Figure 7: Layer-3 multi-tenancy

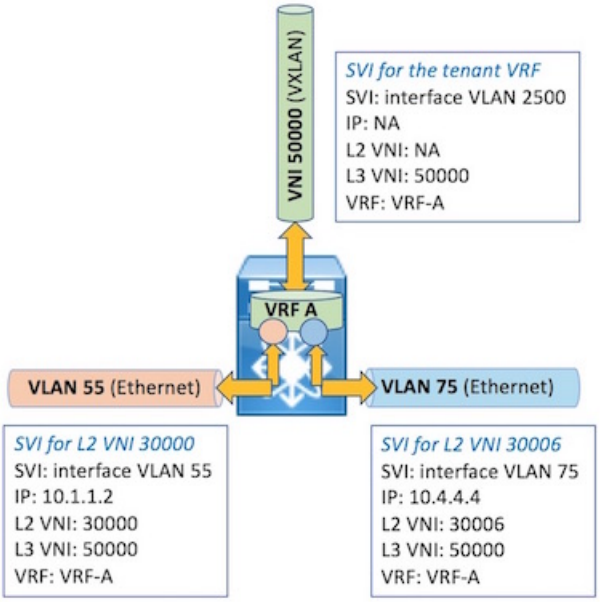

- Tenant VRF A is tied to Layer-3 VNI 50000. On this leaf switch/ToR VTEP, this tenant has the presence of L2 virtual networks 30000 and 30006.
- VLAN 2500 is created for L3 VNI 50000. An SVI, *interface VLAN2500*, is created for the tenant VRF.
- Similarly, SVIs are created for L2 VNI 30000 (VLAN 55), and L2 VNI 30006 (VLAN 75). The SVIs are *interface VLAN 55* and *interface VLAN 75*. An IP address is assigned to each SVI.

Layer-2 multi-tenancy on Cisco Nexus 7000 Series switches

Bridge domain based multi-tenancy can be implemented on the 7000 Series switches.

Figure 8: A depiction of the relationship between bridge domains, VLANs, and ethernet ports

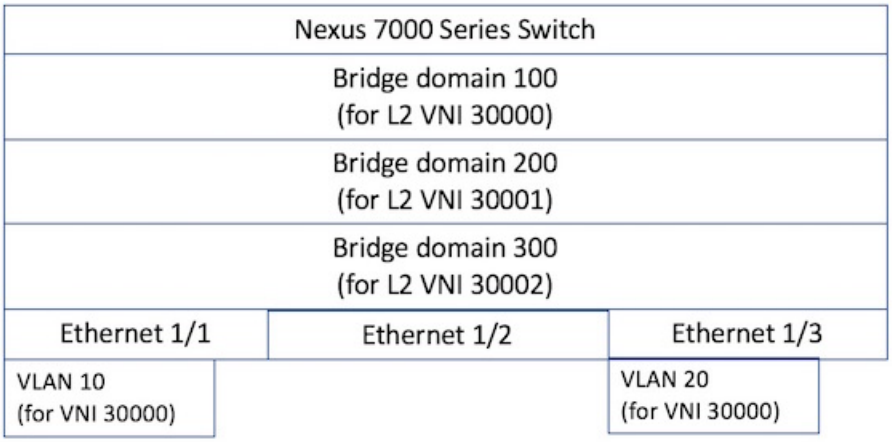

Some pointers about bridge domain based multi-tenancy are given below:

- *Bridge domain*—It is a Layer-2 flood broadcast domain that maps/connects VLANs of a Layer-2 virtual network on the tenant side of a ToR switch to a Layer-2 VNI on the fabric side of the switch. There is a 1:1 mapping between the bridge domain ID and the L2 VNI (one BD for each L2 virtual network). The most common known representation of a bridge domain is a VLAN.

  Note **Important**—Here, VLANs have port level significance. For example, VLAN 10 represents Layer-2 virtual network 30000 on Ethernet 1/1 and VLAN 20 represents it on Ethernet 1/3. So, BD 100 (BD for L2 VNI 30000) bridges L2 end host traffic between Eth 1/1 and Eth 1/3, and from either (VLAN) port to end hosts in the same L2 virtual network across the fabric.

  MAC learning is done on the bridge domain (e.g. BD 100) that maps to VNI 30000.
- *Flexibility*—You can reuse VLAN IDs across ports (see image *Port specific multi-tenancy* below – VLAN 20 ties to VNI 30001 on Eth 1/1 but to VNI 30000 on Eth 1/3), and reuse BD IDs across ToR switches. A ToR switch in this model can accommodate 4K VLANs per 802.1Q trunk port. The VLAN limitation at the network level is removed due to the introduction of VNIs in the fabric.

Figure 9: Port specific multi-tenancy

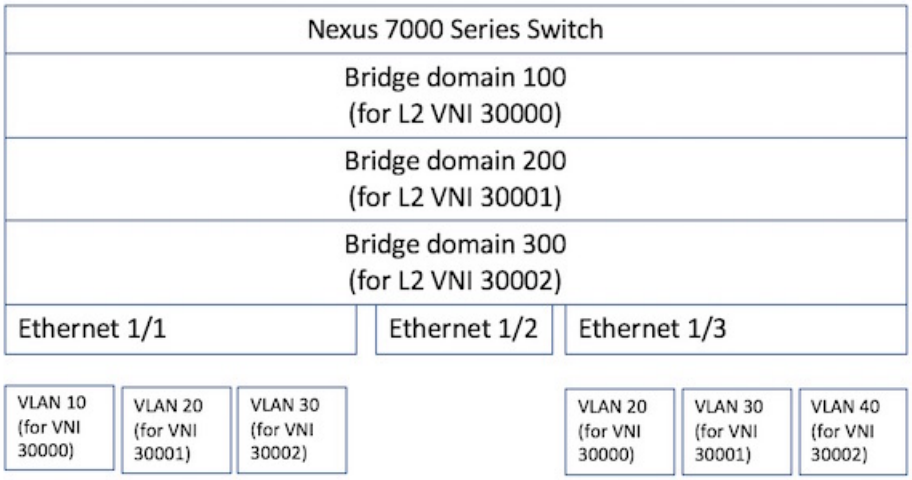

- The L2 VNI is the common binding factor for a L2 virtual network across the fabric.

Note A bridge-domain or VLAN has no significance from a global identifier perspective when using port-VLAN translation.

Figure 10: Layer-2 traffic flow on a ToR switch

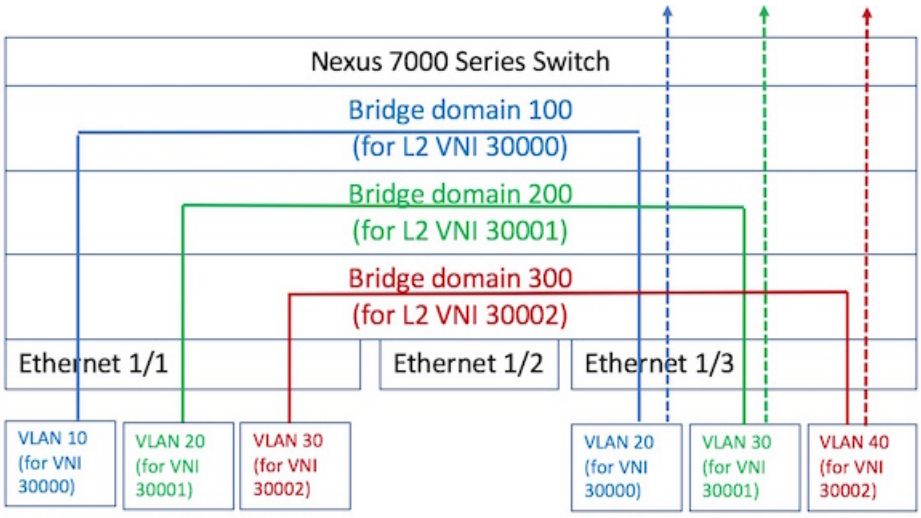

The colored lines represent bridged traffic in a Layer-2 virtual network (blue represents VNI 30000, green VNI 30001 and red VNI 30002), within the ToR switch. The dotted lines represent bridged traffic from the ToR switch to another switch.

Figure 11: Layer-2 multi-tenancy configuration

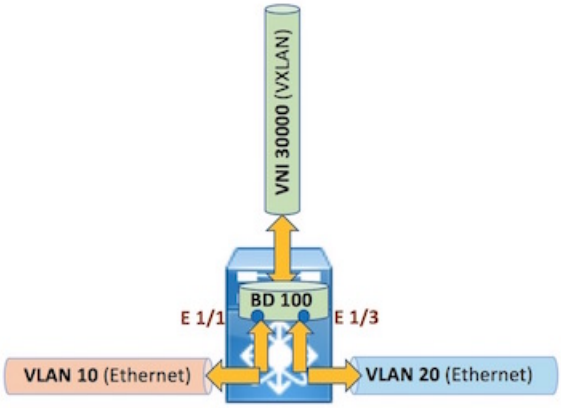

- In the above image, L2 VNI 30000 is represented by BD 100. VLAN 10 on port Ethernet 1/1 and VLAN 20 on port Ethernet 1/3 are tied to BD 100.

Note Though different ToR switches can have different bridge domain IDs for the same Layer-2 virtual network, it is convenient if you assign the same bridge domain ID.

A sample configuration for the above image is given below:

Note **Important**—Configurations in this chapter/white paper are partial configurations in the overall VXLAN BGP EVPN fabric configuration, and should not be done in isolation. For complete configuration, refer the *Forwarding Configurations* chapter, Cisco Nexus 5600 Series switch configuration and Cisco Nexus 7000 Series switch configuration sections.

Create the Layer-2 VNI on the ToR switch

(config) #

```
vni 30000
```

Create a bridge domain to represent the Layer-2 virtual network, and map it to the VNI

(config) #

```
system bridge-domain 100
bridge-domain 100
    member vni 30000
```

The **system bridge-domain** command identifies the bridge domain IDs and the **bridge-domain** command configures the specified bridge domain.

Create a profile that maps the 802.1Q VLANs to the Layer-2 VNI (through bridge domain 100)

(config) #

```
encapsulation profile vni vsi_10
   dot1q 10 vni 30000

 encapsulation profile vni vsi_20
   dot1q 20 vni 30000
```

On Nexus 7000 Series switches, a port is assigned as a trunk port by applying the above profile (mapping) on the specified port.

Assign the 802.1Q VLANs to a designated trunk port

(config) #

```
interface Ethernet 1/1
   no shutdown
   service instance 1 vni
       no shutdown
       encapsulation profile vsi_10 default

interface Ethernet 1/3
   no shutdown
   service instance 1 vni
       no shutdown
       encapsulation profile vsi_20 default
```

In VLAN mode, the **switchport mode trunk** command was used to designate a port as a trunk port.

Use case—Multiple Layer-2 Virtual Networks' configuration

Typically, multiple Layer-2 networks are represented on a ToR switch.

Figure 12: Multiple Layer-2 virtual networks on a ToR switch

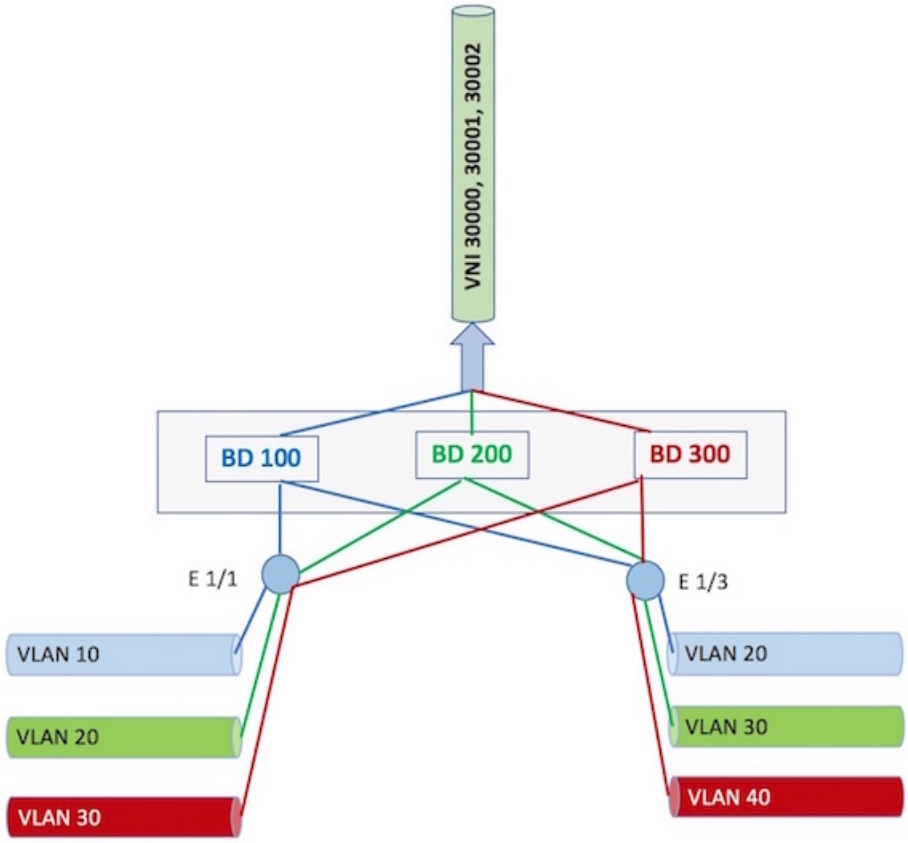

You can see the flexibility of mapping from the above image.

- VLAN 10 on Ethernet 1/1, and VLAN 20 on Ethernet 1/3 are in the same BD/VNI (BD 100/VNI 30000).
- VLAN 20 on Ethernet 1/1, and VLAN 30 on Ethernet 1/3 are in the same BD/VNI (BD 200/VNI 30001).
- VLAN 30 on Ethernet 1/1, and VLAN 40 on Ethernet 1/3 are in the same BD/VNI (BD 300/VNI 30002).

A sample configuration for multiple Layer-2 networks' configuration at the same time is given below:

Create Layer-2 VNIs on the ToR switch

(config) #

```
vni 30000, 30001, 30002
```

Create bridge domains to represent the Layer-2 virtual networks, and map it to the respective VNIs

(config) #

```
system bridge-domain 100, 200, 300
bridge-domain 100, 200, 300
    member vni 30000 - 30002
```

In the above configuration, BD 100 is subscribed to VNI 30000, BD 200 to VNI 30001, and BD 300 to VNI 30002.

Create profiles that map 802.1Q VLANs to respective Layer-2 VNIs

(config) #

```
encapsulation profile vni vsi_10-30
   dot1q 10, 20, 30 vni 30000, 30001, 30002

encapsulation profile vni vsi_20-40
   dot1q 20, 30, 40 vni 30000, 30001, 30002
```

On Nexus 7000 Series switches, a port is assigned as a trunk port by applying the above profiles (and mappings) on the specified ports.

Assign the 802.1Q VLANs to designated trunk ports

(config) #

```
interface Ethernet 1/1
   no shutdown
   service instance 1 vni
       no shutdown
       encapsulation profile vsi_10-30 default

interface Ethernet 1/3
   no shutdown
   service instance 1 vni
       no shutdown
       encapsulation profile vsi_20-40 default
```

Layer 3 Multi-Tenancy on Cisco Nexus 7000 Series Switches

Earlier, the image *L2 traffic flow on a ToR switch* depicted Layer-2 traffic flow on a ToR switch. In the image *Layer-3 traffic flow* (below), the dotted colored lines represent routed traffic between 2 different Layer-2 virtual networks.

Note A bridge domain (BD) representation for Layer-3 VNI 50000 is added in the image.

Figure 13: Layer-3 traffic flow

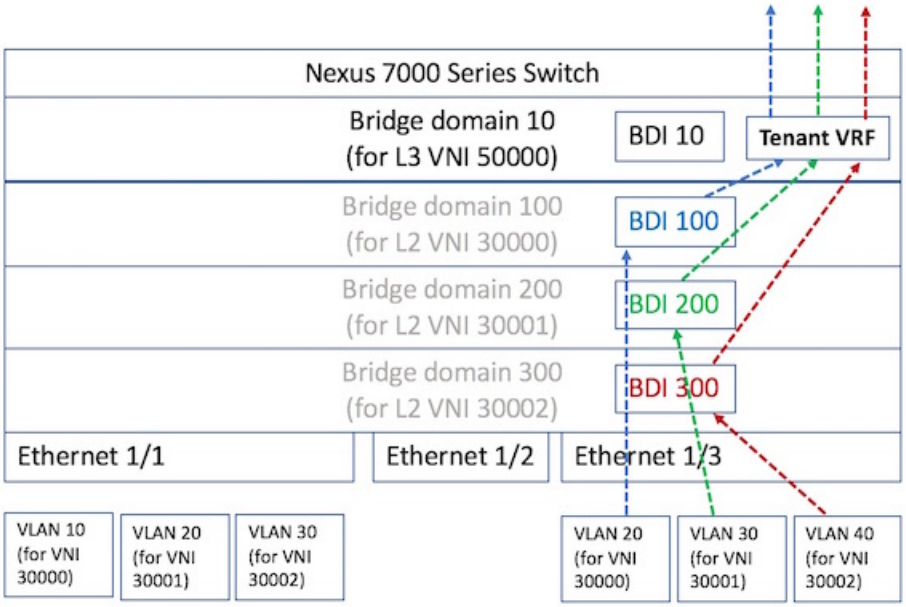

A virtual L3 interface, called a bridge domain interface (BDI), is used to route traffic from a BD. A BDI should be created for each BD (for each L2 VNI) and added to the tenant VRF.

Let us say the source end host in the Layer-2 virtual network 30000 is represented by VLAN 20, on port Eth 1/3. The source sends traffic to an end host on a different L2 virtual network, external to the switch. Since the BDI is added to the VRF table, traffic will be routed through the L3 network VNI that represents the tenant, to the target ToR switch,.

A sample set of configurations is given below:

Figure 14: Layer-3 multi-tenancy configuration

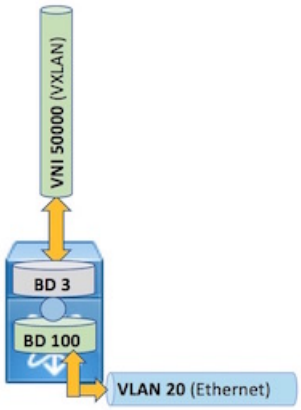

A sample configuration for the above image is given below:

Create the Layer-3 VNI for the tenant VRF

(config) #

```
vni 50000
```

Create a bridge domain to represent the Layer-3 VNI, and map it to the VNI

(config) #

```
system bridge-domain 3
bridge-domain 3
    member vni 50000
```

The **system bridge-domain** command identifies the bridge domain IDs and the **bridge-domain** command configures the specified bridge domain.

Configure a VRF overlay VLAN/BDI for the VRF

(config) #

```
interface BDI3
   no shutdown
   mtu 9216
   vrf member VRF-A
   ip forward
```

Create the tenant VRF and associate the Layer-3 VNI to it

(config) #

```
vrf context VRF-A
   vni 50000
   rd auto
   address-family ipv4 unicast
      route-target both auto evpn
```

Note In the Layer-2 multi-tenancy example, BD 100 was created to represent L2 VNI 30000 on the ToR/leaf switch, as shown below.

(config) #

```
bridge-domain 100
    member vni 30000
```

Create a server facing BDI for BD 100 and associate it with the tenant VRF

(config) #

```
interface BDI100
  no shutdown
  vrf member VRF-A
```

```
ip address 10.2.2.1/24
fabric forwarding mode anycast-gateway
```

The virtual interface BDI100 is the Layer-3 interface that represents BD 100. This configuration updates the BDI information to the VRF table.