

# Troubleshoot EVPN/VxLAN in Multi-Site Setup

## Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Topology](#)

[Verify the Control Plane](#)

[Verify the Data Plane](#)

[Verify the Data Plane](#)

[Take ELAM Captures to Verify Port ASIC, Slice, and SrcId](#)

## Introduction

This document describes the approach to troubleshoot Ethernet VPN/Virtual Extensible LAN (EVPN/VxLAN) in a multi-site setup.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- Multiprotocol Label Switching (MPLS) Layer 3 VPN
- Multiprotocol-Border Gateway Protocol (MP-BGP)
- EVPN


### Components Used

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

All Site Leafs	N9K-C9336C-FX2	NXOS: 10.2(3)
S1_Spine1	N9K-C9364C	NXOS: 10.2(4)
S1_Spine2	N9K-C9364C	NXOS: 9.3(5)
S1_Border Gateway1, S2_Border Gateway2, S2_Border Gateway1	N9K-C9332C	NXOS: 9.3(9)
S1_Border Gateway2	N9K-C9332C	NXOS: 10.2(4)
Route Server	N9K-C9396PX	NXOS: 9.2(2)
Host-1	N3K-C3264C-E	NXOS: 9.3(5)
Host-2 and Host-3	N3K-C3264C-E	NXOS: 9.2(2)

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.

## Topology



### Topology

This document describes where the traffic is originated from DC-2 Host-3 (192.168.200.104/24), and then walks with the packets until the destination DC-1 Host-2 (10.2.3.4).

## Verify the Control Plane

In order to verify the control plane, enter these commands:

```
<#root>
HOST_3#
show ip int brief
□
```

```
10.100.100.2 100 0 300 100 65111 i
```

```
*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224  
10.100.100.2 100 0 300 100 i
```

```
*>i[5]:[0]:[0]:[32]:[10.100.100.2]/224  
10.100.100.2
```

S2-Leaf2#

```
show bgp l2vpn evpn vrf vrf_2
```

BGP routing table information for VRF default, address family L2VPN EVPN  
BGP table version is 4389, Local Router ID is 10.103.0.8  
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, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 10.103.0.8:5 (L3VNI 4000502)					
*>i[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.100.100.2		100	0 300 100	i
*>i[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.100.100.2		100	0 300 100 65111	i
*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.100.100.2		100	0 300 100	i
*>i[5]:[0]:[0]:[32]:[10.100.100.2]/224	10.100.100.2		100	0 300 100	i

S2-Leaf2#

S2-leaf3#

```
show bgp l2vpn evpn vrf vrf_2
```

BGP routing table information for VRF default, address family L2VPN EVPN  
BGP table version is 4196, Local Router ID is 10.103.0.9  
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, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 10.103.0.9:5 (L3VNI 4000502)					
*>i[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.100.100.2		100	0 300 100	i
*>i[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.100.100.2		100	0 300 100 65111	i
*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.100.100.2		100	0 300 100	i
*>i[5]:[0]:[0]:[32]:[10.100.100.2]/224	10.100.100.2		100	0 300 100	i

S2-Leaf4#

S2-Leaf4#

```
show bgp l2vpn evpn vrf vrf_2
```

BGP routing table information for VRF default, address family L2VPN EVPN  
BGP table version is 4381, Local Router ID is 10.102.0.10  
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, 2 - best2

Network Next Hop Metric LocPrf Weight Path  
Route Distinguisher: 10.102.0.10:5 (L3VNI 4000502)

\*>i[5]:[0]:[0]:[24]:[192.168.100.0]/224  
10.100.100.2 100 0 300 100 i

\*>i[5]:[0]:[0]:[32]:[10.2.3.4]/224  
10.100.100.2 100 0 300 100 65111 i

\*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224  
10.100.100.2 100 0 300 100 i

\*>i[5]:[0]:[0]:[32]:[10.100.100.2]/224  
10.100.100.2 100 0 300 100 i

S2-Leaf4#  
S2-Leaf4#

<#root>

S2-Spine1#

show bgp l2vpn evpn

BGP routing table information for VRF default, address family L2VPN EVPN  
BGP table version is 1235, Local Router ID is 10.103.0.3  
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, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 200:4000502					
* i[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.100.100.2		100	0 300 100	
*>i	10.100.100.2		100	0 300 100	i
* i[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.100.100.2		100	0 300 100	65111 i
*>i	10.100.100.2		100	0 300 100	65111 i
* i[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.100.100.2		100	0 300 100	i
*>i	10.100.100.2		100	0 300 100	i
* i[5]:[0]:[0]:[32]:[10.100.100.2]/224	10.100.100.2		100	0 300 100	i
*>i	10.100.100.2		100	0 300 100	i

<#root>

S2-BG1#

show ip int brie

```

IP Interface Status for VRF "default"(1)
Interface          IP Address      Interface Status
Lo0                10.31.1.3      protocol-up/link-up/admin-up
Lo1                10.135.135.135 protocol-up/link-up/admin-up
Lo100              10.100.100.2   protocol-up/link-up/admin-up

Eth1/1            192.168.17.12  protocol-up/link-up/admin-up
Eth1/3            10.150.152.1   protocol-up/link-up/admin-up
S2-BG1#

```

S2-BG1#

```
show ip route 10.2.3.4 vrf vrf_2
```

```

IP Route Table for VRF "vrf_2"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

```

```
10.2.3.4/32, ubest/mbest: 1/0
```

```
*via 10.100.100.1%default, [20/0], 04:09:46, bgp-200, external, tag 300, segid: 4000502 tunnelid: 0xa64
```

S2-BG1#

S2-BG1#

```
show bgp l2vpn evpn
```

```

BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 6206, Local Router ID is 10.31.1.3
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, 2 - best2

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 100:4000502					
*>e[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.100.100.1			0 300 100	i
*>e[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.100.100.1			0 300 100 65111	i
*>e[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.100.100.1			0 300 100	i
*>e[5]:[0]:[0]:[32]:[10.100.100.2]/224	10.100.100.1			0 300 100	i

<#root>

S2-BG2#

```
show ip int brief
```

```
IP Interface Status for VRF "default"(1)
```

Interface	IP Address	Interface Status
Lo0	10.31.1.4	protocol-up/link-up/admin-up
Lo1	10.136.136.136	protocol-up/link-up/admin-up
Lo100	10.100.100.2	protocol-up/link-up/admin-up
Eth1/1	192.168.18.12	protocol-up/link-up/admin-up
Eth1/3	10.150.153.1	protocol-up/link-up/admin-up
S2-BG2#		
S2-BG2#		
S2-BG2#		

```
show ip route 10.2.3.4 vrf vrf_2
```

```
IP Route Table for VRF "vrf_2"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>
```

```
10.2.3.4/32, ubest/mbest: 1/0
  *via 10.100.100.1%default, [20/0], 04:15:13, bgp-200, external, tag 300, segid: 4000502 tunnelid: 0
```

```
S2-BG2#
S2-BG2#
```

```
show bgp l2vpn evpn
```

```
BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 5455, Local Router ID is 10.31.1.4
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, 2 - best2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 100:4000502					
*>e[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.100.100.1			0 300 100	i
*>e[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.100.100.1			0 300 100 65111	i
*>e[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.100.100.1			0 300 100	i
*>e[5]:[0]:[0]:[32]:[10.100.100.2]/224	10.100.100.1			0 300 100	i

```
<#root>
```

```
Router_Server#
```

```
show ip int brief
```

```
IP Interface Status for VRF "default"(1)
Interface      IP Address      Interface Status
Lo0            10.32.1.1       protocol-up/link-up/admin-up
```

```

Eth2/1          10.150.150.2    protocol-up/link-up/admin-up
Eth2/2          10.150.151.2    protocol-up/link-up/admin-up
Eth2/4          10.150.152.2    protocol-up/link-up/admin-up
Eth2/5          10.150.153.2    protocol-up/link-up/admin-up
Router_Server#
Router_Server#

```

```
show ip route 10.100.100.1
```

```

IP Route Table for VRF "default"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

```

```

10.100.100.1/32, ubest/mbest: 2/0
  *via 10.150.150.1, [20/0], 4d22h, bgp-300, external, tag 100
  *via 10.150.151.1, [20/0], 4d22h, bgp-300, external, tag 100

```

```

Router_Server#
Router_Server#
Router_Server#

```

```
show ip route 10.100.100.2
```

```

IP Route Table for VRF "default"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

```

```

10.100.100.2/32, ubest/mbest: 2/0
  *via 10.150.152.1, [20/0], 3w5d, bgp-300, external, tag 200
  *via 10.150.153.1, [20/0], 3w5d, bgp-300, external, tag 200

```

```

Router_Server#
Router_Server#

```

```
show bgp l2vpn evpn
```

```

BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 4574, Local Router ID is 10.32.1.1
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, 2 - best2

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 200:4000100					
* e[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272	10.100.100.2	2000		0	200 i
*>e	10.100.100.2	2000		0	200 i
Route Distinguisher: 100:4000502					
*>e[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.100.100.1	2000		0	100 i
* e	10.100.100.1	2000		0	100 i
* e[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.100.100.1	2000		0	100 65111 i

```

*>e          10.100.100.1          2000          0 100 65111 i

*>e[5]:[0]:[0]:[32]:[10.100.100.1]/224
          10.100.100.1          2000          0 100 i
* e          10.100.100.1          2000          0 100 i
*>e[5]:[0]:[0]:[32]:[10.100.100.2]/224
>          10.100.100.1          2000          0 100 i
* e          10.100.100.1          2000          0 100 i

```

<#root>

S1\_B2#  
S1\_B2#

show ip int brie

```

IP Interface Status for VRF "default"(1)
Interface      IP Address      Interface Status
Lo0            10.31.1.2      protocol-up/link-up/admin-up
Lo1            10.134.134.134 protocol-up/link-up/admin-up
Lo100         10.100.100.1   protocol-up/link-up/admin-up
Eth1/1        192.168.16.12  protocol-up/link-up/admin-up
Eth1/3        10.150.151.1   protocol-up/link-up/admin-up
S1_B2#
S1_B2#

```

sho ip route 192.168.100.103 vrf vrf\_2

```

IP Route Table for VRF "vrf_2"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

```

```

192.168.100.103/32, ubest/mbest: 1/0
  *via 10.100.100.2%default, [20/0], 4d23h, bgp-100, external, tag 300, segid: 4000502 tunnelid: 0xa64

```

S1\_B2#  
S1\_B2#

show ip route 10.2.3.4 vrf vrf\_2

```

IP Route Table for VRF "vrf_2"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

```

```

10.2.3.4/32, ubest/mbest: 1/0
  *via 10.102.1.10%default, [200/0], 05:04:19, bgp-100, internal, tag 65111, segid: 4000502 tunnelid:

```

S1\_B2#  
S1\_B2#  
S1\_B2#

show bgp l2vpn evpn



BGP routing table information for VRF default, address family L2VPN EVPN  
 BGP table version is 5449, Local Router ID is 10.31.1.2  
 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, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 200:4000100					
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.df3b]:[32]:[192.168.100.104]/272	10.100.100.2			0 300 200	i
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272	10.100.100.2			0 300 200	i

<#root>

Route Distinguisher: 200:4000200					
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.df3b]:[32]:[192.168.100.104]/272	10.100.100.2			0 300 200	i
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272	10.100.100.2			0 300 200	i

Route Distinguisher: 10.102.0.9:5					
*>i[2]:[0]:[0]:[48]:[cc7f.76fa.118f]:[0]:[0.0.0.0]/216	10.202.202.202		100	0	i
*>i[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.102.1.10		100	0	i
*>i[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.102.1.10		100	0 65111	i

*>i[5]:[0]:[0]:[32]:[10.100.100.2]/224	10.102.1.10		100	0	i
--	-------------	--	-----	---	---

Route Distinguisher: 10.102.0.10:5					
*>i[2]:[0]:[0]:[48]:[cc7f.76c6.a673]:[0]:[0.0.0.0]/216	10.202.202.202		100	0	i
*>i[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.102.1.6		100	0	i
*>i[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.102.1.6		100	0 65111	i
*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.102.1.6		100	0	i

Route Distinguisher: 10.31.1.2:5 (L3VNI 4000502)					
*>l[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.134.134.134		100	0	i
*>l[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.134.134.134		100	0 65111	i
*>l[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.134.134.134		100	0	i
*>l[5]:[0]:[0]:[32]:[10.100.100.2]/224	10.134.134.134		100	0	i

S1\_B2#

<#root>

S1-Bg1#

show ip int brie

IP Interface Status for VRF "default"(1)

Interface	IP Address	Interface Status
Lo0	10.31.1.1	protocol-up/link-up/admin-up
Lo1	10.133.133.133	protocol-up/link-up/admin-up
Lo100	10.100.100.1	protocol-up/link-up/admin-up
Eth1/1	192.168.15.12	protocol-up/link-up/admin-up
Eth1/3	10.150.150.1	protocol-up/link-up/admin-up

S1-Bg1#

S1-Bg1#

show ip route 10.100.100.2 vrf vrf\_2

IP Route Table for VRF "vrf\_2"

'\*' denotes best ucast next-hop  
'\*\*' denotes best mcast next-hop  
'[x/y]' denotes [preference/metric]  
'%<string>' in via output denotes VRF <string>

10.100.100.2/32, ubest/mbest: 1/0

\*via 10.102.1.10%default, [200/0], 4d23h, bgp-100, internal, tag 100, segid: 4000502 tunnelid: 0xa66

S1-Bg1#

S1-Bg1#

show ip route 192.168.100.103 vrf vrf\_2

IP Route Table for VRF "vrf\_2"

'\*' denotes best ucast next-hop  
'\*\*' denotes best mcast next-hop  
'[x/y]' denotes [preference/metric]  
'%<string>' in via output denotes VRF <string>

192.168.100.103/32, ubest/mbest: 1/0

\*via 10.100.100.2%default, [20/0], 4d23h, bgp-100, external, tag 300, segid: 4000502 tunnelid: 0xa66

S1-Bg1#

S1-Bg1#

show ip route 10.2.3.4 vrf vrf\_2

IP Route Table for VRF "vrf\_2"

'\*' denotes best ucast next-hop  
'\*\*' denotes best mcast next-hop  
'[x/y]' denotes [preference/metric]  
'%<string>' in via output denotes VRF <string>

10.2.3.4/32, ubest/mbest: 1/0

\*via 10.102.1.10%default, [200/0], 05:21:41, bgp-100, internal, tag 65111, segid: 4000502 tunnelid:

S1-Bg1#

S1-Bg1#

show bgp l2vpn evpn

BGP routing table information for VRF default, address family L2VPN EVPN  
BGP table version is 6654, Local Router ID is 10.31.1.1  
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, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 200:4000100					
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.df3b]:[32]:[192.168.100.104]/272	10.100.100.2			0 300 200	i
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272	10.100.100.2			0 300 200	i
Route Distinguisher: 200:4000200					
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.df3b]:[32]:[192.168.200.104]/272	10.100.100.2			0 300 200	i
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.200.103]/272	10.100.100.2			0 300 200	i
Route Distinguisher: 10.31.1.1:32867 (L2VNI 4000100)					
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272	10.100.100.2			0 300 200	i
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ff09]:[32]:[192.168.100.102]/272	10.202.202.202		100	0	i
* i	10.202.202.202		100	0	i
*>i[2]:[0]:[0]:[48]:[6c8b.d3ff.00a7]:[32]:[192.168.100.100]/272	10.201.201.201		100	0	i
* i	10.201.201.201		100	0	i
*>e[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.200.103]/272	10.100.100.2			0 300 200	i
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ff09]:[32]:[192.168.200.102]/272	10.202.202.202		100	0	i
* i	10.202.202.202		100	0	i
*>i[2]:[0]:[0]:[48]:[6c8b.d3ff.00a7]:[32]:[192.168.200.100]/272	10.201.201.201		100	0	i
* i	10.201.201.201		100	0	i
Route Distinguisher: 10.102.0.10:5					
*>i[2]:[0]:[0]:[48]:[cc7f.76c6.a673]:[0]:[0.0.0.0]/216	10.202.202.202		100	0	i
*>i[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.102.1.6		100	0	i
*>i[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.102.1.6		100	0	65111 i
*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.102.1.6		100	0	i

Route Distinguisher: 10.31.1.1:5 (L3VNI 4000502)

\*>l[5]:[0]:[0]:[24]:[192.168.100.0]/224

```

                10.133.133.133                100          0 i
*>l[5]:[0]:[0]:[32]:[10.2.3.4]/224
                10.133.133.133                100          0 65111 i
*>l[5]:[0]:[0]:[32]:[10.100.100.1]/224
                10.133.133.133                100          0 i
*>l[5]:[0]:[0]:[32]:[10.100.100.2]/224
                10.133.133.133                100          0 i
S1-Bg1#

```

<#root>

S1-Leaf1#

show ip int brief

```

IP Interface Status for VRF "default"(1)
Interface          IP Address      Interface Status
Lo0                10.102.0.5     protocol-up/link-up/admin-up
Lo1                10.102.1.8     protocol-up/link-up/admin-up
Eth1/2            192.168.17.12  protocol-up/link-up/admin-up
S1-Leaf1#

```

S1-Leaf1#

show bgp l2vpn evpn vrf vrf\_2

```

BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 918, Local Router ID is 10.102.0.5
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, 2 - best2

```

```

      Network          Next Hop          Metric    LocPrf    Weight Path
Route Distinguisher: 10.102.0.5:5 (L3VNI 4000502)
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272
                10.100.100.1                100          0 300 200 i
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.200.103]/272
                10.100.100.1                100          0 300 200 i
*>i[5]:[0]:[0]:[24]:[192.168.100.0]/224
                10.102.1.10            100          0 i
* i                10.102.1.6                100          0 i
*>i[5]:[0]:[0]:[32]:[10.2.3.4]/224
                10.102.1.10            100          0 65111 i
* i                10.102.1.6                100          0 65111 i

*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224
                10.102.1.6                100          0 i
*>i[5]:[0]:[0]:[32]:[10.100.100.2]/224
                10.102.1.10            100          0 i
S1-Leaf1#

```

S1-Leaf2#

show ip int brie

```

IP Interface Status for VRF "default"(1)
Interface          IP Address      Interface Status

```

```

Lo0          10.102.0.8      protocol-up/link-up/admin-up
Lo1          10.102.1.9      protocol-up/link-up/admin-up
Eth1/2       192.168.18.12     protocol-up/link-up/admin-up
S1-Leaf2#
S1-Leaf2#
S1-Leaf2#

```

```
show bgp l2vpn evpn vrf vrf_2
```

```

BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 680, Local Router ID is 10.102.0.8
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, 2 - best2

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 10.102.0.8:5 (L3VNI 4000502)					
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272	10.100.100.1		100	0 300 200	i
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.200.103]/272	10.100.100.1		100	0 300 200	i
*>i[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.102.1.10	100	0		i
* i	10.102.1.6		100	0	i
* i[5]:[0]:[0]:[32]:[10.2.3.4]/224	10.102.1.6		100	0 65111	i
*>i	10.102.1.10		100	0 65111	i
*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224	10.102.1.6		100	0	i
*>i[5]:[0]:[0]:[32]:[10.100.100.2]/224	10.102.1.10		100	0	i

```

S1-Leaf3#
S1-Leaf3#

```

```
show ip int brie
```

```

IP Interface Status for VRF "default"(1)
Interface      IP Address      Interface Status
Lo0            10.102.0.9     protocol-up/link-up/admin-up
Lo1            10.102.1.10    protocol-up/link-up/admin-up
Eth1/2         192.168.19.12  protocol-up/link-up/admin-up

```

```

S1-Leaf3#
S1-Leaf3#
S1-Leaf3#
S1-Leaf3#

```

```
show bgp l2vpn evpn vrf vrf_2
```

```

BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 5431, Local Router ID is 10.102.0.9
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, 2 - best2

```

Network	Next Hop	Metric	LocPrf	Weight	Path
---------	----------	--------	--------	--------	------

```

Route Distinguisher: 10.102.0.9:5 (L3VNI 4000502)
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272
    10.100.100.1          100          0 300 200 i
* i[5]:[0]:[0]:[24]:[192.168.100.0]/224
    10.102.1.6           100          0 i
*>l          10.102.1.10        100          32768 i
* i[5]:[0]:[0]:[32]:[10.2.3.4]/224
    10.102.1.6           100          0 65111 i
*>l          10.102.1.10        100          0 65111 i
*>i[5]:[0]:[0]:[32]:[10.100.100.1]/224
    10.102.1.6           100          0 i
*>l[5]:[0]:[0]:[32]:[10.100.100.2]/224
    10.102.1.10         100          32768 i

```

S1-Leaf3#

S1\_Leaf4#

S1\_Leaf4#

show ip int brief

IP Interface Status for VRF "default"(1)

Interface	IP Address	Interface Status
Lo0	10.102.0.10	protocol-up/link-up/admin-up
Lo1	10.102.1.6	protocol-up/link-up/admin-up
Eth1/2	192.168.20.12	protocol-up/link-up/admin-up

S1\_Leaf4#

S1\_Leaf4#

S1\_Leaf4#

show bgp l2vpn evpn vrf vrf\_2

BGP routing table information for VRF default, address family L2VPN EVPN

BGP table version is 5118, Local Router ID is 10.102.0.10

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, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 10.102.0.10:5 (L3VNI 4000502)					
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.100.103]/272	10.100.100.1		100	0	300 200 i
*>i[2]:[0]:[0]:[48]:[6c8b.d3fe.ecb5]:[32]:[192.168.200.103]/272	10.100.100.1		100	0	300 200 i
*>i[2]:[0]:[0]:[48]:[6c8b.d3ff.00a7]:[32]:[192.168.100.100]/272	10.201.201.201		100	0	i
* i	10.201.201.201		100	0	i
* i[5]:[0]:[0]:[24]:[192.168.100.0]/224	10.102.1.10		100	0	i
*>l	10.102.1.6		100	32768	i
*>l[5]:[0]:[0]:[32]:[10.2.3.4]/224					

```

10.102.1.6                                0 65111 i

* i          10.102.1.10                    100      0 65111 i

*>l[5]:[0]:[0]:[32]:[10.100.100.1]/224
          10.102.1.6                        100      32768 i
*>i[5]:[0]:[0]:[32]:[10.100.100.2]/224
          10.102.1.10                      100      0 i
S1_Leaf4#

```

## Verify the Data Plane

The data plan verification is tested on multiple devices in order to understand different packet capture methods and variants.

Ping the external router loopback 100 "10.2.3.4" from the source IP address 192.168.100.103 on Host-3.

```

<#root>

HOST_3#
HOST_3#

ping 10.2.3.4 source 192.168.100.103

PING 10.2.3.4 (10.2.3.4) from 192.168.100.103: 56 data bytes
64 bytes from 10.2.3.4: icmp_seq=0 ttl=250 time=1.153 ms
64 bytes from 10.2.3.4: icmp_seq=1 ttl=250 time=0.569 ms
64 bytes from 10.2.3.4: icmp_seq=2 ttl=250 time=0.562 ms
64 bytes from 10.2.3.4: icmp_seq=3 ttl=250 time=0.525 ms
64 bytes from 10.2.3.4: icmp_seq=4 ttl=250 time=0.527 ms
--- 10.2.3.4 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.525/0.667/1.153 ms
HOST_3#

```

Ethalyzer is taken on Site 2 Leaf-1 and Leaf-2 in order to confirm which leaf receives/forwards traffic for external router loopback 10.2.3.4 reachability.

```

<#root>

S2-Leaf1(config-monitor)#

sho clock

Warning: No NTP peer/server configured. Time may be out of sync.
07:11:37.455 UTC Tue Feb 21 2023
Time source is NTP
S2-Leaf1(config-monitor)#
S2-Leaf1(config-monitor)#

show run section monitor

```

```
show running-config | section monitor
icam monitor scale
monitor session 1
  source interface port-channel100 both
  destination interface sup-eth0
  no shut
S2-Leaf1(config-monitor)#
S2-Leaf2(config-monitor)#
S2-Leaf2(config-monitor)#

ethanalyzer local interface inband display-filter "ip.addr==10.2.3.4 && ip.addr==192.168.100.103 && icmp"
```

```
Capturing on 'ps-inb'
1385 2023-02-21 07:10:46.424195144 192.168.100.103 â†' 10.2.3.4 ICMP 102 Echo (ping) request id=0xc
1386 2023-02-21 07:10:46.424818423 10.2.3.4 â†' 192.168.100.103 ICMP 98 Echo (ping) reply id=0xc
1387 2023-02-21 07:10:46.425263621 192.168.100.103 â†' 10.2.3.4 ICMP 102 Echo (ping) request id=0xc
1388 2023-02-21 07:10:46.425486046 10.2.3.4 â†' 192.168.100.103 ICMP 98 Echo (ping) reply id=0xc
1389 2023-02-21 07:10:46.425856150 192.168.100.103 â†' 10.2.3.4 ICMP 102 Echo (ping) request id=0xc
1390 2023-02-21 07:10:46.426095692 10.2.3.4 â†' 192.168.100.103 ICMP 98 Echo (ping) reply id=0xc
1391 2023-02-21 07:10:46.426438174 192.168.100.103 â†' 10.2.3.4 ICMP 102 Echo (ping) request id=0xc
1392 2023-02-21 07:10:46.426642605 10.2.3.4 â†' 192.168.100.103 ICMP 98 Echo (ping) reply id=0xc
1393 2023-02-21 07:10:46.427004108 192.168.100.103 â†' 10.2.3.4 ICMP 102 Echo (ping) request id=0xc
1394 2023-02-21 07:10:46.427210984 10.2.3.4 â†' 192.168.100.103 ICMP 98 Echo (ping) reply id=0xc
10
S2-Leaf2(config-monitor)#
S2-Leaf2(config-monitor)#
```

```
sho clock
```

```
Warning: No NTP peer/server configured. Time may be out of sync.
07:12:31.069 UTC Tue Feb 21 2023
Time source is NTP
S2-Leaf2(config-monitor)#
```

The CLI output confirmed Site 2 Leaf-2 receives and forwards the Internet Control Message Protocol (ICMP) request for the external router 10.2.3.4.

The next CLI example confirms Site 1 verifies which leaf forwards packets towards destination 10.2.3.4.

```
<#root>
```

```
S1-Leaf3(config-monitor)#
S1-Leaf3(config-monitor)#

ethanalyzer local interface inband display-filter "ip.addr==10.2.3.4 && ip.addr==192.168.100.103 && icmp"
```

```
Capturing on 'ps-inb'
253 2023-02-21 07:10:50.379741403 192.168.100.103 â†' 10.2.3.4 ICMP 98 Echo (ping) request id=0xc
254 2023-02-21 07:10:50.380357311 10.2.3.4 â†' 192.168.100.103 ICMP 102 Echo (ping) reply id=0xc
255 2023-02-21 07:10:50.380810012 192.168.100.103 â†' 10.2.3.4 ICMP 98 Echo (ping) request id=0xc
256 2023-02-21 07:10:50.381025676 10.2.3.4 â†' 192.168.100.103 ICMP 102 Echo (ping) reply id=0xc
257 2023-02-21 07:10:50.381401968 192.168.100.103 â†' 10.2.3.4 ICMP 98 Echo (ping) request id=0xc
258 2023-02-21 07:10:50.381631838 10.2.3.4 â†' 192.168.100.103 ICMP 102 Echo (ping) reply id=0xc
259 2023-02-21 07:10:50.381984272 192.168.100.103 â†' 10.2.3.4 ICMP 98 Echo (ping) request id=0xc
260 2023-02-21 07:10:50.382176820 10.2.3.4 â†' 192.168.100.103 ICMP 102 Echo (ping) reply id=0xc
261 2023-02-21 07:10:50.382549820 192.168.100.103 â†' 10.2.3.4 ICMP 98 Echo (ping) request id=0xc
```



262 2023-02-21 07:10:50.382746640 10.2.3.4 â†’ 192.168.100.103 ICMP 102 Echo (ping) reply id=0x

S1-Leaf3(config-monitor)#

sho clock

Warning: No NTP peer/server configured. Time may be out of sync.

07:11:22.514 UTC Tue Feb 21 2023

Time source is NTP

S1-Leaf3(config-monitor)#

S1-Leaf3(config-monitor)#

show run section monitor

show running-config | section monitor

monitor session 1

source interface port-channel2 both

destination interface sup-eth0

no shut

S1-Leaf3(config-monitor)#

S1-Leaf3(config-monitor)#

show moni sess 1

session 1

-----

type : local

state : up

acl-name : acl-name not specified

source intf :

rx : Po2

tx : Po2

both : Po2

source VLANs :

rx :

tx :

both :

filter VLANs : filter not specified

source fwd drops :

destination ports : sup-eth0

source VSANs :

rx :

S1-Leaf3(config-monitor)#

S1-Leaf4(config-monitor)#

ethalyzer local interface inband display-filter "ip.addr==192.168.100.103" limit-captured-frames 0

Capturing on 'ps-inb'

S1-Leaf4(config-monitor)#

S1-Leaf4(config-monitor)#

sho clock

Warning: No NTP peer/server configured. Time may be out of sync.

07:11:15.187 UTC Tue Feb 21 2023

Time source is NTP

S1-Leaf4(config-monitor)#

The customer responds that they face connectivity issues from Host-3 to the external router. The customer wants to confirm everything is fine in the VXLAN fabric and needs confirmation that our leaf forwards traffic toward the external router. The steps to troubleshoot this issue are:

1. Initiate a ping toward the external router and confirm if the IP address 10.2.3.4 is reachable or not.
2. Take Embedded Logic Analyzer Module (ELAM) captures on both S1-Leaf3 and S1-Leaf4 in order to see if it gets triggered (based on the topology and traffic flow).
3. With the ELAM capture, confirm that the packet is forwarded out of the interface and points to the external router.
4. Site 2- With the ethanalyzer we can see the ICMP request and reply. If there is not a reply, the problem is on the remote side .
5. If 10.2.3.4 is reachable from Host-4 and Host-3 has issues, it could be a host specific issue. Check the Access Control List (ACL), Cyclic Redundancy Check (CRC) errors, and hashing link.

```
<#root>
```

```
HOST_3#
```

```
ping 10.2.3.4 source 192.168.100.103
```

```
PING 10.2.3.4 (10.2.3.4) from 192.168.100.103: 56 data bytes
Request 0 timed out
Request 1 timed out
Request 2 timed out
Request 3 timed out
Request 4 timed out
--- 10.2.3.4 ping statistics ---
5 packets transmitted, 0 packets received, 100.00% packet loss
HOST_3#
```

```
Host4#
```

```
ping 10.2.3.4 source 192.168.100.104
```

```
PING 10.2.3.4 (10.2.3.4) from 192.168.100.104: 56 data bytes
64 bytes from 10.2.3.4: icmp_seq=0 ttl=250 time=1.266 ms
64 bytes from 10.2.3.4: icmp_seq=1 ttl=250 time=0.62 m
64 bytes from 10.2.3.4: icmp_seq=2 ttl=250 time=0.603 ms
64 bytes from 10.2.3.4: icmp_seq=3 ttl=250 time=0.474 ms
64 bytes from 10.2.3.4: icmp_seq=4 ttl=250 time=0.457 ms
--- 10.2.3.4 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.457/0.684/1.266 ms
```

## Verify the Data Plane

### Take ELAM Captures to Verify Port ASIC, Slice, and SrcId

```
<#root>
```

```
show hardware internal tah interface <interface id>
```



```
show system internal ethpm info interface ethernet 1/2 | grep slice
```

```
IF_STATIC_INFO: port_name=Ethernet1/2,if_index:0x1a000200,ltl=6140,slot=0, nxos_port=4,  
dmod=1,dpid=76,unit=0,queue=65535,xbar_unitbmp=0x0,ns_pid=255,slice_num=1,port_on_slice=4,src_id=8
```

```
S1_Leaf4(TAH-elam)#
```

```
debug platform internal tah elam asic 0
```

```
S1_Leaf4(TAH-elam)#
```

```
trigger init asic 0 slice 1 in-select 7 out-select 0 use-src-id 8
```

```
Slot 1: param values: asic 0, slice 1, lu-a2d 1, in-select 7, out-select 0, src_id 8
```

```
S1_Leaf4(TAH-elam-insel7)#
```

```
set inner ipv4 src_ip 192.168.100.103
```

```
S1_Leaf4(TAH-elam-insel7)#
```

```
start
```

```
S1_Leaf4(TAH-elam-insel7)#
```

```
report
```

```
ELAM not triggered yet on slot - 1, asic - 0, slice - 1
```

```
S1_Leaf4(TAH-elam-insel7)#
```

The conclusion from the ELAM output is the leaf forwards traffic to the external router, but there is no response from the external router. Therefore, check with the external router team about the ICMP response.