



## VXLAN BGP EVPN and OTV Interoperation

A datacenter fabric can be Virtual eXtensible Local Area Network (VXLAN) Border Gateway Protocol (BGP) Ethernet VPN (EVPN) based, classical ethernet (CE) based, or combined. You can connect different datacenters over an IP WAN with Layer 2 (such as Overlay Transport Virtualization [OTV]) and Layer 3 technologies, and also connect end hosts between Layer 2 CE and VXLAN pods within a datacenter. With this feature, you can:

- Configure VXLAN and OTV (with Bridge Domain Interface [BDI]) as a one-box solution. The VXLAN and OTV overlays or tunnels are stitched together on a VXLAN BGP EVPN fabric border leaf switch, ensuring that the Layer 2 traffic between VXLAN and OTV is within the same bridge domain.
- Configure OTV with BDI in a datacenter as a one-box solution, instead of a two box solution used in a legacy datacenter.

OTV as the overlay enables Layer 2 connectivity between separate VXLAN BGP EVPN or CE Layer 2 domains while maintaining resiliency and load-balancing benefits of the IP based interconnection.

- [Prerequisites for VXLAN BGP EVPN and OTV Interoperation, on page 1](#)
- [Guidelines and Limitations for VXLAN BGP EVPN and OTV Interoperation , on page 2](#)
- [Information About VXLAN BGP EVPN and OTV Interoperation, on page 2](#)
- [How to Configure VXLAN BGP EVPN and OTV Interoperation, and OTV with BDI, on page 9](#)
- [Verifying VXLAN BGP EVPN and OTV Interoperation, and OTV with BDI, on page 13](#)
- [Feature History for VXLAN BGP EVPN and OTV Interoperation , on page 17](#)

## Prerequisites for VXLAN BGP EVPN and OTV Interoperation

- A Nexus 7000 or 7700 Series switch with an M3 line card.
- Conceptual and configuration knowledge about VXLAN BGP EVPN and OTV datacenter fabrics.
- For a functioning VXLAN BGP EVPN datacenter, configurations should be enabled on the leaf and spine switches. For more information see the "Configuring the VXLAN BGP EVPN" chapter, or [Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide](#).

# Guidelines and Limitations for VXLAN BGP EVPN and OTV Interoperation

- A unique, primary gateway IP address should be configured on each BDI for sending ARP requests over OTV. A secondary anycast IP address is configured on each BDI for sending ARP requests over EVPN.
- For Layer 3 multicast traffic, an external, centralized Layer 3 multicast gateway should be enabled. Layer 3 multicast routing support is not available for the VXLAN BGP EVPN and OTV Interoperation feature in release 8.2(1). However, Layer 3 multicast traffic within a VXLAN BGP EVPN fabric is supported as usual.
- For seamless mobility across legacy and VXLAN BGP EVPN fabrics, the anycast gateway MAC address used in the VXLAN fabric should be configured as the HSRP static MAC address in the legacy fabric.
- Only OTV unicast control network (OTV Adjacency Server function) is supported.
- On the same physical Join interface, regular OTV overlays cannot interoperate with OTV overlays that are stitched to VXLAN overlays. However, an OTV overlay and an OTV+VXLAN overlay can be enabled on separate physical Join interfaces.
- In an OTV with BDI single box solution, the ARP proxy function option is not supported in the Cisco NX-OS 8.2(1) release.
- Support for this feature is limited to 3 OTV datacenter sites, for the 8.2(1) release.
- For OTV overlays, only Generic Routing Encapsulation (GRE) encapsulation is supported for the 8.2(1) release.

## Information About VXLAN BGP EVPN and OTV Interoperation

Two VXLAN BGP EVPN datacenters can be connected with each other and to a legacy datacenter (Sample topology 1). A VXLAN BGP EVPN datacenter can be connected to a legacy datacenter and to a datacenter with a OTV+BDI one box solution implemented in the form of a vPC pair of border switches (Sample topology 2). In either topology, the Layer 3 aggregation switch in the legacy datacenter performs the role of an external, centralized multicast gateway for Layer 3 multicast traffic across the datacenters. The two sample topologies and workflow for the VXLAN BGP EVPN and OTV interoperation in them:

# Sample Topologies and Workflow of the VXLAN BGP EVPN and OTV Interoperation

Figure 1: Sample topology 1 - VXLAN BGP EVPN and OTV interoperation

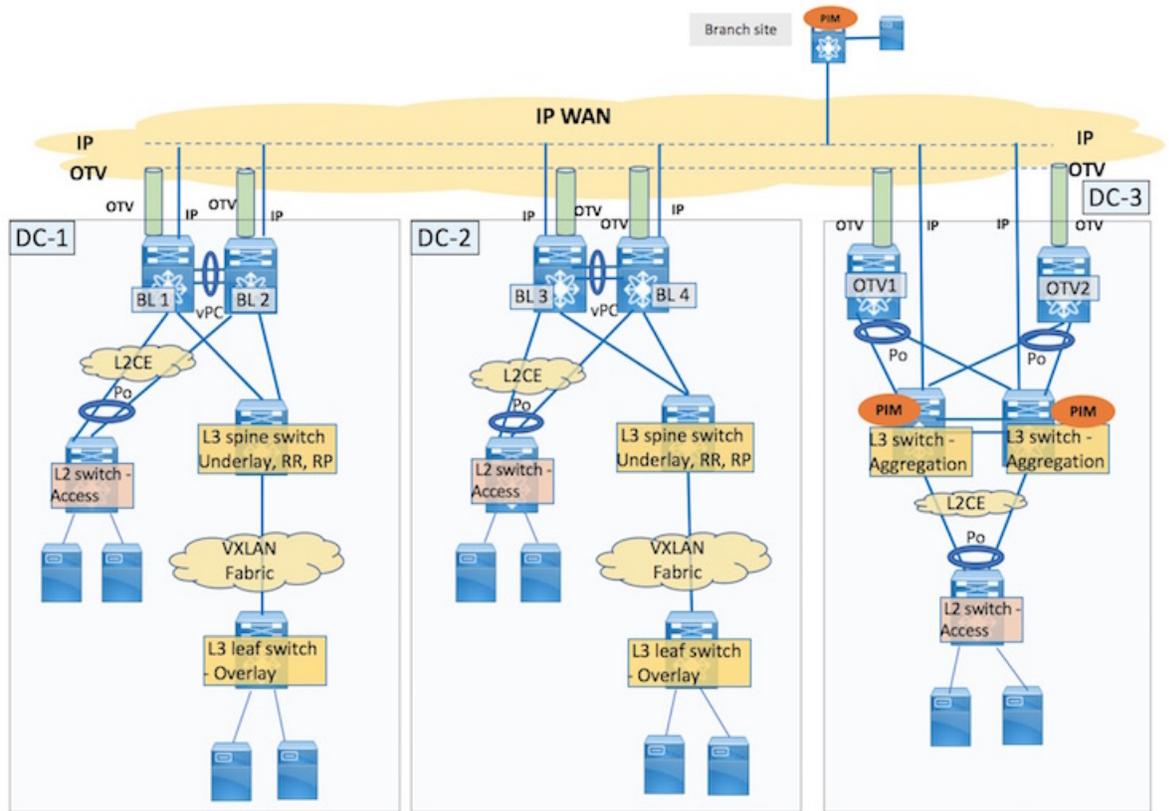
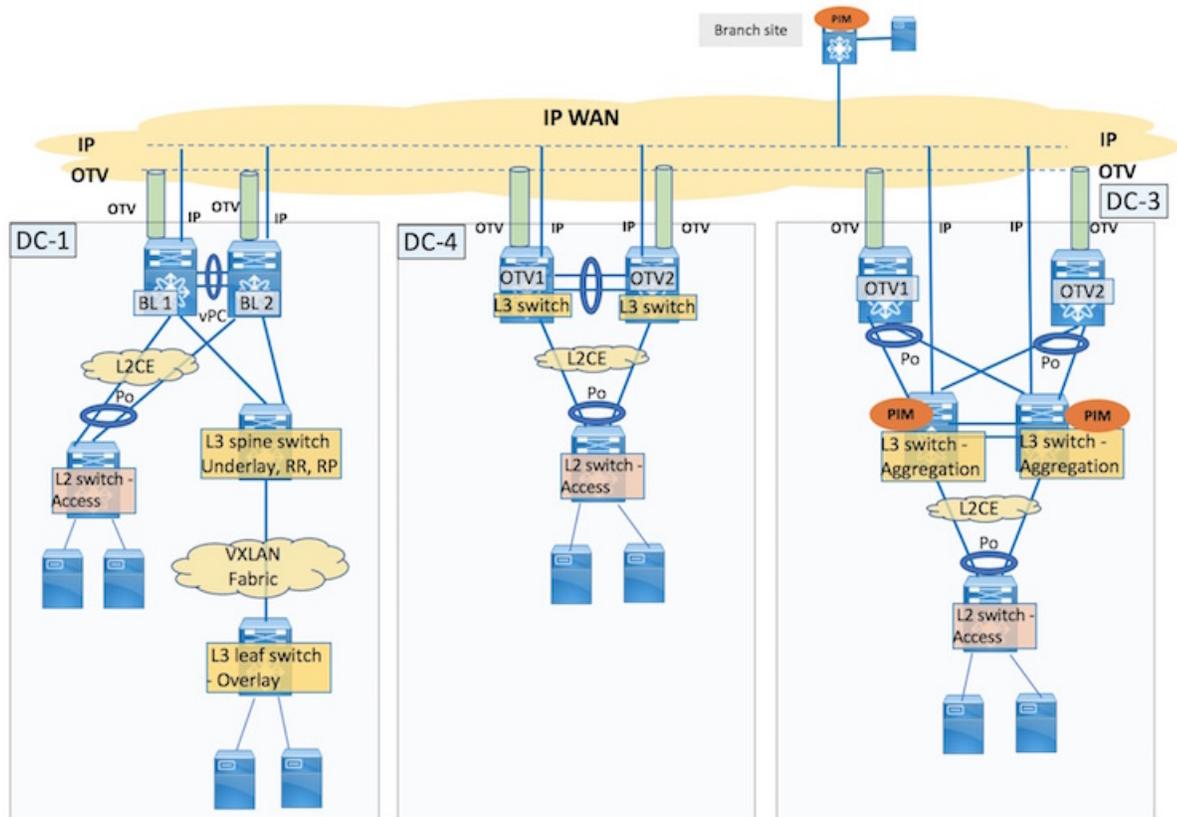


Figure 2: Sample topology 2 - VXLAN BGP EVPN and OTV interoperation



The focus of this feature is the VXLAN+OTV interoperation and OTV+BDI function on a Nexus 7000 or 7700 Series border leaf switch with an M3 line card, and the traffic flow between servers in these and a legacy datacenter. Topology information:



**Note** Support for this feature is limited to 3 OTV datacenter sites, for the 8.2(1) release.

- Topology 1 displays 3 datacenters, DC-1, DC-2 and DC-3. DC-1 and DC-2 have VXLAN BGP EVPN and CE pods, with VXLAN and OTV (with Bridge Domain Interface [BDI]) configurations on the Virtual Port Channel (vPC) pair border leaf switches (BL 1 and BL 2, BL 3 and BL 4). DC-3 is a legacy OTV site with OTV switches in separate VDCs (OTV1 and OTV2) at the border.
- Topology 2 displays 3 datacenters, DC-1, DC-4 and DC-3. DC-1 has VXLAN BGP EVPN and CE pods, with VXLAN and OTV (with BDI) enabled on the vPC border leaf switch pair. DC-4 is a OTV+BDI one box solution with configurations implemented on the vPC border switch pair. DC-3 is a legacy OTV site.

## Layer 2 Switching

Layer 2 traffic is transported between the datacenters through the border leaf switches (in DC-1 and DC-2) and OTV devices (in DC-3 and DC-4) at the site border, over the IP WAN.

In DC-1 and DC-2, there are 2 scenarios where Layer 2 traffic is transported between VXLAN and OTV overlays or tunnels. On the border leaf switch, traffic from the VXLAN fabric is either sent to a server within the datacenter, or towards another datacenter through OTV. The other way around, the border leaf switch receive traffic from OTV towards the VXLAN fabric or CE pod. When you configure this feature, Layer 2 traffic seamlessly passes between the VXLAN and OTV tunnels on the same device. Packet flow details in DC-1 and DC-2:

- Packet flow within the Layer 2 CE pod and packet flow within the VXLAN BGP EVPN pod (see "Configuring VXLAN BGP EVPN" chapter) remains the same. When a Layer 2 CE pod server sends traffic to a server in the VXLAN fabric within the site or to another site, the packets reach the border leaf switch. The bridge domain, Layer 2/Layer 3 VNI mappings, and MAC routes of the VXLAN fabric are available in the border leaf switch. If the destination server is within the fabric, the border leaf switch VXLAN encapsulates the packet and sends it to the corresponding ToR or leaf switch. The leaf switch VXLAN decapsulates the traffic and sends the original packet to the intended server. If the destination server is in another site, the border leaf switch OTV encapsulates the traffic towards the remote site.




---

**Note** Though the simplified sample topology depicts a single switch at the ToR/leaf layer carrying Layer 2 server traffic within the VXLAN BGP EVPN fabric, a real time VXLAN fabric spine-leaf setup will have multiple switches at the ToR/leaf and spine layers, and intra fabric Layer 2 server traffic flows through those non border leaf switches.

---

- When a VXLAN+OTV border leaf switch receives traffic from another site over OTV, it removes the OTV encapsulation, does a lookup to find out where the destination server resides, and VXLAN encapsulates the traffic towards the corresponding ToR/leaf switch. The leaf switch VXLAN decapsulates the traffic and sends the original packet to the intended server. If the destination server is in the Layer 2 CE pod, the border leaf switch OTV decapsulates the traffic and sends the traffic to the destination server without any encapsulation.
- When a server in the VXLAN BGP EVPN fabric sends traffic to a server in the Layer 2 CE pod, the border leaf switch receives the packets. It VXLAN decapsulates the traffic and sends it to the destination server in the CE pod.

Packet flow details in DC-4:

- When a Layer 2 CE pod server in the datacenter with the OTV+BDI one box solution sends traffic, the destination server is either within the datacenter or outside of it. Traffic flow within the datacenter remains the same. If the destination server is in another site, then the packets reach the Layer 3 OTV (with BDI) switch. The switch OTV encapsulates the traffic towards the legacy or VXLAN+OTV datacenter. The border switch in the destination datacenter receives the traffic, OTV decapsulates it, and forwards it as explained in the earlier sections.
- When the Layer 3 OTV (with BDI) switch receives traffic from another datacenter over OTV, it OTV decapsulates the traffic and sends it towards the corresponding Layer 2 access switch. The access switch forwards the packets to the destination server.

## Control Plane

- BGP EVPN is used for advertising MAC and MAC-IP routes across the VXLAN BGP EVPN fabric in DC-1 and DC-2.

- Intermediate-System to Intermediate-System (IS-IS) is used for advertising MAC routes across OTV configured device.
- MAC Route updates in the control plane are reflected across the OTV and VXLAN tunnels.

## Layer 3 Unicast Routing

### DC-1 and DC-2 (VXLAN BGP EVPN fabric)

Layer 3 routing within the VXLAN BGP EVPN fabric takes place through an underlay Interior Gateway Protocol (IGP) such as Intermediate System-to-Intermediate System (IS-IS) or OSPF.

Layer 3 traffic between the border leaf switches and the IP WAN should be through over Multiprotocol Label Switching (MPLS) L3VPN or virtual routing and forwarding (VRF) Lite. The IGP and external connectivity documentation is available in the *Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide*.

A distributed anycast gateway (or BDI) IP address is used for Layer 3 traffic between Layer 2 virtual networks in the VXLAN fabric. This should be configured as the secondary BDI IP address. A unique, primary gateway IP address should be configured on each BDI for sending ARP requests over OTV. Packet flow details in DC-1 and DC-2:



#### Note

- Though the main focus of the feature is one box solutions for VXLAN+OTV with BDI and OTV with BDI functions, traffic flow from/to the Layer 2 CE pod is also explained for completeness.
  - Though the simplified sample topology depicts a border leaf switch carrying Layer 3 server traffic within the VXLAN BGP EVPN fabric, real time intra VXLAN fabric server traffic flow is through a non border leaf switch.
- 
- When a server in the Layer 2 CE pod sends traffic to a server in another VLAN/subnet, either located in the VXLAN BGP EVPN fabric within the site or to a server in a remote site, the traffic first reaches the border leaf switch. The bridge domain, Layer 2/Layer 3 VNI mappings, MAC routes and the appropriate IGP configuration will be available in the border leaf switch. If the destination server is within the fabric, the border leaf switch VXLAN encapsulates the packet and routes it to the corresponding ToR or leaf switch through the fabric underlay routing protocol such as IS-IS or OSPF. The leaf switch VXLAN decapsulates the traffic and sends the original packet to the intended server. If the destination server is in another site, the border leaf switch sends the traffic towards the remote site through a Layer 3 DCI technology like MPLS L3VPN or OTV, depending on how the host route is learnt (through MPLS VPN or OTV).
  - When the Layer 3 DCI enabled border switch receives Layer 3 traffic from another site (through MPLS VPN or OTV), it does a lookup to find out where the destination server resides and routes the packets (across VLANs, through the corresponding destination BDI), to the corresponding Layer 2 switch. The switch forwards the packets to the intended destination server.

### DC-3, the legacy datacenter

- Layer 3 routing within DC-3 is through an IGP implemented on the aggregation switches. A Layer 3 Hot Standby Router Protocol (HSRP) gateway with First Hop Redundancy Protocol (FHRP) filtering should be configured.




---

**Note** FHRP filtering is required to allow for the existence of the same default gateway in different locations and optimize outbound traffic flows (traffic from a server in DC-3 to another datacenter). See [Overlay Transport Virtualization Best Practices Guide](#) for more information.

---

- Configure the anycast gateway MAC address used in the VXLAN fabric as the HSRP MAC address in the legacy fabric for efficient MAC moves.
- Layer 3 traffic between the legacy fabric and the IP WAN should be over OTV, MPLS L3VPN, or VRF Lite.

#### DC-4, the datacenter with the OTV+BDI switch as the one box solution

- Layer 3 routing within this datacenter is through an IGP implemented on the Layer 3 vPC switch pair at the border. When a server in DC-4 sends traffic to a server in another VLAN/subnet, either located within the site or in a remote site, the traffic first reaches the designated Layer-3 switch. If the destination is within the fabric, it routes the packets (across VLANs, through the corresponding BDI) towards the corresponding Layer 2 switch which forwards it to the destination. If the destination is in another site, the Layer-3 border switch routes the traffic towards the remote site, through a Layer 3 DCI technology like MPLS L3VPN or OTV, depending on how the host route was learnt (through MPLS VPN or OTV).
- When the Layer 3 DCI enabled border switch receives Layer 3 traffic from another site, it does a lookup to find out where the destination server resides and routes the traffic through the IGP in the fabric, to the corresponding Layer 2 switch. The switch forwards the packets to the intended destination server.

#### ARP requests

- All ARP requests are resolved through the VXLAN and/or OTV overlays.
- When ARP suppression is enabled on the VXLAN overlay NVE interface (for a specific Layer 2 virtual network[s]), then an ARP request is suppressed if the entry is present in the ARP cache. Else, the request is flooded into the VXLAN fabric.
- A Layer 3 gateway generated ARP request sent over OTV uses the primary gateway IP and virtual device context (VDC) MAC addresses, while an ARP request sent over the VXLAN BGP EVPN fabric uses the secondary BDI or anycast gateway IP and corresponding MAC addresses.




---

**Note** You should enable the ARP proxy function under the OTV overlay and ARP suppression function under the VXLAN overlay at the same time or you should disable both the functions.

---

## Layer 2 Multicast Forwarding and Layer 3 Multicast Routing

Multicast traffic across the datacenters:

#### Layer 2 Multicast Forwarding

- When a server in the VXLAN BGP EVPN fabric sends multicast traffic to the attached ToR/leaf switch, the leaf switch forwards the multicast traffic within the fabric, as explained in the Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide. If there are receivers in the Layer 2 CE pod

within DC-1, the leaf switch sends traffic towards the border switch, which VXLAN decapsulates the traffic and sends bridged traffic to relevant Layer 2 switches (or receivers in the Layer 2 CE pod). For receivers in DC-2, DC-3, or DC-4, the border leaf switch VXLAN decapsulates the traffic and sends traffic over OTV to DC-2, DC-3 (or DC-4).

- When the DC-2 border leaf switch receives the packets, it has to send traffic to receivers in the VXLAN fabric and in the Layer 2 CE pod. For receivers within the VXLAN fabric, the border leaf switch OTV decapsulates the traffic, VXLAN encapsulates it, and sends it towards relevant ToR/leaf switches. For receivers in the CE pod, the traffic is OTV decapsulated and sent towards relevant Layer 2 CE switches.
- When a DC-3 OTV switch receives the packets, the switch OTV decapsulates the traffic and sends a copy of the bridged traffic towards relevant Layer 2 CE switches/receivers.
- When the designated DC-4 OTV+BDI border vPC switch receives the packets, it OTV decapsulates the traffic and sends a copy of the bridged traffic towards relevant Layer 2 CE switches/receivers.
- Similarly, when a server in the Layer 2 CE pod (in DC-1 or DC-2) sends multicast traffic, traffic flow to receivers within the CE pod remains the same. If there are receivers in the VXLAN fabric in DC-1, and in DC-2 and DC-3, the attached access switch sends it towards the border leaf switch. The border leaf switch sends traffic to receivers in the VXLAN fabric (in DC-1), and towards DC-2 and DC-3, as explained in the previous points. If there are receivers in DC-4 (per topology 2), a copy of the packets is sent towards DC-4.
- When a server in DC-4 (OTV+BDI one box solution at the border) sends Layer 2 multicast traffic, traffic flow to receivers within the CE pod remains the same. If there are receivers in other datacenters, the switch at the datacenter border receives the packets, and sends it over OTV towards the destination datacenters (DC-1 and DC-3). If DC-4 has multicast receivers and receives multicast traffic from a sender in another datacenter, it OTV decapsulates the traffic and sends a copy of the bridged traffic towards relevant Layer 2 CE switches/receivers. The Layer 2 switch(es) forwards the packets to intended destination servers.

### Layer 3 multicast routing

When a server in DC-1 or DC-2 (in topology 1), or DC-1/DC-4 (in topology 2) sends Layer 3 multicast traffic across the fabrics, it is routed through the external, centralized multicast gateway in the legacy datacenter, DC-3. Layer 3 multicast routing support is not available for the VXLAN BGP EVPN and OTV Interoperation feature in release 8.2(1). On the border leaf switch, ensure that you do not enable the **ip pim sparse-mode** command on OTV enabled bridge domains.




---

**Attention** Layer 3 multicast traffic within a VXLAN BGP EVPN fabric (for VXLAN enabled bridge domains) is supported as usual. This is the use case wherein a sender within the VXLAN fabric sends Layer 3 multicast traffic to receivers located within the fabric. For more details, see "Multicast Routing in the VXLAN Underlay" section, "IP Fabric Underlay" chapter in the *Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide*.

---

Layer 3 multicast packet flow in DC-1 and DC-2:

- If a server in the VXLAN fabric in DC-1 sends multicast traffic to receivers in different VLANs/subnets outside the VXLAN fabric, then the traffic is sent to the Layer 3 switch in DC-3 where the centralized multicast gateway is enabled. First, the traffic reaches the border leaf switch in DC-1. Assuming that a receiver is located in the CE pod in DC-1 and another receiver in the VXLAN fabric in DC-2, the border leaf switch forwards the traffic towards DC-3. On receiving the request, the Layer 3 switch (or centralized

multicast gateway) in DC-3 sends a copy of the routed multicast traffic to the border leaf switch in DC-1 and DC-2. On receiving the traffic, the border leaf switch in DC-2 VXLAN encapsulates the traffic and sends it towards the corresponding ToR/leaf switch. The leaf switch VXLAN decapsulates the packets and sends it to the destination server. The border leaf switch in DC-1 forwards the traffic to the corresponding Layer 2 CE switch, which forwards it to the attached destination server.

Layer 3 multicast packet flow in DC-4:

- If a server in this datacenter sends multicast traffic to receivers in different VLANs/subnets, the traffic first reaches the Layer 3 OTV+BDI switch at the border. The Layer 3 switch forwards it towards DC-3, (the external, centralized multicast gateway) which sends a copy of the routed multicast traffic towards other datacenter fabrics that have receivers. The rest of the flow is the same as explained in the previous points.

## How to Configure VXLAN BGP EVPN and OTV Interoperation, and OTV with BDI



**Note** Type the `switch# configure terminal` command to enter global configuration mode (`config`)#

### Configure VXLAN BGP EVPN and OTV Interoperation on the Border Leaf Switches in DC-1 and DC-2

Only configurations relevant to the VXLAN BGP EVPN and OTV Interoperation feature are noted here. If VXLAN BGP EVPN fabric configurations are not enabled on the fabric's leaf and spine switches, enable them. See the "Configuring the VXLAN BGP EVPN" chapter, or [Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide](#).

VXLAN BGP EVPN and OTV configurations should be configured on a single box, the border leaf vPC switch pair. BL1 and BL2 configurations on DC-1:

#### Step 1 Configure tunnel stitching in the VXLAN overlay

The VXLAN BGP EVPN configurations are documented in the *Configuring VXLAN BGP EVPN* chapter. Only VXLAN configurations required for the VXLAN and OTV interoperation for the one-box solution are given.

```
BL1(config)# feature nve
              vni 40000
              system bridge-domain 2500-3500
              bridge-domain 3500
                member vni 40000
              exit
              interface nve1
                source-interface loopback0
                tunnel-stitching enable
                member vni 40000
                no suppress-arp
                mcast-group 239.1.1.65
```

- The **tunnel-stitching enable** command is the VXLAN command for connecting VXLAN and OTV tunnels.

### Step 2 Configure the VXLAN overlay loopback interface

```
BL1(config)# interface loopback0
                ip address 209.165.200.25/32
                ip address 203.0.113.1/32 secondary
```

- Primary and secondary IP addresses are assigned to the source loopback interface on the switch.

### Step 3 Configure tunnel switching in the OTV overlay

```
BL1(config)# feature otv
                otv site-vni 40000
                interface Overlay1
                    otv join-interface Ethernet5/5
                    otv extend-vni 10000, 20000
                    otv vni mapping 10000, 20000, 30000 to vlan 1000, 2000, 3000
                    otv use-adjacency-server 10.0.0.1 unicast-only
                    no otv suppress-arp-nd
                    otv tunnel-stitch
                    no shutdown
                    exit
                otv site-identifier 0000.0000.000A
                otv encapsulation-format ip gre
```

- The **otv site-vni** command enables the OTV site specific VNI. This VNI should not be extended over any overlay interface and should be operationally up before it can be configured as the OTV site VNI. At least one interface should be present where the VNI is up. On a VXLAN + OTV pod, the VNI can be configured under the NVE interface. If the **otv site-vlan** configuration is enabled, then you need to remove it before configuring the **otv site-vni** command.
- The **otv tunnel-stitch** command is the OTV command for connecting VXLAN and OTV tunnels.
- For OTV overlays, only Generic Routing Encapsulation (GRE) encapsulation is supported for the 8.2(1) release.




---

**Note** You should enable ARP proxy under the OTV overlay and ARP suppression under the VXLAN overlay at the same time or you should disable both the functions.

---

### Step 4 Configure the VXLAN and OTV tunnel stitching configurations on BL2.

### Step 5 Configure vPC function on BL1 and BL2

#### vPC Peer 1 (BL1) configuration

```
BL1(config)# interface Bdi3500
                no shutdown
                vrf member cust1
                no ip redirects
                ip address 198.51.100.20/24
                ip address 198.51.100.1/24 secondary anycast-primary
                ipv6 address 2001:DB8:1::1/64
```

```
no ipv6 redirects
fabric forwarding mode anycast-gateway
```

### vPC Peer 2 (BL2) configuration

```
BL2(config)# interface Bdi3500
no shutdown
vrf member cust1
no ip redirects
ip address 198.51.100.30/24
ip address 198.51.100.1/24 secondary anycast-primary
ipv6 address 2001:DB8:1::1/64
no ipv6 redirects
fabric forwarding mode anycast-gateway
```

- The unique, primary gateway IP address (198.51.100.20 for vPC peer 1, and 198.51.100.30 for vPC peer 2 switches) will be used for sending ARP requests over OTV. The common, secondary anycast gateway IP address (198.51.100.1) will be used for sending ARP requests on the VXLAN side.

**Step 6** The configurations are for DC-1. Similarly, configure for DC-2 too.

## Configure OTV with BDI Configuration on the Border Switches in DC-4




---

**Note** These configurations are relevant to the OTV+BDI single box solution (comprising the vPC switch pair at the border, OTV1 and OTV2) and not to the legacy datacenter, where the two box solution comprises of a separate OTV and aggregation switch.

---

OTV1 and OTV2 configurations:

### Step 1 Configure OTV on OTV1

```
OTV1(config)# feature otv
feature nv overlay
feature-set fabric
feature fabric forwarding
feature vni
vni 40000
system bridge-domain 2500-3500
bridge-domain 3500
member vni 40000
exit
encapsulation profile vni vsi_cisco
dot1q 60 vni 40000
exit
interface Ethernet 1/12
no shutdown
service instance 1 vni
encapsulation profile vsi_cisco default
no shutdown
```

The **feature nv overlay** command is a required command for this feature.

```
OTV1(config)# no otv site-vlan 111
```

```

otv site-vni 40000
interface Overlay1
  otv join-interface Ethernet2/3
  otv extend-vni 10000, 20000, 30000
  otv vni mapping 10000, 20000, 30000 to vlan 1000, 2001, 3000
  otv use-adjacency-server 10.0.0.1 192.0.2.1 unicast-only
  no otv suppress-arp-nd
  otv adjacency-server unicast-only
  no shutdown
  exit
otv site-identifier 0000.0000.000A
otv encapsulation-format ip gre

```

- The **otv site-vni** command enables the OTV site specific VNI. This VNI should not be extended over any overlay interface and should be operationally up before it can be configured as the OTV site VNI. At least one interface should be present where the VNI is up. On a OTV with BDI pod, the VNI can be brought up on a VSI interface, using an encapsulation profile, as shown in the example. If the **otv site-vlan** configuration is enabled, then you need to remove it before configuring the **otv site-vni** command.
- In an OTV with BDI single box solution, the **otv suppress-arp-nd** option is not supported in the Cisco NX-OS 8.2(1) release.
- For OTV overlays, only Generic Routing Encapsulation (GRE) encapsulation is supported for the 8.2(1) release.

### Step 2 Similarly, enable OTV configurations on OTV2.

### Step 3 Configure vPC function on OTV1 and OTV2

#### vPC Peer 1 (OTV1) configuration

```

OTV1(config)# interface Bdi3500
  no shutdown
  vrf member cust1
  no ip redirects
  ip address 209.165.201.10/24
  ip address 209.165.201.20/24 secondary anycast-primary
  fabric forwarding mode anycast-gateway

```

#### vPC Peer 1 (OTV2) configuration

```

OTV1(config)# interface Bdi3500
  no shutdown
  vrf member cust1
  no ip redirects
  ip address 209.165.201.12/24
  ip address 209.165.201.20/24 secondary anycast-primary
  fabric forwarding mode anycast-gateway

```

- The primary, unique IP address (209.165.201.10 for vPC peer 1 switch [OTV1], and 209.165.201.12 for vPC peer 2 switch [OTV2]) is created for ARP requests over OTV. The common, secondary anycast gateway IP address (209.165.201.20) is used as a Layer 3 gateway for traffic, including ARP requests, from the specified bridge domain Bdi3500 within the datacenter.

# Verifying VXLAN BGP EVPN and OTV Interoperation, and OTV with BDI

Verify VXLAN and OTV configurations on BL1:

## Display VXLAN Configuration on the Border Leaf Switch BL1

In the following example, you can see the VXLAN overlay information such as the loopback interface information, and the state of the overlay. For vPC switches, primary and secondary IP addresses are configured:

```
BL1# show nve interface

Interface: nve1, State: Up, encapsulation: VXLAN
VPC Capability: VPC-VIP-Only [not-notified]
Local Router MAC: 1005.caf4.cdd9
Host Learning Mode: Control-Plane
Source-Interface: loopback1 (primary: 209.165.200.25, secondary: 203.0.113.1)
```

In the following example, you can view bridge domain information for the specified bridge domain:

```
BL1# show bridge-domain 100

Bridge-domain 100 (2 ports in all)
Name: Bridge-Domain100
Administrative State: UP Operational State: UP
vni10000
Overlay1
nve1
```

In the following example, you can see VNI information for the VXLAN overlay, including Layer 2/Layer 3 VNI, multicast group mapping, and so on:

```
BL1# show nve vni

Codes: CP - Control Plane          DP - Data Plane
       UC - Unconfigured           SA - Suppress ARP

Interface VNI      Multicast-group  State Mode Type [BD/VRF]      Flags
-----
nve1      5000            224.1.1.1        Up   CP   L3 [cust1]
```

In the following example, you can see attached VXLAN overlay peer information such as the peer IP address, state, and so on:

```
BL1# show nve peers

Interface Peer-IP      State  LearnType  Uptime  Router-Mac
-----
nve1      10.1.1.2      Up     CP         4d19h   8c60.4f04.21c2
```

## Display OTV Configuration on the Border Leaf Switch BL1

In the following example, you can see the extended VLANs and VNIs in the **Extended vlans** and **Extended vni** fields. Also, the OTV join interface and adjacency server information are displayed:

```
BL1# show otv

OTV Overlay Information
Site Identifier 0000.0000.000a
Encapsulation-Format ip - gre

Overlay interface Overlay1

VPN name           : Overlay1
VPN state          : UP
Extended vlans     : 50 60 100-101 300-301 (Total:6)
Extended vni      : 10000, 20000, 30000, 40000(total:4)
Join interface(s) : Eth5/5 (10.3.3.3)
Site vlan         : 101 (up)
AED-Capable       : Yes
Capability         : Unicast-Only
Is Adjacency Server : No
Adjacency Server(s) : 10.0.0.1 / [None]
```

In the following example, site VLAN, site VNI, and other site details are displayed. You should remove the site VLAN using the **no otv site-vlan** command.

```
BL1# show otv site detail

Full      - Both site and overlay adjacency up
Partial   - Either site/overlay adjacency down
Down      - Both adjacencies are down (Neighbor is down/unreachable)
(!)       - Site-ID mismatch detected

Local Edge Device Information:
  Hostname BL1
  System-ID 002a.6a65.09c1
  Site-Identifier 0000.0000.000a
  Site-VLAN 101 State is Up

  Site-VLAN-Cfg 0
  Site-VNI 10001

Site Information for Overlay1:

Bridge-domain 101 (1 ports in all)
Name:: Bridge-Domain101
Administrative State: UP                               Operational State: UP
      vni10001
      nve1
```

## Display OTV Overlay State

In the following example, the **Tunnel Stitch** field displays **enabled** state, indicating that the OTV part of the VXLAN OTV tunnel stitching is enabled:

```
BL1# show otv internal overlay detail
```

```

Overlay interface Overlay1 idx: 1635021663
Overlay interface idx: 0x22080001
  VPN name       : Overlay1
  Device ID      : 8c60.4f04.21c6
  Protocol state : UP           Interface state      : UP
  Tunnel Stitch  : enabled
  Tier Id        : 1
  Pend Tier Id   : 0
  Tier Id State  : Done(4)
  Refcount      : 6
.
.

```

In the following example, the overlay ID and the VXLAN-to-OTV tunnel stitching status are displayed. In the **Enable** column, **Y** indicates that the OTV part of the tunnel is enabled:

```
BL1# show otv internal tunnel-stitch
```

Overlay	Enable	Tier-id	Tier-pend-id	State	Detail
Overlay1	Y	1	0	4	Done

In the following example, you can see that ARP proxy is enabled under the OTV overlay:



**Note** Ensure that you also enable ARP suppression under the VXLAN overlay at the same time. Alternatively, disable the ARP proxy/suppression function under both the overlays.

```
BL1# show otv internal arp-nd status
```

```

Overlay: Overlay1
  Suppress arp-nd: Enabled
  VNI Suppress ARP:
  50000      1000      : Enabled
  50001      1001      : Enabled
  50002      1002      : Enabled
  50003      1003      : Enabled
  50004      1004      : Enabled
  50005      1005      : Enabled
  50006      1006      : Enabled
  50007      1007      : Enabled

```

In the following example, for the specified OTV overlay (Overlay1), VNI, BD, and VLAN mapping information is displayed. The **State** field indicates that the overlay is functional.

```
BL1# show otv vni Overlay1
```

Overlay	VNI	BD	VLAN	Flag	State
Overlay1	10000	100	1000	0x1	(12)Overlay Logical up
Overlay1	20000	101	2000	0x1	(12)Overlay Logical up
Overlay1	30000	300	3000	0x1	(12)Overlay Logical up
Overlay1	40000	301	3500	0x1	(12)Overlay Logical up

## Display OTV Adjacencies

In the following example, OTV adjacencies towards the legacy datacenter switches OTV1 and OTV2 in DC-3, BL2 in DC-1, and BL3 in DC-2 are displayed:

```
BL1# show otv adjacency

Overlay Adjacency database

Overlay-Interface Overlay1 :

Hostname          System-ID          Dest Addr          Up Time  State
-----
LEGACY-OTV1       002a.6a65.0045    209.165.200.25    3d01h   UP
LEGACY-OTV2       002a.6a65.0046    203.0.113.1       3d01h   UP
core-BL2          8c60.4f04.21c7    203.0.113.200     3d01h   UP
BL3               8c60.4f04.43c1    203.0.113.250     3d01h   UP
```

## Display Tier IDs

In the following example, you can see tier IDs allocated to the VXLAN and OTV overlays. Unique tier IDs are automatically allocated so that packets that are sent from a tunnel are not sent back to the same tunnel, such as multicast packets being sent back to the multicast source. OTV and VXLAN overlay tier IDs are unique:

```
BL1# show forwarding distribution tierpeerid otv

Tier-Peer-id allocations:

Type: Tier ID          Tier ID: 0x1  Tier Peer ID: 0x1000          Data: 0x0

Type: Tier Peer ID     Tier ID: 0x0  Tier Peer ID: 0x3fd  App: OTV  Data: 0x0
Type: Tier Peer ID     Tier ID: 0x1  Tier Peer ID: 0x1001  App: OTV  Data: 0x1
Type: Tier Peer ID     Tier ID: 0x1  Tier Peer ID: 0x1002  App: OTV  Data: 0x2
Type: Tier Peer ID     Tier ID: 0x1  Tier Peer ID: 0x13fd  App: OTV  Data: 0x0

BL1# show forwarding distribution tierpeerid nve

Tier-Peer-id allocations:

Type: Tier ID          Tier ID: 0x2  Tier Peer ID: 0x2000          Data: 0x0

Type: Tier Peer ID     Tier ID: 0x2  Tier Peer ID: 0x2001  App: NVE  Data: 0x1c1c1c1c
Type: Tier Peer ID     Tier ID: 0x2  Tier Peer ID: 0x2002  App: NVE  Data: 0x1b1b1b1b
Type: Tier Peer ID     Tier ID: 0x2  Tier Peer ID: 0x23fd  App: NVE  Data: 0x0
```

## Troubleshooting VXLAN BGP EVPN and OTV Interoperation, and OTV with BDI

```
BL1# show tech-support arp
BL1# show tech-support otv
BL1# show tech-support lim
BL1# show tech-support oim
BL1# show tech-support nve
```

```
BL1# show tech-support vxlan-evpn
```

## Feature History for VXLAN BGP EVPN and OTV Interoperation

This table lists the release history for this feature.

**Table 1: Feature History for VXLAN BGP EVPN and OTV interoperation**

Feature Name	Release	Feature Information
VXLAN BGP EVPN and OTV interoperation	8.2(1)	<p>This feature was introduced to enable VXLAN and OTV (with BDI) configurations on a single switch (as a one box solution), and to enable OTV+BDI configurations on a single switch (as a one box solution).</p> <p>The following commands were introduced or modified for the OTV overlay:</p> <p><b>otv tunnel-stitch</b></p> <p><b>otv vni mapping</b></p> <p><b>show otv</b></p> <p><b>show otv internal overlay detail</b></p> <p><b>show otv internal tunnel-stitch</b></p> <p><b>show otv vni overlay</b></p> <p>The following command was introduced for the VXLAN overlay:</p> <p><b>tunnel-stitching enable</b></p> <p>The following command was introduced for the BDI function:</p> <p><b>ip address secondary anycast-primary</b></p>

