

Implement and Verify BGP-Only VxLAN EVPN on Catalyst 9000

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[BGP-only EVPN Feature of Use](#)

[BGP-only EVPN Comparisons and Considerations](#)

[EBGP Comparisons](#)

[Two-AS Model](#)

[Underlay BGP IPv4 Routing Consideration](#)

[Underlay BGP IPv4 Allowed AS IN](#)

[Underlay BGP IPv4 Maximum-paths](#)

[Overlay BGP EVPN Routing Consideration](#)

[Overlay BGP EVPN Allowed AS IN](#)

[Overlay BGP EVPND o not Change Next-Hop](#)

[Overlay BGP EVPNDisable RT Filter](#)

[Configure](#)

[Network Diagram](#)

[Configurations](#)

[Underlay BGP IPv4 Routing](#)

[Configure BGP IPv4 Routing](#)

[Configure BGP IPv4 Allowed AS In](#)

[Configure BGP Maximum-Paths](#)

[Underlay Multicast](#)

[Overlay BGP](#)

[Configure BGP L2VPN EVPN](#)

[Configure BGP EVPN Allowed AS In](#)

[Configure BGP EVPN do not Change Next-Hop](#)

[Configure BGP EVPN Disable RT Filter](#)

[VRF Configuration on Leaf](#)

[EVPN L2](#)

[EVPN L3](#)

[Verify](#)

[Verify BGP Sessions are Established](#)

[Verify EVPN L2](#)

[Verify EVPN L3](#)

[Related Information](#)

Introduction

This document describes how to implement and verify VXLAN Ethernet EVPN on Cisco Catalyst 9000 Series Switches with BGP only.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Border Gateway Protocol (BGP) Ethernet VPN (EVPN)
- Virtual Extensible LAN (VXLAN) Overlay
- Software Configuration Guide, Cisco IOS® XE

Components Used

The information in this document is based on these software and hardware versions:

- Catalyst 9600X
- Catalyst 9500X
- Cisco IOS XE 17.12 and later

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

Designing a next-generation campus network involves adopting modern technologies and architectures to meet the evolving demands of users, applications, and devices. VXLAN with BGP EVPN solution can provide a fabric-based architecture for simplicity, scalability, and ease of management. This document describes the BGP EVPN solution for users who prefer to use BGP for both IPv4 and EVPN routing for any reason.

BGP-only EVPN Feature of Use


VXLAN with BGP EVPN utilizes a spine-leaf architecture instead of the traditional 3-Tier network model. With a spine-leaf architecture, the spine acts as a high-speed conduit between access switches. The spine model enables a scale-out model where bandwidth between leaves can be increased by adding additional spines or endpoint capacity can be increased by adding more leaves.

For users who prefer to use BGP for both IPv4 and EVPN routing information, include these considerations:

- **Simplified Configuration:** With a single BGP session, the configuration and management of routing information become streamlined. There is no need to deploy and maintain separate routing protocols for IPv4 and EVPN, reducing complexity.
- **Unified Control Plane:** By utilizing BGP as the sole routing protocol, there is a unified control plane for both IPv4 and EVPN routes. This facilitates efficient route propagation, convergence, and route advertisement throughout the data center network.
- **Scalability:** BGP is well-suited for handling large-scale networks and offers robust scalability. Using a single BGP session for IPv4 and EVPN routing information ensures efficient scaling as the network grows, without the need for multiple routing protocol instances. At the same time, for large-scale fabric, BGP convergence time is shorter.

- **Interoperability:** BGP is a widely adopted industry-standard routing protocol. Utilizing BGP exclusively simplifies interoperability with various networking equipment and vendors, ensuring compatibility and seamless integration within the data center environment.

This topology shows a common C9K EVPN Single Fabric design.



C9K EVPN Single Fabric Design


BGP-only EVPN Comparisons and Considerations

EBGP Comparisons

For BGP-only design, the first issue to consider is whether to use Internal BGP (IBGP) or External BGP (EBGP). The case of using IBGP, which is common in VxLAN EVPN of traditional DC. Compared with using IBGP as Underlay, when using EBGP, Spine no longer needs to be configured as a route reflector, but functions as a traditional Router Server to exchange routes. So the prerequisite for this document is the case of using EBGP.


Option 1. Two-AS: Spine uses one AS, and Leaf and Border Leaf use another AS.

Two-AS Model



Two-AS Model

Option 2. Multi-AS: Spine, Leaf, and Border Leaf each use one AS.



Multi-AS Model


Comparing the two designs, a common problem is scalability, because for option 2, each time a spine or leaf is added, a new AS number needs to be added, which brings more complex configuration changes in the future, which is not conducive to expansion and maintenance. Therefore, this document uses option 1. for discussion.

Compared with using IBGP as Underlay, when using EBGP, Spine no longer needs to be configured as a route reflector, but functions as a traditional Router Server to exchange routes.

Underlay BGP IPv4 Routing Consideration

These are key points that need to be considered in the underlay plane.

Underlay BGP IPv4 Allowed AS IN



Underlay BGP IPv4 Allowed AS IN

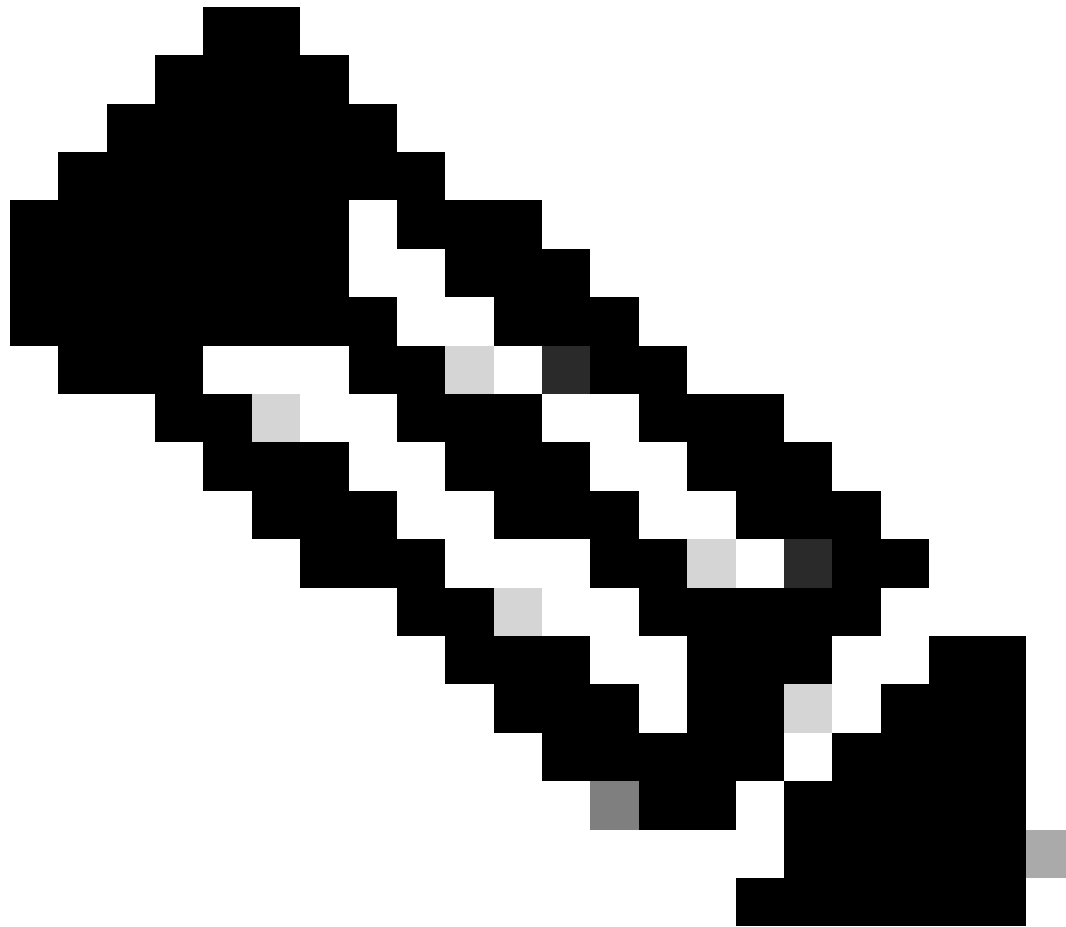
AS loop detection is done by scanning the full AS path (as specified in the AS_PATH attribute), and checking that the autonomous system number of the local system does not appear in the AS path.

According to the diagram above, the BGP AS Loop is formed - the same AS number in the as-path in this scenario:

- On Leaf and Border Leaf devices, the as-path is {#1, #2, #1}.
- On Spine devices, the as-path is {#2, #1, #2}.

To solve this problem, allow-as-in is configured in the BGP IPv4 address family, with the instructions outlined here:


- Allowed AS In to appear only once on all the Leaf and Border Leaf devices (**Leaf > Spine > Leaf**) as all the Leaf switches run in the same AS.
- Allowed AS In to appear only once on all the Spine devices (**Spine > BL > Spine**) or (**Spine > Leaf > Spine**) as all the Spine devices run in the same AS.



Note: When Single-Fabric is used with DGW, it is unlikely that routing is required from one spine to another. However, considering topology changes, such as super-spine, it is recommended to disable AS check on Spine devices also.

Underlay BGP IPv4 Maximum-paths

BGP chooses a route based on its criteria, and it is unlikely to appear 2 ECMP routes in the BGP table by default. To achieve ECMP for bandwidth optimization, **maximum-paths X** must be configured in the BGP IPv4 address family in all BGP running devices. Meanwhile, we suggest keeping the same link bandwidth between spine and leaf as a best practice.




Note: The maximum-paths depend on the design of the topology. With two spine switches, you can configure 'maximum-paths 2'.

Overlay BGP EVPN Routing Consideration

These key points need to be considered in the overlay plane.

Overlay BGP EVPN Allowed AS IN



Overlay BGP IPv4 Allowed AS IN


AS loop detection is done by scanning the full AS path (as specified in the AS_PATH attribute), and checking that the autonomous system number of the local system does not appear in the AS path.

According to the image, the BGP AS Loop is formed - the same AS number in the as-path in this scenario:

- On Leaf and Border Leaf devices, the as-path is {#1, #2, #1}
- On Spine devices, the as-path is {#2, #1, #2}

To solve this problem, allow-as-in must be configured in the BGP IPv4 address family, with the instructions outlined:

- Allowed AS In to appear only once on all the Leaf and Border Leaf devices (**Leaf > Spine > Leaf**) as all the Leaf switches run in the same AS.
- Allowed AS In to appear only once on all the Spine devices (**Spine > BL > Spine**) or (**Spine > Leaf > Spine**) as all the Spine devices run in the same AS.



Note: When Single-Fabric is used with DGW, it is unlikely that routing is required from one spine to another. However, considering topology changes, such as super-spine, it is recommended to disable AS check on Spine devices also.

Overlay BGP EVPN Do not Change Next-Hop

BGP changes the next-hop attribute of Network layer reachability information (NLRI) advertised from EBGp neighbor by default. Leaf/VXLAN Tunnel End Point (VTEP) uses its NVE source address as the next-hop attribute of the EVPN routes, and this address is used to determine the destination of the VXLAN tunnel (Network Virtual Interface/NVE peer). If Spine nodes change the next-hop, the VXLAN tunnel cannot be correctly established.

To solve this problem, these instructions are applied:

- On all Spine nodes, need to configure route-map with action next-hop unchanged

Overlay BGP EVPN Disable RT Filter

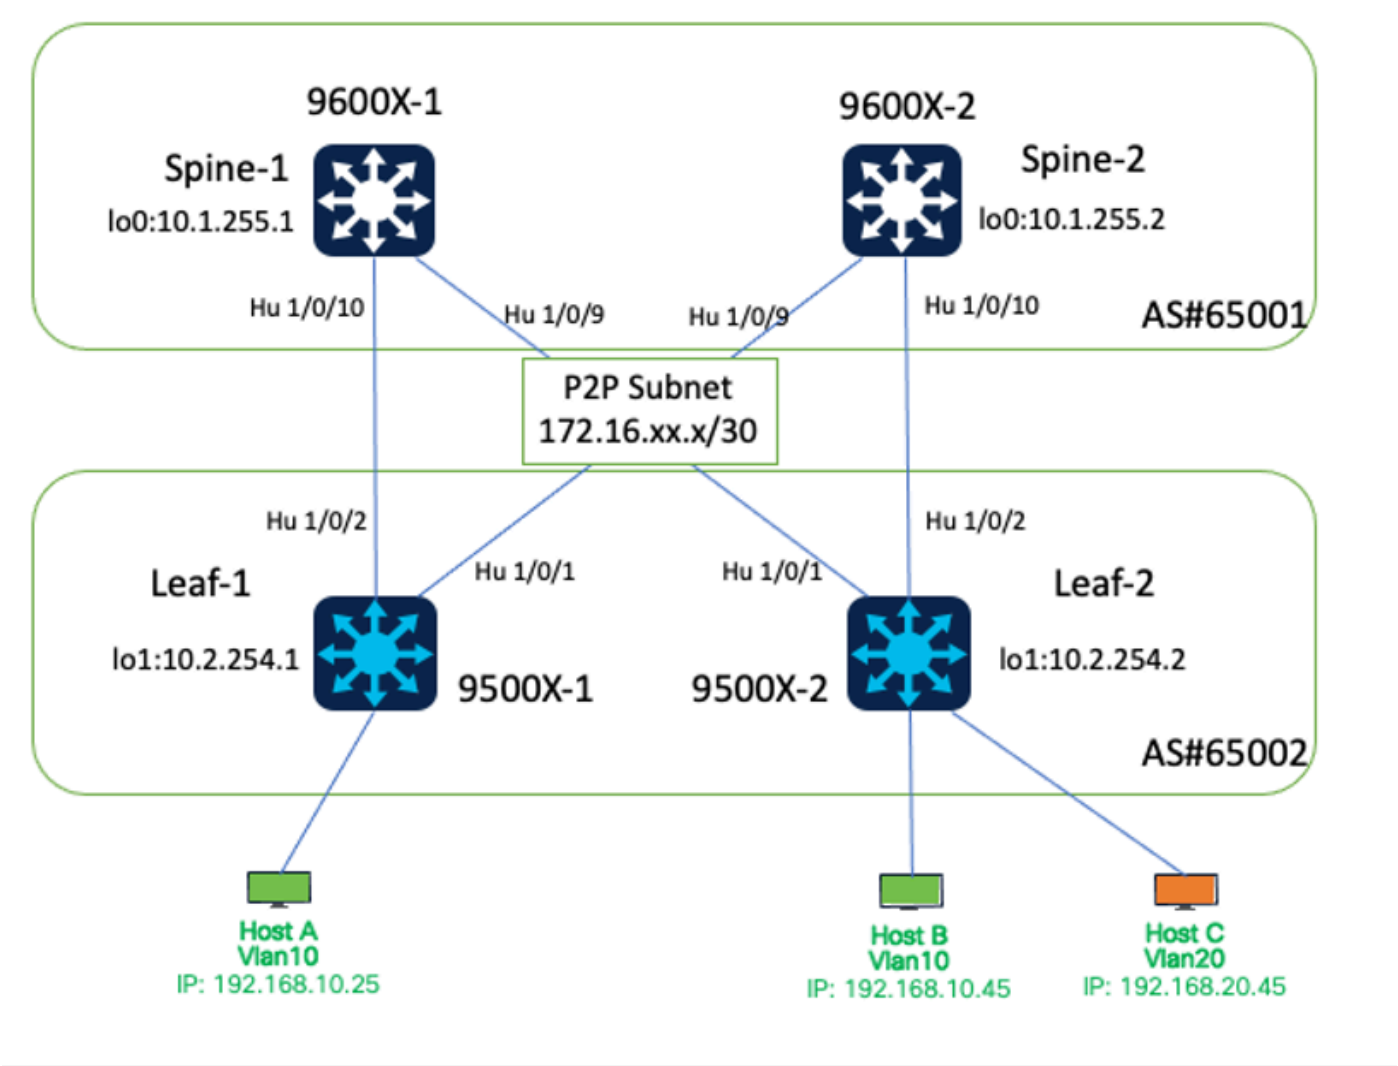
The EVPN routes from the Leaf devices are advertised with the Route Target (RT) community. Routers

without the corresponding RT configuration drop the routes with the RT community by default. Whereas all the spine devices have no Virtual Routing and Forwarding (VRF) configured. It means that the spine devices drop all the EVPN routes advertised from the Leaf devices by default.

To solve this problem, on all Spine nodes, the default route-target filter needs to be disabled.

Configure

Network Diagram




Network Diagram

Interface details for this lab environment are as follows.

Device Name	Software Version	Interface#	IP Address
Spine-1	IOS-XE 17.12.1	Hu 1/0/9	172.16.12.1/30
		Hu 1/0/10	172.16.11.1/30

		Lo 0	10.1.255.1/32
Spine-2	IOS-XE 17.12.1	Hu 1/0/9	172.16.21.1/30
		Hu 1/0/10	172.16.22.1/30
		Lo 0	10.1.255.2/32
Leaf-1	IOS-XE 17.12.1	Hu 1/0/1	172.16.21.2/30
		Hu 1/0/2	172.16.11.2/30
		Lo 1	10.2.254.1/32
Leaf-2	IOS-XE 17.12.1	Hu 1/0/1	172.16.12.2/30
		Hu 1/0/2	172.16.22.2/30
		Lo 1	10.2.254.2/32



Note: The IP Address assignment in this lab is for testing purposes only. Subnet Mask (that is, /30, /31) for Point-to-Point connections could be considered based on your actual design requirements.

Configurations

Underlay BGP IPv4 Routing

In this example, the physical interfaces are used to establish BGP connections.

- Configure BGP IPv4 Routing
- Configure BGP IPv4 Allowed AS In
- Configure BGP maximum-paths

Configure BGP IPv4 Routing

Configuration on Spine:

```
router bgp 65001
bgp log-neighbor-changes
bgp listen range 172.16.0.0/16 peer-group Leaf-Peers
no bgp default ipv4-unicast
neighbor Leaf-Peers peer-group
neighbor Leaf-Peers remote-as 65002
!
address-family ipv4
redistribute connected
neighbor Leaf-Peers activate
neighbor Leaf-Peers allowas-in 1
maximum-paths 2
exit-address-family
```

Configuration on Leaf-1:

```
router bgp 65002
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.11.1 remote-as 65001
neighbor 172.16.21.1 remote-as 65001
!
address-family ipv4
redistribute connected
neighbor 172.16.11.1 activate
neighbor 172.16.21.1 activate
exit-address-family
```

Configuration on Leaf-2:

```
router bgp 65002
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.12.1 remote-as 65001
neighbor 172.16.22.1 remote-as 65001
!
address-family ipv4
redistribute connected
neighbor 172.16.12.1 activate
neighbor 172.16.22.1 activate
exit-address-family
```

Configure BGP IPv4 Allowed AS In

Configuration on Spine:

```
router bgp 65001
address-family ipv4
neighbor Leaf-Peers allowas-in 1
```

Configuration on Leaf-1:

```
router bgp 65002
address-family ipv4
neighbor 172.16.11.1 allowas-in 1
neighbor 172.16.21.1 allowas-in 1
```

Configuration on Leaf-2:

```
router bgp 65002
address-family ipv4
neighbor 172.16.12.1 allowas-in 1
neighbor 172.16.22.1 allowas-in 1
```

Configure BGP Maximum-Paths

Configuration on Spine:

```
router bgp 65001
address-family ipv4
maximum-paths 2
```

Configuration on Leaf:

```
router bgp 65002
address-family ipv4
maximum-paths 2
```

Underlay Multicast

To enable Multicast Replication (MR) to handle Broadcast, Unknown Unicast and Link-Local Multicast (BUM) traffic, multicast routing is required on all Spine and Leaf devices. All Spine and Leaf connection interfaces and related loopbacks must have PIM enabled.

Example of underlay multicast on Spine 1:

```
ip multicast-routing
ip pim rp-address 10.1.255.1 //configure Spine loopback as RP
interface Loopback0
ip pim sparse-mode
interface HundredGigE1/0/9
ip pim sparse-mode
interface HundredGigE1/0/10
ip pim sparse-mode
```

Overlay BGP

- Configure BGP L2VPN EVPN
- Configure BGP EVPN Allowed AS In
- Configure BGP EVPN Do not Change Next-Hop
- Configure BGP EVPN Disable RT Filter

Configure BGP L2VPN EVPN

Configuration on Spine:

```
router bgp 65001
neighbor Leaf-Peers ebgp-multihop 255
address-family l2vpn evpn
neighbor Leaf-Peers activate
neighbor Leaf-Peers send-community both
```

Configuration on Leaf-1:

```
router bgp 65002
neighbor 172.16.11.1 ebgp-multihop 255
neighbor 172.16.21.1 ebgp-multihop 255
address-family l2vpn evpn
neighbor 172.16.11.1 activate
neighbor 172.16.11.1 send-community both
neighbor 172.16.21.1 activate
neighbor 172.16.21.1 send-community both
```

Configuration on Leaf-2:

```
router bgp 65002
neighbor 172.16.12.1 ebgp-multihop 255
neighbor 172.16.22.1 ebgp-multihop 255
address-family l2vpn evpn
neighbor 172.16.12.1 activate
neighbor 172.16.12.1 send-community both
neighbor 172.16.22.1 activate
neighbor 172.16.22.1 send-community both
```


Configure BGP EVPN Allowed AS In

Configuration on Leaf-1:

```
router bgp 65002
address-family l2vpn evpn
neighbor 172.16.11.1 allowas-in 1
neighbor 172.16.21.1 allowas-in 1
```

Configuration on Leaf-2:

```
router bgp 65002
address-family l2vpn evpn
neighbor 172.16.12.1 allowas-in 1
neighbor 172.16.22.1 allowas-in 1
```



Note: When Single-Fabric is used with DGW, it is unlikely that routing is required from one spine to another. However, considering topology changes, such as super-spine, it is recommended to disable AS check on Spine devices also.

Configure BGP EVPN do not Change Next-Hop

Configuration on Spine:

```
route-map BGP-NHU permit 10
set ip next-hop unchanged
!
```

```
router bgp 65001
address-family l2vpn evpn
neighbor Leaf-Peers route-map BGP-NHU out
```

Configure BGP EVPN Disable RT Filter

Configuration on Spine:

```
router bgp 65001
no bgp default route-target filter
```

VRF Configuration on Leaf

```
vrf definition S1-EVPN
rd 1:1
!
address-family ipv4
route-target export 1:1
route-target import 1:1
route-target export 1:1 stitching
route-target import 1:1 stitching
exit-address-family
router bgp 65002
address-family ipv4 vrf S1-EVPN
advertise l2vpn evpn
redistribute connected
maximum-paths 2
exit-address-family
```

EVPN L2

Enable L2VPN EVPN and multicast replication on Leaf:

```
l2vpn evpn
replication-type static
```

Create EVPN instances(EVI) on Leaf:

```
l2vpn evpn instance 10 vlan-based
encapsulation vxlan
l2vpn evpn instance 20 vlan-based
encapsulation vxlan
```

Create VLANs and VNI for user traffic on Leaf:

```
vlan configuration 10
member evpn-instance 10 vni 10010
vlan configuration 20
member evpn-instance 20 vni 10020
```

Create NVE interface and stitch VNI to mcast groups on Leaf.

```
interface nve1
no ip address
source-interface Loopback1
host-reachability protocol bgp
member vni 10010 mcast-group 225.0.0.10
member vni 10020 mcast-group 225.0.0.20
```

EVPN L3

Create VLAN for L3VNI on Leaf. EVI is not required for L3VNI.

```
vlan configuration 3000
member vni 33000
```

Configure SVI for L2VNI on Leaf.

```
interface Vlan10
mac-address 0010.0010.0010
vrf forwarding S1-EVPN
ip address 192.168.10.254 255.255.255.0
```

Configure SVI for L3VNI on Leaf. “no autostate” is configured to bring the SVI up when no active interface is assigned to that VLAN.

```
interface Vlan3000
vrf forwarding S1-EVPN
ip unnumbered Loopback1
no autostate
```

On Leaf, stitch L3VNI to the VRF under NVE configuration.

```
interface nve1
member vni 33000 vrf S1-EVPN
```

Verify

Verify BGP Sessions are Established

C9600X-SPINE-1#show ip bgp all summary

For address family: IPv4 Unicast

BGP router identifier 10.1.255.1, local AS number 65001

BGP table version is 23, main routing table version 23

12 network entries using 2976 bytes of memory

22 path entries using 2992 bytes of memory

2 multipath network entries and 4 multipath paths

4/3 BGP path/bestpath attribute entries using 1184 bytes of memory

3 BGP AS-PATH entries using 104 bytes of memory

8 BGP extended community entries using 400 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 7656 total bytes of memory

BGP activity 7259/7235 prefixes, 13926/13892 paths, scan interval 60 secs

12 networks peaked at 07:06:41 Dec 5 2023 UTC (2w1d ago)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
*172.16.11.2	4	65002	138	130	23	0	0	01:38:17	9
*172.16.12.2	4	65002	138	130	23	0	0	01:38:11	9

* Dynamically created based on a listen range command

Dynamically created neighbors: 2, Subnet ranges: 1

BGP peergroup Leaf-Peers listen range group members:

172.16.0.0/16

For address family: L2VPN E-VPN

BGP router identifier 10.1.255.1, local AS number 65001

BGP table version is 27, main routing table version 27

10 network entries using 3840 bytes of memory

12 path entries using 2784 bytes of memory

8/6 BGP path/bestpath attribute entries using 2368 bytes of memory

3 BGP AS-PATH entries using 104 bytes of memory

8 BGP extended community entries using 400 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 9496 total bytes of memory

BGP activity 7259/7235 prefixes, 13926/13892 paths, scan interval 60 secs

12 networks peaked at 07:38:03 Dec 6 2023 UTC (2w0d ago)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
*172.16.11.2	4	65002	138	130	27	0	0	01:38:17	6
*172.16.12.2	4	65002	138	130	27	0	0	01:38:11	6

* Dynamically created based on a listen range command

Dynamically created neighbors: 2, Subnet ranges: 1

BGP peergroup Leaf-Peers listen range group members:

172.16.0.0/16

Total dynamically created neighbors: 2/(100 max), Subnet ranges: 1

C9500X-LEAF-1#show ip bgp all summary

For address family: IPv4 Unicast

BGP router identifier 10.2.255.1, local AS number 65002

BGP table version is 19, main routing table version 19
 12 network entries using 2976 bytes of memory
 22 path entries using 2992 bytes of memory
 2 multipath network entries and 4 multipath paths
 4/3 BGP path/bestpath attribute entries using 1184 bytes of memory
 3 BGP AS-PATH entries using 104 bytes of memory
 8 BGP extended community entries using 384 bytes of memory
 0 BGP route-map cache entries using 0 bytes of memory
 0 BGP filter-list cache entries using 0 bytes of memory
 BGP using 7640 total bytes of memory
 BGP activity 577/545 prefixes, 4021/3975 paths, scan interval 60 secs
 12 networks peaked at 07:10:16 Dec 5 2023 UTC (1d18h ago)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.11.1	4	65001	2427	3100	19	0	0	20:39:49	9
172.16.21.1	4	65001	2430	3094	19	0	0	20:39:49	9

For address family: L2VPN E-VPN
 BGP router identifier 10.2.255.1, local AS number 65002
 BGP table version is 5371, main routing table version 5371
 16 network entries using 6144 bytes of memory
 20 path entries using 4640 bytes of memory
 9/9 BGP path/bestpath attribute entries using 2664 bytes of memory
 3 BGP AS-PATH entries using 104 bytes of memory
 8 BGP extended community entries using 384 bytes of memory
 0 BGP route-map cache entries using 0 bytes of memory
 0 BGP filter-list cache entries using 0 bytes of memory
 BGP using 13936 total bytes of memory
 BGP activity 577/545 prefixes, 4021/3975 paths, scan interval 60 secs
 16 networks peaked at 07:36:38 Dec 6 2023 UTC (18:16:58.620 ago)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.11.1	4	65001	2427	3100	5371	0	0	20:39:49	4
172.16.21.1	4	65001	2430	3094	5371	0	0	20:39:49	4

Initiate traffic between hosts, verify IP Multicast and PIM configuration, and mroute table.

Please note that on IOS-XE platform, (*, G) entry should always present, and (S, G) entry presents only

C9600X-SPINE-1#show ip mroute

IP Multicast Routing Table

<snip>

Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 t - LISP transit group

Timers: Uptime/Expires

Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.20), 16:51:00/stopped, RP 10.1.255.1, flags: SJCx

Incoming interface: HundredGigE1/0/2, RPF nbr 172.16.11.1

Outgoing interface list:

Tunnel0, Forward/Sparse-Dense, 16:51:00/00:02:58, flags:

(*, 225.0.0.10), 16:51:14/stopped, RP 10.1.255.1, flags: SJCx

Incoming interface: HundredGigE1/0/2, RPF nbr 172.16.11.1

Outgoing interface list:

Tunnel0, Forward/Sparse-Dense, 16:51:14/00:02:45, flags:

(10.2.254.1, 225.0.0.10), 00:00:01/00:02:57, flags: FTx

Incoming interface: Loopback1, RPF nbr 0.0.0.0, Registering
Outgoing interface list:
HundredGigE1/0/2, Forward/Sparse, 00:00:01/00:03:27, flags:

(*, 224.0.1.40), 1d18h/00:02:42, RP 10.1.255.1, flags: SJCL
Incoming interface: HundredGigE1/0/2, RPF nbr 172.16.11.1
Outgoing interface list:
Loopback0, Forward/Sparse, 1d18h/00:02:42, flags

Verify EVPN L2

```
C9500X-LEAF-1#show l2vpn evpn evi 10 detail
EVPN instance:      10 (VLAN Based)
RD:                  10.2.254.1:10 (auto)
Import-RTs:          65002:10
Export-RTs:           65002:10
<snip>
```

```
C9500X-LEAF-1#show nve peers
'M' - MAC entry download flag  'A' - Adjacency download flag
'4' - IPv4 flag  '6' - IPv6 flag
```

Interface	VNI	Type	Peer-IP	RMAC/Num_RT	eVNI	state	flags	UP time
nve1	33000	L3CP	10.2.254.2	242a.0412.0102	33000	UP	A/M/4	18:11:35
nve1	10010	L2CP	10.2.254.2	2	10010	UP	N/A	00:36:00
nve1	10020	L2CP	10.2.254.2	2	10020	UP	N/A	00:01:17

```
C9500X-LEAF-1#show bgp l2vpn evpn
BGP table version is 5475, local router ID is 10.2.254.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 10.2.254.1:10					
> [2][10.2.254.1:10][0][48][683B78FC8C9F][0][]/20	10.2.254.2	0	65001	65002	?
*> [2][10.2.254.1:10][0][48][683B78FC8C9F][32][192.168.10.45]/24	10.2.254.2	0	65001	65002	?

<snip>

```
C9500X-LEAF-1#show bgp l2vpn evpn detail [2][10.2.254.1:10][0][48][683B78FC8C9F][32][192.168.10.45]/24
BGP routing table entry for [2][10.2.254.1:10][0][48][683B78FC8C9F][32][192.168.10.45]/24, version 5371
Paths: (1 available, best #1, table evi_10)
Not advertised to any peer
Refresh Epoch 12
65001 65002, imported path from [2][10.2.254.2:10][0][48][683B78FC8C9F][32][192.168.10.45]/24 (global)
10.2.254.2 (via default) from 172.16.21.1 (10.1.255.2)
Origin incomplete, localpref 100, valid, external, best
EVPN ESI: 00000000000000000000, Label1 10010, Label2 33000
Extended Community: RT:1:1 RT:65002:10 ENCAP:8
Router MAC:242A.0412.0102
rx pathid: 0, tx pathid: 0x0
Updated on Dec 7 2023 01:52:33 UTC
```

C9500X-LEAF-1#show device-tracking database
<snip>

	Network Layer Address	Link Layer Address	Interface	vlan	prlv1	ag
ARP	192.168.20.25	3c13.cc01.a7df	Hu1/0/7	20	0005	3m
ARP	192.168.10.25	3c13.cc01.a7df	Hu1/0/7	10	0005	20

C9500X-LEAF-1#show l2vpn evpn mac ip

IP Address	EVI	VLAN	MAC Address	Next Hop(s)
192.168.10.25	10	10	3c13.cc01.a7df	Hu1/0/7:10
192.168.10.45	10	10	683b.78fc.8c9f	10.2.254.2

Verify EVPN L3

C9500X-LEAF-1#show nve peers
'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag

Interface	VNI	Type	Peer-IP	RMAC/Num_RT	eVNI	state	flags	UP time
nve1	33000	L3CP	10.2.254.2	242a.0412.0102	33000	UP	A/M/4	18:50:51
nve1	10010	L2CP	10.2.254.2	2	10010	UP	N/A	01:15:16
nve1	10020	L2CP	10.2.254.2	2	10020	UP	N/A	00:31:39

9500X-LEAF-1#sh bgp l2vpn evpn
BGP table version is 5523, local router ID is 10.2.255.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
x best-external, a additional-path, c RIB-compressed,
t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

Network	Next Hop	Metric	LocPrf	Weight	Path
<snip>					
Route Distinguisher: 1:1 (default for vrf S1-EVPN)					
*> [5][1:1][0][24][192.168.10.0]/17	0.0.0.0	0		32768	?
*> [5][1:1][0][24][192.168.20.0]/17	0.0.0.0	0		32768	?

C9500X-LEAF-1#sh ip ro vrf S1-EVPN

Routing Table: S1-EVPN
<snip>

192.168.10.0/24 is variably subnetted, 4 subnets, 2 masks
C 192.168.10.0/24 is directly connected, Vlan10
S 192.168.10.25/32 is directly connected, Vlan10
B 192.168.10.45/32 [20/0] via 10.2.254.2, 00:00:56, Vlan3000
L 192.168.10.254/32 is directly connected, Vlan10
192.168.20.0/24 is variably subnetted, 4 subnets, 2 masks
C 192.168.20.0/24 is directly connected, Vlan20
S 192.168.20.25/32 is directly connected, Vlan20

B 192.168.20.45/32 [20/0] via 10.2.254.2, 00:49:54, Vlan3000
L 192.168.20.254/32 is directly connected, Vlan20

Related Information

- [BGP EVPN VXLAN Configuration Guide, Cisco IOS XE Dublin 17.12.x:](#)
- [Cisco Technical Support & Downloads](#)