



Configuring VXLAN BGP EVPN

This chapter contains the following sections:

- [Information About VXLAN BGP EVPN, on page 1](#)
- [Configuring VXLAN BGP EVPN, on page 4](#)
- [Configuring Anycast Gateway for VXLAN Routing, on page 6](#)
- [Configuring the NVE Interface and VNIs, on page 6](#)
- [Configuring BGP on the VTEP, on page 7](#)
- [Configuring RD and Route Targets for VXLAN Bridging, on page 8](#)
- [Configuring BGP for EVPN on the Spine, on page 8](#)
- [Disabling VXLANs, on page 9](#)
- [Duplicate Detection for IP and MAC Addresses, on page 10](#)
- [Verifying the VXLAN Configuration, on page 11](#)
- [Example of VXLAN BGP EVPN \(EBGP\), on page 12](#)
- [Example of VXLAN BGP EVPN \(IBGP\), on page 21](#)
- [Example Show Commands, on page 28](#)

Information About VXLAN BGP EVPN

Guidelines and Limitations for VXLAN BGP EVPN

VXLAN BGP EVPN has the following guidelines and limitations:

- The Cisco Nexus 3600 platform switch is not supported as a border spine or border leaf.
- SVI and sub-interfaces as core links are not supported along with Layer 2 GW configurations.
- In a VXLAN EVPN setup, border leaves must use unique route distinguishers, preferably using **auto rd** command. It is not supported to have same route distinguishers in different border leaves.
- ARP suppression is only supported for a VNI if the VTEP hosts the First-Hop Gateway (Distributed Anycast Gateway) for this VNI. The VTEP and the SVI for this VLAN have to be properly configured for the distributed anycast gateway operation, for example, global anycast gateway MAC address configured and anycast gateway feature with the virtual IP address on the SVI.
- The **show** commands with the **internal** keyword are not supported.

- DHCP snooping (Dynamic Host Configuration Protocol snooping) is not supported on VXLAN VLANs.
- SPAN for VXLAN uplink interface is not supported.
- ACLs are not supported on Layer 3 uplinks for VXLAN traffic.
- RACLs and PACLs are not supported for VXLAN VLANs.
- QoS classification is not supported for VXLAN VLANs.
- Uplink ports can be of type Layer 3 interface, sub-interface, or a Layer 3 port-channel interface. However with Layer 2 GW sub-interface uplink ports are not supported.
- For EBGp, it is recommended to use a single overlay EBGp EVPN session between loopbacks.
- Bind NVE to a loopback address that is separate from other loopback addresses that are required by Layer 3 protocols. A best practice is to use a dedicated loopback address for VXLAN.
- VXLAN BGP EVPN does not support an NVE interface in a non-default VRF.
- It is recommended to configure a single BGP session over the loopback for an overlay BGP session.
- The VXLAN UDP port number is used for VXLAN encapsulation. For Cisco Nexus NX-OS, the UDP port number is 4789. It complies with IETF standards and is not configurable.
- VXLAN does not support co-existence with the MPLS feature.
- VXLAN with Layer 3 VPN is not supported.
- VXLAN with ingress replication is not supported.
- MLD snooping is not supported on VXLAN VLANs.
- DHCP snooping is not supported on VXLAN VLANs.

Considerations for VXLAN BGP EVPN Deployment

- A loopback address is required when using the **source-interface config** command. The loopback address represents the local VTEP IP.
- To establish IP multicast routing in the core, IP multicast configuration, PIM configuration, and RP configuration is required.
- VTEP to VTEP unicast reachability can be configured through any IGP/BGP protocol.
- As a best practice when changing the IP address of a VTEP device, enter the **shut** command on the loopback interface used by the NVE interface and then enter the **no shut** command before changing the IP address.
- Every tenant VRF needs a VRF overlay VLAN and SVI for VXLAN routing.

Network Considerations for VXLAN Deployments

- MTU Size in the Transport Network

Due to the MAC-to-UDP encapsulation, VXLAN introduces 50-byte overhead to the original frames. Therefore, the maximum transmission unit (MTU) in the transport network needs to be increased by 50

bytes. If the overlays use a 1500-byte MTU, the transport network needs to be configured to accommodate 1550-byte packets at a minimum. Jumbo-frame support in the transport network is required if the overlay applications tend to use larger frame sizes than 1500 bytes.

- **ECMP and LACP Hashing Algorithms in the Transport Network**

As described in a previous section, Cisco Nexus 3600 platform switches introduce a level of entropy in the source UDP port for ECMP and LACP hashing in the transport network. As a way to augment this implementation, the transport network uses an ECMP or LACP hashing algorithm that takes the UDP source port as an input for hashing, which achieves the best load-sharing results for VXLAN encapsulated traffic.

- **Multicast Group Scaling**

The VXLAN implementation on Cisco Nexus 3600 platform switches uses multicast tunnels for broadcast, unknown unicast, and multicast traffic forwarding. Ideally, one VXLAN segment mapping to one IP multicast group is the way to provide the optimal multicast forwarding. It is possible, however, to have multiple VXLAN segments share a single IP multicast group in the core network. VXLAN can support up to 16 million logical Layer 2 segments, using the 24-bit VNID field in the header. With one-to-one mapping between VXLAN segments and IP multicast groups, an increase in the number of VXLAN segments causes a parallel increase in the required multicast address space and the amount of forwarding states on the core network devices. At some point, multicast scalability in the transport network can become a concern. In this case, mapping multiple VXLAN segments to a single multicast group can help conserve multicast control plane resources on the core devices and achieve the desired VXLAN scalability. However, this mapping comes at the cost of suboptimal multicast forwarding. Packets forwarded to the multicast group for one tenant are now sent to the VTEPs of other tenants that are sharing the same multicast group. This causes inefficient utilization of multicast data plane resources. Therefore, this solution is a trade-off between control plane scalability and data plane efficiency.

Despite the suboptimal multicast replication and forwarding, having multiple-tenant VXLAN networks to share a multicast group does not bring any implications to the Layer 2 isolation between the tenant networks. After receiving an encapsulated packet from the multicast group, a VTEP checks and validates the VNID in the VXLAN header of the packet. The VTEP discards the packet if the VNID is unknown to it. Only when the VNID matches one of the VTEP's local VXLAN VNIDs, does it forward the packet to that VXLAN segment. Other tenant networks will not receive the packet. Thus, the segregation between VXLAN segments is not compromised.

Considerations for the Transport Network

The following are considerations for the configuration of the transport network:

- On the VTEP device:
 - Enable and configure IP multicast.
 - Create and configure a loopback interface with a /32 IP address.
 - Enable IP multicast on the loopback interface.
 - Advertise the loopback interface /32 addresses through the routing protocol (static route) that runs in the transport network.
 - Enable IP multicast on the uplink outgoing physical interface.
- Throughout the transport network:

- Enable and configure IP multicast.

BGP EVPN Considerations for VXLAN Deployment

Configuring VXLAN BGP EVPN

Enabling VXLAN

Enable VXLAN and the EVPN.

Procedure

	Command or Action	Purpose
Step 1	feature vn-segment	Enable VLAN-based VXLAN
Step 2	feature nv overlay	Enable VXLAN
Step 3	nv overlay evpn	Enable the EVPN control plane for VXLAN.

Configuring VLAN and VXLAN VNI

Procedure

	Command or Action	Purpose
Step 1	vlan <i>number</i>	Specify VLAN.
Step 2	vn-segment <i>number</i>	Map VLAN to VXLAN VNI to configure Layer 2 VNI under VXLAN VLAN.

Configuring VRF for VXLAN Routing

Configure the tenant VRF.

Procedure

	Command or Action	Purpose
Step 1	vrf context <i>vxlan</i>	Configure the VRF.
Step 2	vni <i>number</i>	Specify VNI.
Step 3	rd auto	Specify VRF RD (route distinguisher).
Step 4	address-family ipv4 unicast	Configure address family for IPv4.

	Command or Action	Purpose
Step 5	route-target both auto	Note Specifying the auto option is applicable only for IBGP. Manually configured route targets are required for EBGp.
Step 6	route-target both auto evpn	Note Specifying the auto option is applicable only for IBGP. Manually configured route targets are required for EBGp.
Step 7	address-family ipv6 unicast	Configure address family for IPv6.
Step 8	route-target both auto	Note Specifying the auto option is applicable only for IBGP. Manually configured route targets are required for EBGp.
Step 9	route-target both auto evpn	Note Specifying the auto option is applicable only for IBGP. Manually configured route targets are required for EBGp.

Configuring SVI for Hosts for VXLAN Routing

Configure the SVI for hosts.

Procedure

	Command or Action	Purpose
Step 1	vlan <i>number</i>	Specify VLAN
Step 2	interface <i>vlan-number</i>	Specify VLAN interface.
Step 3	vrf member <i>vxlان-number</i>	Configure SVI for host.
Step 4	ip address <i>address</i>	Specify IP address.

Configuring VRF Overlay VLAN for VXLAN Routing

Procedure

	Command or Action	Purpose
Step 1	vlan <i>number</i>	Specify VLAN.

	Command or Action	Purpose
Step 2	vn-segment <i>number</i>	Specify vn-segment.

Configuring VNI Under VRF for VXLAN Routing

Configures a Layer 3 VNI under a VRF overlay VLAN. (A VRF overlay VLAN is a VLAN that is not associated with any server facing ports. All VXLAN VNIs that are mapped to a VRF, need to have their own internal VLANs allocated to it.)

Procedure

	Command or Action	Purpose
Step 1	vrf context <i>vxlan</i>	Create a VXLAN Tenant VRF
Step 2	vni <i>number</i>	Configure Layer 3 VNI under VRF.

Configuring Anycast Gateway for VXLAN Routing

Procedure

	Command or Action	Purpose
Step 1	fabric forwarding anycast-gateway-mac <i>address</i>	Configure distributed gateway virtual MAC address Note One virtual MAC per VTEP Note All VTEPs should have the same virtual MAC address
Step 2	fabric forwarding mode anycast-gateway	Associate SVI with anycast gateway under VLAN configuration mode.

Configuring the NVE Interface and VNIs

Procedure

	Command or Action	Purpose
Step 1	interface <i>nve-interface</i>	Configure the NVE interface.
Step 2	host-reachability protocol bgp	This defines BGP as the mechanism for host reachability advertisement

	Command or Action	Purpose
Step 3	member vni <i>vni</i> associate-vrf	Add Layer-3 VNIs, one per tenant VRF, to the overlay. Note Required for VXLAN routing only.
Step 4	member vni <i>vni</i>	Add Layer 2 VNIs to the tunnel interface. switch# member vni 900001 associate-vrf
Step 5	mcast-group <i>address</i>	Configure the mcast group on a per-VNI basis

Configuring BGP on the VTEP

Procedure

	Command or Action	Purpose
Step 1	router bgp <i>number</i>	Configure BGP.
Step 2	router-id <i>address</i>	Specify router address.
Step 3	neighbor <i>address</i> remote-as <i>number</i>	Define MP-BGP neighbors. Under each neighbor define l2vpn evpn.
Step 4	address-family ipv4 unicast	Configure address family for IPv4.
Step 5	address-family l2vpn evpn	Configure address family Layer 2 VPN EVPN under the BGP neighbor. Note Address-family ipv4 evpn for vxlan host-based routing
Step 6	(Optional) Allowas-in	Allows duplicate AS numbers in the AS path. Configure this parameter on the leaf for eBGP when all leaves are using the same AS, but the spines have a different AS than leaves.
Step 7	send-community extended	Configures community for BGP neighbors.
Step 8	vrf <i>vrf-name</i>	Specify VRF.
Step 9	address-family ipv4 unicast	Configure address family for IPv4.
Step 10	advertise <i>l2vpn evpn</i>	Enable advertising EVPN routes.
Step 11	address-family ipv6 unicast	Configure address family for IPv6.
Step 12	advertise <i>l2vpn evpn</i>	Enable advertising EVPN routes.

Configuring RD and Route Targets for VXLAN Bridging

Procedure

	Command or Action	Purpose
Step 1	evpn	Configure VRF.
Step 2	vni number l2	Note Only Layer 2 VNIs need to be specified.
Step 3	rd auto	Define VRF RD (route distinguisher) to configure VRF context.
Step 4	route-target import auto	Define VRF Route Target and import policies.
Step 5	route-target export auto	Define VRF Route Target and export policies.

Configuring BGP for EVPN on the Spine

Procedure

	Command or Action	Purpose
Step 1	route-map permitall permit 10	Configure route-map. Note The route-map keeps the next-hop unchanged for EVPN routes. • Required for eBGP. • Optional for iBGP.
Step 2	set ip next-hop unchanged	Set next-hop address. Note The route-map keeps the next-hop unchanged for EVPN routes. • Required for eBGP. • Optional for iBGP.
Step 3	router bgp autonomous system number	Specify BGP.
Step 4	address-family l2vpn evpn	Configure address family Layer 2 VPN EVPN under the BGP neighbor.

	Command or Action	Purpose
Step 5	retain route-target all	Configure retain route-target all under address-family Layer 2 VPN EVPN [global]. Note Required for eBGP. Allows the spine to retain and advertise all EVPN routes when there are no local VNI configured with matching import route targets.
Step 6	neighbor address remote-as number	Define neighbor.
Step 7	address-family l2vpn evpn	Configure address family Layer 2 VPN EVPN under the BGP neighbor.
Step 8	disable-peer-as-check	Disables checking the peer AS number during route advertisement. Configure this parameter on the spine for eBGP when all leafs are using the same AS but the spines have a different AS than leafs. Note Required for eBGP.
Step 9	send-community extended	Configures community for BGP neighbors.
Step 10	route-map permitall out	Applies route-map to keep the next-hop unchanged. Note Required for eBGP.

Disabling VXLANs

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters configuration mode.
Step 2	no nv overlay evpn	Disables EVPN control plane.
Step 3	no feature vn-segment-vlan-based	Disables the global mode for all VXLAN bridge domains
Step 4	no feature nv overlay	Disables the VXLAN feature.
Step 5	(Optional) copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Duplicate Detection for IP and MAC Addresses

Cisco NX-OS supports duplicate detection for IP and MAC addresses. This enables the detection of duplicate IP or MAC addresses based on the number of moves in a given time-interval (seconds).

The default is 5 moves in 180 seconds. (Default number of moves is 5 moves. Default time-interval is 180 seconds.)

- For IP addresses:
 - After the 5th move within 180 seconds, the switch starts a 30 second lock (hold down timer) before checking to see if the duplication still exists (an effort to prevent an increment of the sequence bit). This 30 second lock can occur 5 times (this means 5 moves in 180 seconds for 5 times) before the switch permanently locks or freezes the duplicate entry.
- For MAC addresses:
 - After the 5th move within 180 seconds, the switch starts a 30 second lock (hold down timer) before checking to see if the duplication still exists (an effort to prevent an increment of the sequence bit). This 30 second lock can occur 3 times (this means 5 moves in 180 seconds for 3 times) before the switch permanently locks or freezes the duplicate entry.

The following are example commands to help the configuration of the number of VM moves in a specific time interval (seconds) for duplicate IP-detection:

Command	Description
<pre>switch(config)# fabric forwarding ? anycast-gateway-mac dup-host-ip-addr-detection</pre>	Available sub-commands: • Anycast gateway MAC of the switch. • To detect duplicate host addresses in n seconds.
<pre>switch(config)# fabric forwarding dup-host-ip-addr-detection ? <1-1000></pre>	The number of host moves allowed in n seconds. The range is 1 to 1000 moves; default is 5 moves.
<pre>switch(config)# fabric forwarding dup-host-ip-addr-detection 100 ? <2-36000></pre>	The duplicate detection timeout in seconds for the number of host moves. The range is 2 to 36000 seconds; default is 180 seconds.
<pre>switch(config)# fabric forwarding dup-host-ip-addr-detection 100 10</pre>	Detects duplicate host addresses (limited to 100 moves) in a period of 10 seconds.

The following are example commands to help the configuration of the number of VM moves in a specific time interval (seconds) for duplicate MAC-detection:

Command	Description
<pre>switch(config)# l2rib dup-host-mac-detection ? <1-1000> default</pre>	Available sub-commands for L2RIB: • The number of host moves allowed in n seconds. The range is 1 to 1000 moves. • Default setting (5 moves in 180 in seconds).
<pre>switch(config)# l2rib dup-host-mac-detection 100 ? <2-36000></pre>	The duplicate detection timeout in seconds for the number of host moves. The range is 2 to 36000 seconds; default is 180 seconds.
<pre>switch(config)# l2rib dup-host-mac-detection 100 10</pre>	Detects duplicate host addresses (limited to 100 moves) in a period of 10 seconds.

Verifying the VXLAN Configuration


To display the VXLAN configuration information, enter one of the following commands:

Command	Purpose
show tech-support vxlan	Displays related VXLAN tech-support information.
show logging level nve	Displays logging level.
show tech-support nve	Displays related NVE tech-support information.
show tech-support vxlan-evpn	Displays related VXLAN EVPN tech-support information.
show tech-support vxlan platform	Displays VXLAN platform related tech-support information.
show run interface nve	Displays NVE overlay interface configuration.
show nve interface	Displays NVE overlay interface status.
show nve peers	Displays NVE peer status.
show nve peers <i>peer_IP_address</i> interface <i>interface_ID</i> counters	Displays per NVE peer statistics.
clear nve peers <i>peer_IP_address</i> interface <i>interface_ID</i> counters	Clears per NVE peer statistics.
show nve vni	Displays VXLAN VNI status.
show nve vxlan-params	Displays VXLAN parameters, such as VXLAN destination or UDP port.

Example of VXLAN BGP EVPN (EBGP)

An example of a VXLAN BGP EVPN (EBGP):

Figure 1: VXLAN BGP EVPN Topology (EBGP)



EBGP between Spine and Leaf

- Spine (9504-A)

- Enable the EVPN control plane

```
nv overlay evpn
```

- Enable the relevant protocols

```
feature bgp
feature pim
```

- Configure Loopback for local VTEP IP, and BGP

```
interface loopback0
 ip address 10.1.1.1/32
 ip pim sparse-mode
```

- Configure Loopback for Anycast RP

```
interface loopback1
```

```
ip address 100.1.1.1/32
ip pim sparse-mode
```

- Configure Anycast RP

```
ip pim rp-address 100.1.1.1 group-list 225.0.0.0/8
ip pim rp-candidate loopback1 group-list 225.0.0.0/8
ip pim log-neighbor-changes
ip pim ssm range 232.0.0.0/8
ip pim anycast-rp 100.1.1.1 10.1.1.1
ip pim anycast-rp 100.1.1.1 20.1.1.1
```

- Configure route-map used by EBGp for Spine

```
route-map permitall permit 10
set ip next-hop unchanged
```

- Enable OSPF for underlay routing

```
router ospf 1
log-adjacency-changes detail
```

- Configure interfaces for Spine-leaf interconnect

```
interface Ethernet4/2
ip address 192.168.1.42/24
ip pim sparse-mode
no shutdown

interface Ethernet4/3
ip address 192.168.2.43/24
ip pim sparse-mode
no shutdown
```

- Configure the BGP overlay for the EVPN address family.

```
router bgp 100
router-id 10.1.1.1
address-family l2vpn evpn
nexthop route-map permitall
retain route-target all
neighbor 30.1.1.1 remote-as 200
update-source loopback0
ebgp-multihop 3
address-family l2vpn evpn
disable-peer-as-check
send-community extended
route-map permitall out
neighbor 40.1.1.1 remote-as 200
update-source loopback0
ebgp-multihop 3
address-family l2vpn evpn
disable-peer-as-check
send-community extended
route-map permitall out
```

- Configure the BGP underlay.

```
neighbor 192.168.1.43 remote-as 200
```

```

address-family ipv4 unicast
  allowas-in
  disable-peer-as-check

```

- Spine (9504-B)

- Enable the EVPN control plane and the relevant protocols

```

feature telnet
feature nxapi
feature bash-shell
feature scp-server
nv overlay evpn
feature bgp
feature pim
feature lldp

```

- Configure Anycast RP

```

ip pim rp-address 100.1.1.1 group-list 225.0.0.0/8
ip pim rp-candidate loopback1 group-list 225.0.0.0/8
ip pim log-neighbor-changes
ip pim ssm range 232.0.0.0/8
ip pim anycast-rp 100.1.1.1 10.1.1.1
ip pim anycast-rp 100.1.1.1 20.1.1.1
vlan 1-1002
route-map permitall permit 10
  set ip next-hop unchanged

```

- Configure interfaces for Spine-leaf interconnect

```

interface Ethernet4/2
  ip address 192.168.4.42/24
  no shutdown

interface Ethernet4/3
  ip address 192.168.3.43/24
  no shutdown

```

- Configure Loopback for local VTEP IP, and BGP

```

interface loopback0
  ip address 20.1.1.1/32

```

- Configure the BGP overlay for the EVPN address family.

```

router bgp 100
  router-id 20.1.1.1
  address-family l2vpn evpn
    retain route-target all
  neighbor 30.1.1.1 remote-as 200
  update-source loopback0
  ebgp-multihop 3
  address-family l2vpn evpn
    disable-peer-as-check
    send-community extended
    route-map permitall out
  neighbor 40.1.1.1 remote-as 200

```

```
ebgp-multihop 3
address-family l2vpn evpn
disable-peer-as-check
send-community extended
route-map permitall out
```

- Configure the BGP underlay.

```
neighbor 192.168.1.43 remote-as 200
address-family ipv4 unicast
allowas-in
disable-peer-as-check
```

- Leaf (9396-A)

- Enable the EVPN control plane

```
nv overlay evpn
```

- Enable the relevant protocols

```
feature bgp
feature interface-vlan
feature dhcp
```

- Enable VxLAN with distributed anycast-gateway using BGP EVPN

```
feature vn-segment-vlan-based
feature nv overlay
fabric forwarding anycast-gateway-mac 0000.2222.3333
```

- Enable PIM RP

```
ip pim rp-address 100.1.1.1 group-list 225.0.0.0/8
```

- Configure Loopback for BGP

```
interface loopback0
ip address 30.1.1.1/32
```

- Configure Loopback for local VTEP IP

```
interface loopback1
ip address 50.1.1.1/32
```

- Configure interfaces for Spine-leaf interconnect

```
interface Ethernet2/2
no switchport
load-interval counter 1 5
ip address 192.168.1.22/24
no shutdown

interface Ethernet2/3
no switchport
load-interval counter 1 5
```

```
ip address 192.168.3.23/24
no shutdown
```

- Create the VRF overlay VLAN and configure the vn-segment.

```
vlan 101
  vn-segment 900001
```

- Configure VRF overlay VLAN/SVI for the VRF

```
interface Vlan101
  no shutdown
  vrf member vxlan-900001
```

- Create VLAN and provide mapping to VXLAN

```
vlan 1001
  vn-segment 2001001
vlan 1002
  vn-segment 2001002
```

- Create VRF and configure VNI

```
vrf context vxlan-900001
  vni 900001
```

```
rd auto
address-family ipv4 unicast
  route-target import 65535:101 evpn
  route-target export 65535:101 evpn
  route-target import 65535:101
  route-target export 65535:101
address-family ipv6 unicast
  route-target import 65535:101 evpn
  route-target export 65535:101 evpn
  route-target import 65535:101
  route-target export 65535:101
```

- Create server facing SVI and enable distributed anycast-gateway

```
interface Vlan1001
  no shutdown
  vrf member vxlan-900001
  ip address 4.1.1.1/24
  ipv6 address 4:1:0:1::1/64
  fabric forwarding mode anycast-gateway
  ip dhcp relay address 192.168.100.1 use-vrf default

interface Vlan1002
  no shutdown
  vrf member vxlan-900001
  ip address 4.2.2.1/24
  ipv6 address 4:2:0:1::1/64
  fabric forwarding mode anycast-gateway
```

- Create the network virtualization endpoint (NVE) interface


```

interface nve1
  no shutdown
  source-interface loopback1
  host-reachability protocol bgp
  member vni 10000 associate-vrf
  mcast-group 224.1.1.1
  member vni 10001 associate-vrf
  mcast-group 224.1.1.1
  member vni20000
  suppress-arp
  mcast-group 225.1.1.1
  member vni 20001
  suppress-arp
  mcast-group 225.1.1.1

```

- Configure interfaces for hosts/servers.

```

interface Ethernet1/47
  switchport access vlan 1002
interface Ethernet1/48
  switchport access vlan 1001

```

- Configure BGP

```

router bgp 200
router-id 30.1.1.1
  neighbor 10.1.1.1 remote-as 100
    update-source loopback0
    ebgp-multihop 3
    allowas-in
    send-community extended
  address-family l2vpn evpn
    allowas-in
    send-community extended
  neighbor 20.1.1.1 remote-as 100
    update-source loopback0
    ebgp-multihop 3
    allowas-in
    send-community extended
  address-family l2vpn evpn
    allowas-in
    send-community extended
vrf vxlan-900001

  advertise l2vpn evpn

evpn
  vni 2001001 l2
  vni 2001002 l2

rd auto
route-target import auto
route-target export auto

```

```

router bgp 200
router-id 30.1.1.1
  neighbor 10.1.1.1 remote-as 100
    update-source loopback0
    ebgp-multihop 3

```

```

        allowas-in
        send-community extended
    address-family l2vpn evpn
        allowas-in
        send-community extended
    neighbor 20.1.1.1 remote-as 100
    update-source loopback0
    ebgp-multihop 3
        allowas-in
        send-community extended
    address-family l2vpn evpn
        allowas-in
        send-community extended
    vrf vxlan-900001
    advertise l2vpn evpn

evpn
    vni 2001001 l2
    vni 2001002 l2

```

- Leaf (9396-B)

- Enable the EVPN control plane functionality and the relevant protocols

```

feature telnet
feature nxapi
feature bash-shell
feature scp-server
nv overlay evpn
feature bgp
feature pim
feature interface-vlan
feature vn-segment-vlan-based
feature lldp
feature nv overlay

```

- Enable VxLAN with distributed anycast-gateway using BGP EVPN

```

fabric forwarding anycast-gateway-mac 0000.2222.3333

```

- Create the VRF overlay VLAN and configure the vn-segment

```

vlan 1-1002
vlan 101
    vn-segment 900001

```

- Create VLAN and provide mapping to VXLAN

```

vlan 1001
    vn-segment 2001001
vlan 1002
    vn-segment 2001002

```

- Create VRF and configure VNI

```

vrf context vxlan-900001
    vni 900001

```

```

rd auto
address-family ipv4 unicast
  route-target import 65535:101 evpn
  route-target export 65535:101 evpn
  route-target import 65535:101
  route-target export 65535:101
address-family ipv6 unicast
  route-target import 65535:101 evpn
  route-target export 65535:101 evpn
  route-target import 65535:101 evpn
  route-target export 65535:101 evpn

```

- Configure internal control VLAN/SVI for the VRF

```

interface Vlan1

interface Vlan101
  no shutdown
  vrf member vxlan-900001

```

- Create server facing SVI and enable distributed anycast-gateway

```

interface Vlan1001
  no shutdown
  vrf member vxlan-900001
  ip address 4.1.1.1/24
  ipv6 address 4:1:0:1::1/64
  fabric forwarding mode anycast-gateway

interface Vlan1002
  no shutdown
  vrf member vxlan-900001
  ip address 4.2.2.1/24
  ipv6 address 4:2:0:1::1/64
  fabric forwarding mode anycast-gateway

```

- Create the network virtualization endpoint (NVE) interface

```

interface nve1
  no shutdown
  source-interface loopback1
  host-reachability protocol bgp
  member vni 10000 associate-vrf
  mcast-group 224.1.1.1
  member vni 10001 associate-vrf
  mcast-group 224.1.1.1
  member vni20000
  suppress-arp
  mcast-group 225.1.1.1
  member vni 20001
  suppress-arp
  mcast-group 225.1.1.1

```

- Configure interfaces for hosts/servers

```

interface Ethernet1/47
  switchport access vlan 1002

```

```
interface Ethernet1/48
  switchport access vlan 1001
```

- Configure interfaces for Spine-leaf interconnect

```
interface Ethernet2/1

interface Ethernet2/2
  no switchport
  load-interval counter 1 5
  ip address 192.168.4.22/24
  ip pim sparse-mode
  no shutdown

interface Ethernet2/3
  no switchport
  load-interval counter 1 5
  ip address 192.168.2.23/24
  ip pim sparse-mode
  no shutdown
```

- Configure Loopback for BGP

```
interface loopback0
  ip address 40.1.1.1/32
```

- Configure Loopback for local VTEP IP

```
interface loopback1
  ip address 51.1.1.1/32
  ip pim sparse-mode
```


- Configure BGP

```
router bgp 200
router-id 40.1.1.1
  neighbor 10.1.1.1 remote-as 100
    update-source loopback0
    ebgp-multihop 3
    allowas-in
    send-community extended
  address-family l2vpn evpn
    allowas-in
    send-community extended
  neighbor 20.1.1.1 remote-as 100
    update-source loopback0
    ebgp-multihop 3
    allowas-in
    send-community extended
  address-family l2vpn evpn
    allowas-in
    send-community extended
vrf vxlan-900001
  advertise l2vpn evpn
```

Example of VXLAN BGP EVPN (IBGP)

An example of a VXLAN BGP EVPN (IBGP):

Figure 2: VXLAN BGP EVPN Topology (IBGP)



IBGP between Spine and Leaf

- Spine (9504-A)
 - Enable the EVPN control plane


```
nv overlay evpn
```
 - Enable the relevant protocols


```
feature ospf
feature bgp
```
 - Configure Loopback for local VTEP IP, and BGP


```
interface loopback0
ip address 10.1.1.1/32
ip router ospf 1 area 0.0.0.0
```
 - Enable OSPF for underlay routing

```
router ospf 1
```

- Configure interfaces for Spine-leaf interconnect

```
interface Ethernet4/2
 ip address 192.168.1.42/24
 ip router ospf 1 area 0.0.0.0
 no shutdown
```

```
interface Ethernet4/3
 ip address 192.168.2.43/24
 ip router ospf 1 area 0.0.0.0
 no shutdown
```

- Configure BGP

```
router bgp 65535
router-id 10.1.1.1
 neighbor 30.1.1.1 remote-as 65535
   update-source loopback0
   address-family l2vpn evpn
     send-community both
     route-reflector-client
 neighbor 40.1.1.1 remote-as 65535
   update-source loopback0
   address-family l2vpn evpn
     send-community both
     route-reflector-client
```

- Spine (9504-B)

- Enable the EVPN control plane and the relevant protocols

```
feature telnet
feature nxapi
feature bash-shell
feature scp-server
nv overlay evpn
feature ospf
feature bgp
feature lldp
```

- Configure interfaces for Spine-leaf interconnect

```
interface Ethernet4/2
 ip address 192.168.4.42/24
 ip router ospf 1 area 0.0.0.0
 no shutdown
```

```
interface Ethernet4/3
 ip address 192.168.3.43/24
 ip router ospf 1 area 0.0.0.0
 no shutdown
```

- Configure Loopback for local VTEP IP, and BGP

```
interface loopback0
 ip address 20.1.1.1/32
```

```
ip router ospf 1 area 0.0.0.0
```

- Configure Loopback for Anycast RP

```
interface loopback1
ip address 100.1.1.1/32
ip router ospf 1 area 0.0.0.0
```

- Enable OSPF for underlay routing

```
router ospf 1
```

- Configure BGP

```
router bgp 65535
router-id 20.1.1.1
neighbor 30.1.1.1 remote-as 65535
update-source loopback0
address-family l2vpn evpn
send-community both
route-reflector-client
neighbor 40.1.1.1 remote-as 65535
update-source loopback0
address-family l2vpn evpn
send-community both
route-reflector-client
```

- Leaf (9396-A)

- Enable the EVPN control plane

```
nv overlay evpn
```

- Enable the relevant protocols

```
feature ospf
feature bgp
feature interface-vlan
```

- Enabling OSPF for underlay routing

```
router ospf 1
```

- Configure Loopback for local VTEP IP, and BGP

```
interface loopback0
ip address 30.1.1.1/32
ip router ospf 1 area 0.0.0.0
```

- Configure interfaces for Spine-leaf interconnect

```
interface Ethernet2/2
no switchport
ip address 192.168.1.22/24
ip router ospf 1 area 0.0.0.0
```

```

no shutdown

interface Ethernet2/3
no switchport
ip address 192.168.3.23/24
ip router ospf 1 area 0.0.0.0
no shutdown

```

- Create overlay VRF VLAN and configure vn-segment

```

vlan 101
vn-segment 900001

```

- Configure VRF overlay VLAN/SVI for the VRF

```

interface Vlan101
no shutdown
vrf member vxlan-900001

```

- Create VLAN and provide mapping to VXLAN

```

vlan 1001
vn-segment 2001001
vlan 1002
vn-segment 2001002

```

- Create VRF and configure VNI

```

vrf context vxlan-900001
vni 900001

rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn

```

- Create server facing SVI and enable distributed anycast-gateway

```

interface Vlan1001
no shutdown
vrf member vxlan-900001
ip address 4.1.1.1/24
ipv6 address 4:1:0:1::1/64
fabric forwarding mode anycast-gateway

interface Vlan1002
no shutdown
vrf member vxlan-900001
ip address 4.2.2.1/24
ipv6 address 4:2:0:1::1/64
fabric forwarding mode anycast-gateway

```

- Create the network virtualization endpoint (NVE) interface

```

interface nve1
no shutdown
source-interface loopback0

```



```

host-reachability protocol bgp
member vni 900001 associate-vrf
member vni 2001001
    suppress-arp
    mcast-group 225.4.0.1
member vni 2001002
    suppress-arp
    mcast-group 225.4.0.1

```

- Configure BGP

```

router bgp 65535
router-id 30.1.1.1
    neighbor 10.1.1.1 remote-as 65535
        update-source loopback0
        address-family l2vpn evpn
            send-community both
    neighbor 20.1.1.1 remote-as 65535
        update-source loopback0
        address-family l2vpn evpn
            send-community both
vrf vxlan-900001
    address-family ipv4 unicast
        advertise l2vpn evpn

evpn
    vni 2001001 l2
    vni 2001002 l2

rd auto
    route-target import auto
    route-target export auto

```

- Leaf (9396-B)

- Enable the EVPN control plane functionality and the relevant protocols

```

feature telnet
feature nxapi
feature bash-shell
feature scp-server
nv overlay evpn
feature ospf
feature bgp
feature interface-vlan
feature vn-segment-vlan-based
feature lldp
feature nv overlay

```

- Enable VxLAN with distributed anycast-gateway using BGP EVPN

```

fabric forwarding anycast-gateway-mac 0000.2222.3333

```

- Create overlay VRF VLAN and configure vn-segment

```

vlan 1-1002
vlan 101
    vn-segment 900001

```

- Create VLAN and provide mapping to VXLAN

```
vlan 1001
  vn-segment 2001001
vlan 1002
  vn-segment 2001002
```

- Create VRF and configure VNI

```
vrf context vxlan-900001
  vni 900001

rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  address-family ipv6 unicast
    route-target both auto
    route-target both auto evpn
```

- Configure internal control VLAN/SVI for the VRF

```
interface Vlan101
  no shutdown
  vrf member vxlan-900001
```

- Create server facing SVI and enable distributed anycast-gateway

```
interface Vlan1001
  no shutdown
  vrf member vxlan-900001
  ip address 4.1.1.1/24
  ipv6 address 4:1:0:1::1/64
  fabric forwarding mode anycast-gateway

interface Vlan1002
  no shutdown
  vrf member vxlan-900001
  ip address 4.2.2.1/24
  ipv6 address 4:2:0:1::1/64
  fabric forwarding mode anycast-gateway
```

- Create the network virtualization endpoint (NVE) interface

```
interface nve1
  no shutdown
  source-interface loopback0
  host-reachability protocol bgp
  member vni 900001 associate-vrf
  member vni 2001001
    suppress-arp
    mcast-group 225.4.0.1
  member vni 2001002
    suppress-arp
    mcast-group 225.4.0.1
```

- Configure interfaces for hosts/servers

```
interface Ethernet1/47
  switchport access vlan 1002
```

```
interface Ethernet1/48
  switchport access vlan 1001
```

- Configure interfaces for Spine-leaf interconnect

```
interface Ethernet2/1

interface Ethernet2/2
  no switchport
  ip address 192.168.4.22/24
  ip router ospf 1 area 0.0.0.0
  no shutdown

interface Ethernet2/3
  no switchport
  ip address 192.168.2.23/24
  ip router ospf 1 area 0.0.0.0
  no shutdown
```

- Configure Loopback for local VTEP IP, and BGP

```
interface loopback0
  ip address 40.1.1.1/32
  ip router ospf 1 area 0.0.0.0
```

- Enabling OSPF for underlay routing

```
router ospf 1
```

- Configure BGP

```
router bgp 65535
router-id 40.1.1.1
  neighbor 10.1.1.1 remote-as 65535
  update-source loopback0
  address-family l2vpn evpn
    send-community both
  neighbor 20.1.1.1 remote-as 65535
  update-source loopback0
  address-family l2vpn evpn
    send-community both
vrf vxlan-900001
  address-family ipv4 unicast
    advertise l2vpn evpn
evpn
  vni 2001001 l2
    rd auto
    route-target import auto
    route-target export auto
  vni 2001002 l2
    rd auto
    route-target import auto
    route-target export auto
evpn
  vni 2001001 l2
    rd auto
    route-target import auto
    route-target export auto
```

```
vni 2001002 12
rd auto
route-target import auto
route-target export auto
```

Example Show Commands

• show nve peers

```
9396-B# show nve peers
Interface Peer-IP Peer-State
-----
nve1      30.1.1.1      Up
```

• show nve vni

```
9396-B# 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	900001	n/a	Up	CP	L3	[vxlan-900001]	
nve1	2001001	225.4.0.1	Up	CP	L2	[1001]	SA
nve1	2001002	225.4.0.1	Up	CP	L2	[1002]	SA

• show vxlan interface

```
9396-B# show vxlan interface
Interface      Vlan    VPL Ifindex    LTL      HW VP
=====
Eth1/47        1002    0x4c07d22e     0x10000  5697
Eth1/48        1001    0x4c07d02f     0x10001  5698
```

• show bgp l2vpn evpn summary

```
leaf3# show bgp l2vpn evpn summary
BGP summary information for VRF default, address family L2VPN EVPN
BGP router identifier 40.0.0.4, local AS number 10
BGP table version is 60, L2VPN EVPN config peers 1, capable peers 1
21 network entries and 21 paths using 2088 bytes of memory
BGP attribute entries [8/1152], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [1/4]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
State/PfxRcd								
40.0.0.1	4	10	8570	8565	60	0	0	5d22h 6

• show bgp l2vpn evpn

```
leaf3# show bgp l2vpn evpn
BGP routing table information for VRF default, address family L2VPN EVPN
```

```

BGP table version is 60, local router ID is 40.0.0.4
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid,
>-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist,
I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 40.0.0.2:32868					
*>i[2]:[0]:[10001]:[48]:[0000.8816.b645]:[0]:[0.0.0.0]/216	40.0.0.2		100	0	i
*>i[2]:[0]:[10001]:[48]:[0011.0000.0034]:[0]:[0.0.0.0]/216	40.0.0.2		100	0	i

• show l2route evpn mac all

```

leaf3# show l2route evpn mac all
Topology  Mac Address  Prod  Next Hop (s)
-----
101      0000.8816.b645 BGP   40.0.0.2
101      0001.0000.0033 Local Ifindex 4362086
101      0001.0000.0035 Local Ifindex 4362086
101      0011.0000.0034 BGP   40.0.0.2

```

• show l2route evpn mac-ip all

```

leaf3# show l2route evpn mac-ip all
Topology ID Mac Address  Prod Host IP  Next Hop (s)
-----
101      0011.0000.0034 BGP  5.1.3.2     40.0.0.2
102      0011.0000.0034 BGP  5.1.3.2     40.0.0.2

```

