

排除多站点环境中的EVPN/VxLAN故障

目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[拓扑](#)

[控制平面验证](#)

[检验Source-Leaf中的ARP表](#)

[DC-1 Leaf1到Spine1的路由验证](#)

[DC-1 Leaf1到Spine2的路由验证](#)

[DC-1 BGW-1](#)

[DC-2 BGW-2](#)

[DC-2主干到枝叶路由通告](#)

[DC-2枝叶-1验证](#)

[数据平面故障排除](#)

[对Source-Host4的基本检查](#)

[目标枝叶1检查](#)

[配置监控器会话](#)

[第一跳DC-2、枝叶1的基本检查](#)

[监控源枝叶上的会话](#)

[监控目标枝叶上的会话](#)

[在源枝叶上捕获](#)

[对最后一跳DC-1、枝叶-2的基本检查](#)

简介

本文档介绍如何在多站点环境中排除以太网VPN/虚拟可扩展局域网(EVPN/VxLAN)故障。

先决条件

要求

Cisco 建议您了解以下主题：

- 多协议标签交换(MPLS)第3层VPN
- 多协议边界网关协议(MP-BGP)
- EVPN

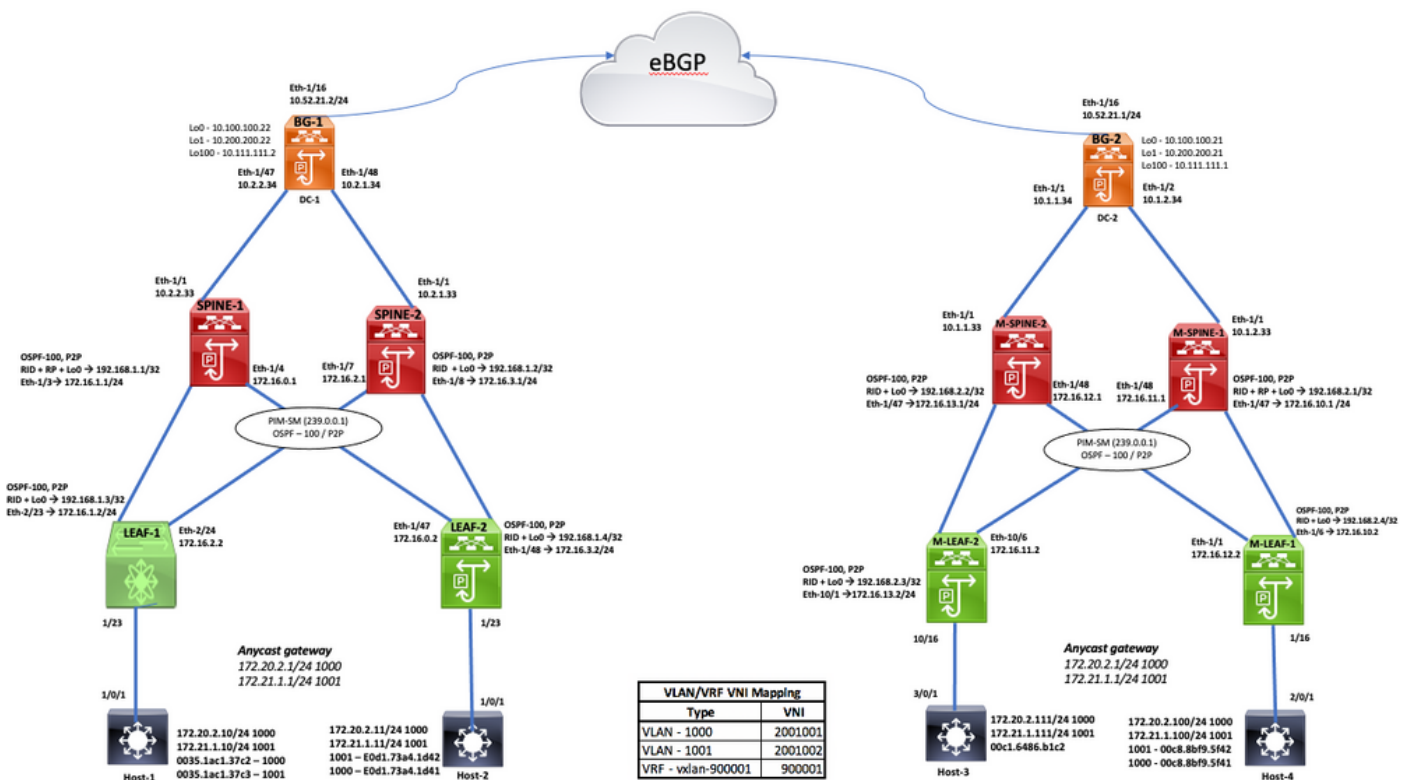
使用的组件

本文档中的信息基于以下软件和硬件版本：

leaf1#	N5K-C5672UP-16G-SUP	系统：版本7.3(0)N1(1)
leaf2#	N9K-C92160YC-X	NXOS：版本9.2(3)
spine1#	N9K-C9396PX	NXOS：版本9.2(3)
spine2#	N9K-C9396PX	NXOS：版本9.2(3)
MultisiteBG1#	N9K-C93108TC-EX	NXOS：版本9.2(3)
MultisiteBG2#	N9K-C-93108-FX	NXOS：版本9.3(1)
multisitespine2#	N9K-C9372TX-E	NXOS：版本9.2(3)
Multisastine1#	N9K-C92160YC-X	NXOS：版本9.2(3)
MultisetLeaf1#	N9K-C93108TC-EX	NXOS：版本7.0(3)I7(5)

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

拓扑



本文档介绍流量源自(DC-1、Host1/2 - 172.20.2.10/11)的位置，然后随数据包一起到达目的地DC-2、Host4(172.20.2.100)。

控制平面验证



通信流：
点上检查正确的IP、VLAN。

步骤1:在源节

ToLeaf1#show ip interface brief | exclude down

Interface	IP-Address	OK?	Method	Status	Protocol
Vlan1000	172.20.2.10	YES	NVRAM	up	up
Vlan1001	172.21.1.10	YES	NVRAM	up	up
GigabitEthernet1/0/1	unassigned	YES	unset	up	up

ToLeaf1#

现在，您可以从源节点访问其SVI-GW(Vlan-1000)。

ToLeaf1#**ping 172.20.2.1**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 172.20.2.1, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/205/1006 ms

ToLeaf1#

第二步：转到leaf1的第一跳，并检验第2层和第3层级别的内容。

现在查看leaf1是否从源节点获取mac和arp。

是，源节点的mac-address(0035.1ac1.37c2)和IP 172.20.2.10通过Vlan1000在eth1/23上获知。

leaf1#**show mac address-table**

Legend:

* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC

age - seconds since last seen,+ - primary entry using vPC Peer-Link

VLAN	MAC Address	Type	age	Secure	NTFY	Ports/SWID.SSID.LID
* 1001	0000.2222.3333	static	0	F	F	sup-eth2
* 1001	0035.1ac1.37c3	dynamic	500	F	F	Eth1/23
* 1001	005d.738e.a337	static	0	F	F	nve1/10.111.111.2
* 1001	00c8.8bf9.5f42	dynamic	0	F	F	nve1/10.111.111.2
* 1001	6cb2.ae91.38bf	static	0	F	F	nve1/10.200.200.22
* 1001	e0d1.73a4.1d42	dynamic	0	F	F	nve1/192.168.1.4
* 1000	0000.2222.3333	static	0	F	F	sup-eth2
* 1000	0035.1ac1.37c2	dynamic	70	F	F	Eth1/23

```

* 1000    005d.738e.a337    static    0          F    F    nve1/10.111.111.2
* 1000    00c8.8bf9.5f41    dynamic   0          F    F    nve1/10.111.111.2
* 1000    6cb2.ae91.38bf    static    0          F    F    nve1/10.200.200.22
* 1000    e0d1.73a4.1d41    dynamic   0          F    F    nve1/192.168.1.4

leaf1#

```

检验Source-Leaf中的ARP表

```
leaf1#show ip arp vrf all
```

```

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSOE
       # - Adjacencies Throttled for Glean
       D - Static Adjacencies attached to down interface

```

```
IP ARP Table for all contexts
```

```
Total number of entries: 5
```

Address	Age	MAC Address	Interface
172.21.1.10	00:08:14	0035.1ac1.37c3	Vlan1001
172.20.2.10	00:00:58	0035.1ac1.37c2	Vlan1000
10.31.121.1	00:08:14	2c31.24b0.bf46	mgmt0
172.16.1.1	00:07:51	0081.c41c.f007	Ethernet2/23
172.16.2.1	00:08:14	cc46.d68f.d74b	Ethernet2/24

```
leaf1#
```

这显示了leaf1如何获取消息以生成源节点的IP/mac条目。

```
leaf1#show system internal l2rib event-history mac | i 0035.1ac1.37c2 | be create
```

```

[04/24/20 13:10:09.721 UTC 6 4173] (1000,0035.1ac1.37c2,3):MAC route created with seq num:0,
flags:L (), soo:0, peerid:0

[04/24/20 13:10:09.732 UTC c 4173] (1000,0035.1ac1.37c2,3):Encoding MAC best route (ADD, client
id 4)

[04/24/20 13:10:09.871 UTC e 4173] (1000,0035.1ac1.37c2):Bound MAC-IP(172.20.2.10) to MAC, Total
MAC-IP linked: 1

[04/24/20 13:10:42.651 UTC 1a 4173] Received MAC ROUTE msg: addr: (1000-0035.1ac1.37c2) vni: 0

```


Route Distinguisher: 10.100.100.22:33767

Route Distinguisher: 10.100.100.22:33768

Route Distinguisher: 192.168.1.3:33767 (L2VNI 2001002)

*>1[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216

192.168.1.3 100 32768 i

*>1[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[32]:[172.20.2.10]/272

192.168.1.3 100 32768 i

Route Distinguisher: 192.168.1.3:33768 (L2VNI 2001001)

*>1[2]:[0]:[0]:[48]:[0035.1ac1.37c3]:[0]:[0.0.0.0]/216

192.168.1.3 100 32768 i

*>1[2]:[0]:[0]:[48]:[0035.1ac1.37c3]:[32]:[172.21.1.10]/272

192.168.1.3 100 32768 i

Route Distinguisher: 192.168.1.4:33767

Route Distinguisher: 192.168.1.4:33768

Route Distinguisher: 192.168.2.4:33767

Route Distinguisher: 192.168.2.4:33768

Route Distinguisher: 192.168.1.3:3 (L3VNI 900001)

leaf1#

DC-1 Leaf1到Spine2的路由验证

leaf1#show bgp l2vpn evpn neighbors 192.168.1.2 advertised-routes

Peer 192.168.1.2 routes for address family L2VPN EVPN:

BGP table version is 191, local router ID is 192.168.1.3

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

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

Route Distinguisher: 10.100.100.21:33767

Route Distinguisher: 10.100.100.21:33768

Route Distinguisher: 10.100.100.22:33767

Route Distinguisher: 10.100.100.22:33768

Route Distinguisher: 192.168.1.3:33767 (L2VNI 2001002)

***>1[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216**

192.168.1.3	100	32768	i
-------------	-----	-------	---

***>1[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[32]:[172.20.2.10]/272**

192.168.1.3	100	32768	i
-------------	-----	-------	---

Route Distinguisher: 192.168.1.3:33768 (L2VNI 2001001)

***>1[2]:[0]:[0]:[48]:[0035.1ac1.37c3]:[0]:[0.0.0.0]/216**

192.168.1.3	100	32768	i
-------------	-----	-------	---

***>1[2]:[0]:[0]:[48]:[0035.1ac1.37c3]:[32]:[172.21.1.10]/272**

192.168.1.3	100	32768	i
-------------	-----	-------	---

Route Distinguisher: 192.168.1.4:33767

Route Distinguisher: 192.168.1.4:33768

Route Distinguisher: 192.168.2.4:33767

Route Distinguisher: 192.168.2.4:33768

Route Distinguisher: 192.168.1.3:3 (L3VNI 900001)

leaf1#

由于有多条主干，请检查不同主干上的不同命令以确认（保持路由更新不变）。

在DC-1中，SPINE-1和SPINE-2与Leaf1、Leaf2和BGW-1具有EVPN邻居关系

spine1#**show bgp l2vpn evpn summary**

BGP summary information for VRF default, address family L2VPN EVPN

BGP router identifier 192.168.1.1, local AS number 200

BGP table version is 31, L2VPN EVPN config peers 3, capable peers 3

19 network entries and 19 paths using 4256 bytes of memory

BGP attribute entries [17/2788], BGP AS path entries [1/6]

BGP community entries [0/0], BGP clusterlist entries [0/0]

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.100.100.22	4	200	44002	43993	31	0	0	4w2d	11
192.168.1.3	4	200	43991	43989	31	0	0	4w2d	4
192.168.1.4	4	200	43996	43992	31	0	0	4w2d	4

spine1#

spine2#**show bgp l2vpn evpn summary**

BGP summary information for VRF default, address family L2VPN EVPN

BGP router identifier 192.168.1.2, local AS number 200

BGP table version is 65, L2VPN EVPN config peers 3, capable peers 3

19 network entries and 19 paths using 4256 bytes of memory

BGP attribute entries [17/2788], BGP AS path entries [1/6]

BGP community entries [0/0], BGP clusterlist entries [0/0]

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.100.100.22	4	200	47140	47115	65	0	0	4w4d 11	
192.168.1.3	4	200	47115	47112	65	0	0	4w4d 4	
192.168.1.4	4	200	47121	47116	65	0	0	4w4d 4	

spine2#

到目前为止，您已经到达了SPINE层，现在请看这是否传递到边界网关(DC-1、BGW-1)。

spine2#show bgp l2vpn evpn neighbors 10.100.100.22 advertised-routes

Peer 10.100.100.22 routes for address family L2VPN EVPN:

BGP table version is 65, Local Router ID is 192.168.1.2

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: 10.100.100.21:33767					
Route Distinguisher: 10.100.100.21:33768					
Route Distinguisher: 10.100.100.22:27001					
Route Distinguisher: 10.100.100.22:33767					
Route Distinguisher: 10.100.100.22:33768					
Route Distinguisher: 192.168.1.3:33767					
*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216					
	192.168.1.3		100	0	i
*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[32]:[172.20.2.10]/272					
	192.168.1.3		100	0	i

Route Distinguisher: 192.168.1.3:33768

*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c3]:[0]:[0.0.0.0]/216

192.168.1.3 100 0 i

*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c3]:[32]:[172.21.1.10]/272

192.168.1.3 100 0 i

Route Distinguisher: 192.168.1.4:33767

*>i[2]:[0]:[0]:[48]:[e0d1.73a4.1d41]:[0]:[0.0.0.0]/216

192.168.1.4 100 0 i

*>i[2]:[0]:[0]:[48]:[e0d1.73a4.1d41]:[32]:[172.20.2.11]/272

192.168.1.4 100 0 i

Route Distinguisher: 192.168.1.4:33768

*>i[2]:[0]:[0]:[48]:[e0d1.73a4.1d42]:[0]:[0.0.0.0]/216

192.168.1.4 100 0 i

*>i[2]:[0]:[0]:[48]:[e0d1.73a4.1d42]:[32]:[172.21.1.11]/272

192.168.1.4 100 0 i

Route Distinguisher: 192.168.2.4:33767

Route Distinguisher: 192.168.2.4:33768

spine2#

<<<<<<<<<

spine2#show bgp l2vpn evpn neighbors 10.100.100.22 advertised-routes | i 0035.1ac1.37c2
p 1 n 1

Route Distinguisher: 192.168.1.3:33767

*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216

192.168.1.3 100 0 i

```
*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[32]:[172.20.2.10]/272
192.168.1.3 100 0 i
spine2#
```

在DC-1的BORDER层进行验证，路由将从该层传播到其他站点。

然后，您可以看到在BGW-1中交换了哪些路由器、哪些路由类型以及有多少路由

DC-1 BGW-1

```
MultisiteBG1#show bgp l2vpn evpn summary
```

```
BGP summary information for VRF default, address family L2VPN EVPN
BGP router identifier 10.100.100.22, local AS number 200
BGP table version is 233, L2VPN EVPN config peers 3, capable peers 3
37 network entries and 45 paths using 7296 bytes of memory
BGP attribute entries [37/6068], BGP AS path entries [1/6]
BGP community entries [0/0], BGP clusterlist entries [4/16]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.100.100.21	4	100	47145	47121	233	0	0	4w4d	8
192.168.1.1	4	200	47153	47125	233	0	0	18:52:35	8
192.168.1.2	4	200	47139	47119	233	0	0	4w4d	8

Neighbor	T	AS	PfxRcd	Type-2	Type-3	Type-4	Type-5
10.100.100.21	E	100	8	6	2	0	0
192.168.1.1	I	200	8	8	0	0	0
192.168.1.2	I	200	8	8	0	0	0

```
MultisiteBG1#
```

```
MultisiteBG1#show bgp l2vpn evpn neighbors 10.100.100.21 advertised-routes
```

```
Peer 10.100.100.21 routes for address family L2VPN EVPN:
```

```
BGP table version is 233, Local Router ID is 10.100.100.22
```

```
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-redirect, I-injected

Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

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

Route Distinguisher: 10.100.100.21:33767

Route Distinguisher: 10.100.100.21:33768

Route Distinguisher: 10.100.100.22:27001 (ES [0300.0000.0000.c800.0309 0])

*>l[4]:[0300.0000.0000.c800.0309]:[32]:[10.200.200.22]/136

10.200.200.22	100	32768	i
---------------	-----	-------	---

Route Distinguisher: 10.100.100.22:33767 (L2VNI 2001002)

*>l[2]:[0]:[0]:[48]:[6cb2.ae91.38bf]:[0]:[0.0.0.0]/216

10.200.200.22	100	32768	i
---------------	-----	-------	---

*>l[3]:[0]:[32]:[10.200.200.22]/88

10.200.200.22	100	32768	i
---------------	-----	-------	---

Route Distinguisher: 10.100.100.22:33768 (L2VNI 2001001)

*>l[2]:[0]:[0]:[48]:[6cb2.ae91.38bf]:[0]:[0.0.0.0]/216

10.200.200.22	100	32768	i
---------------	-----	-------	---

*>l[3]:[0]:[32]:[10.200.200.22]/88

10.200.200.22	100	32768	i
---------------	-----	-------	---

Route Distinguisher: 192.168.1.3:33767

*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216

192.168.1.3	100	0	i
-------------	-----	---	---

*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[32]:[172.20.2.10]/272

192.168.1.3	100	0	i
-------------	-----	---	---

Route Distinguisher: 192.168.1.3:33768

*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c3]:[0]:[0.0.0.0]/216

```
192.168.1.3 100 0 i
```

```
*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c3]:[32]:[172.21.1.10]/272
```

```
192.168.1.3 100 0 i
```

```
Route Distinguisher: 192.168.1.4:33767
```

```
*>i[2]:[0]:[0]:[48]:[e0d1.73a4.1d41]:[0]:[0.0.0.0]/216
```

```
192.168.1.4 100 0 i
```

```
*>i[2]:[0]:[0]:[48]:[e0d1.73a4.1d41]:[32]:[172.20.2.11]/272
```

```
192.168.1.4 100 0 i
```

```
Route Distinguisher: 192.168.1.4:33768
```

```
*>i[2]:[0]:[0]:[48]:[e0d1.73a4.1d42]:[0]:[0.0.0.0]/216
```

```
192.168.1.4 100 0 i
```

```
*>i[2]:[0]:[0]:[48]:[e0d1.73a4.1d42]:[32]:[172.21.1.11]/272
```

```
192.168.1.4 100 0 i
```

```
Route Distinguisher: 192.168.2.4:33767
```

```
Route Distinguisher: 192.168.2.4:33768
```

```
MultisiteBG1#
```

```
MultisiteBG1#show bgp l2vpn evpn neighbors 10.100.100.21 advertised-routes | i 0035.1ac1.37c2  
p 1 n 1
```

```
Route Distinguisher: 192.168.1.3:33767
```

```
*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216
```

```
192.168.1.3 100 0 i
```

```
*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[32]:[172.20.2.10]/272
```

```
192.168.1.3 100 0 i
```

```
MultisiteBG1#
```

在DC-2/BGW-2上收到路由后，您确实可以看到路由通告到其连接的SPINE-1(192.168.2.1)

DC-2 BGW-2

```
MultisiteBG2#show bgp l2vpn evpn summary
```

```
BGP summary information for VRF default, address family L2VPN EVPN
```

```
BGP router identifier 10.100.100.21, local AS number 100
```

```
BGP table version is 142, L2VPN EVPN config peers 3, capable peers 2
```

```
43 network entries and 43 paths using 7680 bytes of memory
```

```
BGP attribute entries [33/5412], BGP AS path entries [1/6]
```

```
BGP community entries [0/0], BGP clusterlist entries [1/4]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.100.100.22	4	200	47169	47124	142	0	0	4w4d	12
192.168.2.1	4	100	47136	47124	142	0	0	4w4d	4
192.168.2.2	4	100	45969	45963	0	0	0	19:21:40	Idle

Neighbor	T	AS	PfxRcd	Type-2	Type-3	Type-4	Type-5
10.100.100.22	E	200	12	10	2	0	0
192.168.2.1	I	100	4	4	0	0	0
192.168.2.2	I	100	Idle	0	0	0	0

```
MultisiteBG2#
```

```
MultisiteBG2#show bgp l2vpn evpn neighbors 192.168.2.1 advertised-routes | i 0035.1ac1.37c2 p 1 n 1
```

```
Route Distinguisher: 192.168.1.3:33767
```

```
*>e[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216
```

```
10.111.111.2 2000 0 200 i
```

```
*>e[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[32]:[172.20.2.10]/272
```

```
10.111.111.2 2000 0 200 i
```

```
MultisiteBG2#
```

路由会进一步通告到目标所连接的DC-2、枝叶-1。

DC-2主干到枝叶路由通告

```
Multistespine1#show bgp l2vpn evpn summary
```

```
BGP summary information for VRF default, address family L2VPN EVPN
```

```
BGP router identifier 192.168.2.1, local AS number 100
```

```
BGP table version is 50, L2VPN EVPN config peers 3, capable peers 2
```

```
19 network entries and 19 paths using 4256 bytes of memory
```

```
BGP attribute entries [15/2460], BGP AS path entries [1/6]
```

```
BGP community entries [0/0], BGP clusterlist entries [0/0]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.100.100.21	4	100	47152	47131	50	0	0	4w4d 15	
192.168.2.3	4	100	0	0	0	0	0	4w4d Idle	
192.168.2.4	4	100	47135	47131	50	0	0	4w4d 4	

```
Multistespine1#
```

```
Multistespine1#show bgp l2vpn evpn neighbors 192.168.2.4 advertised-routes | i 0035.1ac1.37c2 p 1 n 1
```

```
Route Distinguisher: 192.168.1.3:33767
```

```
*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216
```

```
10.111.111.1 2000 100 0 200 i
```

```
*>i[2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[32]:[172.20.2.10]/272
```

```
10.111.111.1 2000 100 0 200 i
```

```
Multistespine1#
```

现在，您可以看到，第2层拓扑已从DC-1、枝叶-1扩展到DC-2、枝叶-1。

DC-2枝叶-1验证

```
MultisteLeaf1#show bgp l2vpn evpn summary
```

```
BGP summary information for VRF default, address family L2VPN EVPN
```

```
BGP router identifier 192.168.2.4, local AS number 100
```

```
BGP table version is 188, L2VPN EVPN config peers 2, capable peers 1
```



```
MultisteLeaf1#show ip arp vrf vxlan-900001
```

```
Flags: * - Adjacencies learnt on non-active FHRP router
+ - Adjacencies synced via CFSOE
# - Adjacencies Throttled for Glean
CP - Added via L2RIB, Control plane Adjacencies
PS - Added via L2RIB, Peer Sync
RO - Re-Originated Peer Sync Entry
D - Static Adjacencies attached to down interface
```

```
IP ARP Table for context vxlan-900001
```

```
Total number of entries: 2
```

Address	Age	MAC Address	Interface	Flags
172.21.1.100	00:04:09	00c8.8bf9.5f42	Vlan1001	
172.20.2.100	00:04:09	00c8.8bf9.5f41	Vlan1000	

```
MultisteLeaf1#
```

数据平面故障排除

在多台设备上测试数据计划验证，以了解不同的数据包捕获方法和变体。



在本例中，放大作为目标枝叶的N5K。

```
SOURCE: 172.20.2.100 (Host-4)
DESTINATION: 172.20.2.10 (Host-1)
```

对Source-Host4的基本检查

```
toMultisiteLeaf1#show ip interface brief | exclude down
Interface IP-Address OK? Method Status Protocol
Vlan1000 172.20.2.100 YES NVRAM up up
Vlan1001 172.21.1.100 YES NVRAM up up
```



```
toMultisiteLeaf1#ping 172.20.2.11 repeat 2
Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 172.20.2.11, timeout is 2 seconds:
!!
Success rate is 100 percent (2/2), round-trip min/avg/max = 1/1/1 ms
toMultisiteLeaf1#
```

```
leaf2#ethanalyzer local interface inband mirror
```

```
Capturing on inband
2020-05-27 12:20:57.081654 172.20.2.100 -> 172.20.2.11 ICMP Echo (ping) request
2020-05-27 12:20:57.082193 172.20.2.11 -> 172.20.2.100 ICMP Echo (ping) reply
2020-05-27 12:20:57.084902 172.20.2.100 -> 172.20.2.11 ICMP Echo (ping) request
2020-05-27 12:20:57.087406 172.20.2.11 -> 172.20.2.100 ICMP Echo (ping) reply
```

```
4 packets captured
leaf2#
```

```
leaf2#ethanalyzer local interface inband display-filter "ip.addr==172.20.2.100 &&
ip.addr==172.20.2.11 && icmp" limit-captured-frames 0
```

```
Capturing on inband
2020-05-27 12:20:57.081654 172.20.2.100 -> 172.20.2.11 ICMP Echo (ping) request
2020-05-27 12:20:57.082193 172.20.2.11 -> 172.20.2.100 ICMP Echo (ping) reply
2020-05-27 12:20:57.084902 172.20.2.100 -> 172.20.2.11 ICMP Echo (ping) request
2020-05-27 12:20:57.087406 172.20.2.11 -> 172.20.2.100 ICMP Echo (ping) reply
```

```
4 packets captured
leaf2#
```

如果要从源发送某些数据并在目标枝叶上捕获，请执行以下步骤：

1. 从Host-4向Host-2发起ping。
2. 源：172.20.2.100 //源MAC地址：00:C8:8B:F9:5F:41
3. 目的：172.20.2.11 //目的MAC地址：E0:D1:73:A4:1D:41
4. 数据包大小：777
5. Do-Not-Fragment：是
6. 数据模式字符串：牛肉

在本例中，您拥有足够的数据包计数，因此您可以在源枝叶和目的枝叶上捕获数据包。

```
toMultisiteLeaf1#ping 172.20.2.11 repeat 200000 data beef df-bit validate size 777
Type escape sequence to abort.
Sending 200000, 777-byte ICMP Echos to 172.20.2.11, timeout is 2 seconds:
Packet sent with the DF bit set
Packet has data pattern 0xBEEF
Reply data will be validated
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

<.....>
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (7376/7376), round-trip min/avg/max = 1/5/151 ms
toMultisiteLeaf1#
```


PortOnSlice: 11
Table entries for interface Ethernet1/16
<.....snipped for brevity.....>

```
MultisteLeaf1#show system internal ethpm info interface ethernet 1/16 | grep slice
IF_STATIC_INFO: port_name=Ethernet1/16,if_index:0x1a001e00,ltl=6084,slot=0,
nxos_port=60,dmod=1,dpid=11,unit=0,queue=65535,xbar_unitbmp=0x0,ns_pid=255,slice_num=0,port_on_s
lice=11,src_id=22
MultisteLeaf1#
```

在硬件级别检查是否真正接收了来自Host-4的数据包。

```
MultisteLeaf1#attach module 1
module-1#debug platform internal tah elam asic 0

module-1(TAH-elam)#trigger init asic 0 slice 0 in-select 6 out-select 0 use-src-id 22

param values: start asic 0, start slice 0, lu-a2d 1, in-select 6, out-select 0, src_id 22
```

```
module-1(TAH-elam-insel6)#start
module-1(TAH-elam-insel6)#report
Initting block addresses
```

```
SUGARBOWL ELAM REPORT SUMMARY
slot - 1, asic - 0, slice - 0
=====
```

```
Incoming Interface: Eth1/16
Src Idx : 0x3d, Src BD : 1000
Outgoing Interface Info: dmod 1, dpid 17
Dst Idx : 0x15, Dst BD : 1000
```

Packet Type: IPv4

```
Dst MAC address: E0:D1:73:A4:1D:41
Src MAC address: 00:C8:8B:F9:5F:41
.lq Tag0 VLAN: 1000, cos = 0x0
```

```
Dst IPv4 address: 172.20.2.11
Src IPv4 address: 172.20.2.100
Ver = 4, DSCP = 0, Don't Fragment = 1
Proto = 1, TTL = 255, More Fragments = 0
Hdr len = 20, Pkt len = 777, Checksum = 0xcffe
```

```
L4 Protocol : 1
ICMP type : 8
ICMP code : 0
```

Drop Info:

LUA:
LUB:
LUC:
LUD:
Final Drops:

```
module-1(TAH-elam-insel6)# exit
module-1(TAH-elam)# exit
module-1# exit
MultisteLeaf1#
```

如果要将其视为pcap，请让其持续ping，监控会话在源枝叶和目标枝叶上配置。

0110 be ef be ef be ef be ef be ef be ef be ef be ef
0120 be ef be ef be ef be ef be ef be ef be ef be ef
0130 be ef be ef be ef be ef be ef be ef be ef be ef
0140 be ef be ef be ef be ef be ef be ef be ef be ef
0150 be ef be ef be ef be ef be ef be ef be ef be ef
0160 be ef be ef be ef be ef be ef be ef be ef be ef
0170 be ef be ef be ef be ef be ef be ef be ef be ef
0180 be ef be ef be ef be ef be ef be ef be ef be ef
0190 be ef be ef be ef be ef be ef be ef be ef be ef
01a0 be ef be ef be ef be ef be ef be ef be ef be ef
01b0 be ef be ef be ef be ef be ef be ef be ef be ef
01c0 be ef be ef be ef be ef be ef be ef be ef be ef
01d0 be ef be ef be ef be ef be ef be ef be ef be ef
01e0 be ef be ef be ef be ef be ef be ef be ef be ef
01f0 be ef be ef be ef be ef be ef be ef be ef be ef
0200 be ef be ef be ef be ef be ef be ef be ef be ef
0210 be ef be ef be ef be ef be ef be ef be ef be ef
0220 be ef be ef be ef be ef be ef be ef be ef be ef
0230 be ef be ef be ef be ef be ef be ef be ef be ef
0240 be ef be ef be ef be ef be ef be ef be ef be ef
0250 be ef be ef be ef be ef be ef be ef be ef be ef
0260 be ef be ef be ef be ef be ef be ef be ef be ef
0270 be ef be ef be ef be ef be ef be ef be ef be ef
0280 be ef be ef be ef be ef be ef be ef be ef be ef
0290 be ef be ef be ef be ef be ef be ef be ef be ef
02a0 be ef be ef be ef be ef be ef be ef be ef be ef
02b0 be ef be ef be ef be ef be ef be ef be ef be ef
02c0 be ef be ef be ef be ef be ef be ef be ef be ef
02d0 be ef be ef be ef be ef be ef be ef be ef be ef
02e0 be ef be ef be ef be ef be ef be ef be ef 00

Data: 0000000014675F78BEEFBEEFBEEFBEEFBEEFBEEFBEEFBEEFBEEF...

[Length: 749]

Frame 4 (795 bytes on wire, 795 bytes captured)

Arrival Time: May 31, 2020 15:44:46.888728000

[Time delta from previous captured frame: 0.047867000 seconds]

[Time delta from previous displayed frame: 0.047867000 seconds]

[Time since reference or first frame: 0.121317000 seconds]

Frame Number: 4

Frame Length: 795 bytes

Capture Length: 795 bytes

[Frame is marked: False]

[Protocols in frame: eth:vlan:ip:icmp:data]

Ethernet II, Src: 00:c8:8b:f9:5f:41 (00:c8:8b:f9:5f:41), Dst: e0:d1:73:a4:1d:41 (e0:d1:73:a4:1d:41)

Destination: e0:d1:73:a4:1d:41 (e0:d1:73:a4:1d:41)

Address: e0:d1:73:a4:1d:41 (e0:d1:73:a4:1d:41)

.... 0 = IG bit: Individual address (unicast)

.... .0. = LG bit: Globally unique address (factory default)

Source: 00:c8:8b:f9:5f:41 (00:c8:8b:f9:5f:41)

Address: 00:c8:8b:f9:5f:41 (00:c8:8b:f9:5f:41)

.... 0 = IG bit: Individual address (unicast)

.... .0. = LG bit: Globally unique address (factory default)

Type: 802.1Q Virtual LAN (0x8100)

802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 1000

000. = Priority: 0

...0 = CFI: 0

.... 0011 1110 1000 = ID: 1000

Type: IP (0x0800)

Internet Protocol, Src: 172.20.2.100 (172.20.2.100), Dst: 172.20.2.11 (172.20.2.11)

Version: 4

Header length: 20 bytes

Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)

0000 00.. = Differentiated Services Codepoint: Default (0x00)

```
.... ..0. = ECN-Capable Transport (ECT): 0
.... ...0 = ECN-CE: 0
Total Length: 777
Identification: 0xaf65 (44901)
Flags: 0x02 (Don't Fragment)
0.. = Reserved bit: Not Set
.1. = Don't fragment: Set
..0 = More fragments: Not Set
Fragment offset: 0
Time to live: 255
Protocol: ICMP (0x01)
Header checksum: 0xbd1b [correct]
[Good: True]
[Bad : False]
Source: 172.20.2.100 (172.20.2.100)
Destination: 172.20.2.11 (172.20.2.11)
Internet Control Message Protocol
Type: 8 (Echo (ping) request)
Code: 0 ()
Checksum: 0x704a [correct]
Identifier: 0x001c
Sequence number: 7430 (0x1d06)
Data (749 bytes)
```

```
0000 00 00 00 00 14 67 5f aa be ef be ef be ef be ef .....g_.....
0010 be ef be ef be ef be ef be ef be ef be ef be ef .....
0020 be ef be ef be ef be ef be ef be ef be ef be ef .....
0030 be ef be ef be ef be ef be ef be ef be ef be ef .....
0040 be ef be ef be ef be ef be ef be ef be ef be ef .....
0050 be ef be ef be ef be ef be ef be ef be ef be ef .....
0060 be ef be ef be ef be ef be ef be ef be ef be ef .....
0070 be ef be ef be ef be ef be ef be ef be ef be ef .....
0080 be ef be ef be ef be ef be ef be ef be ef be ef .....
0090 be ef be ef be ef be ef be ef be ef be ef be ef .....
00a0 be ef be ef be ef be ef be ef be ef be ef be ef .....
00b0 be ef be ef be ef be ef be ef be ef be ef be ef .....
00c0 be ef be ef be ef be ef be ef be ef be ef be ef .....
00d0 be ef be ef be ef be ef be ef be ef be ef be ef .....
00e0 be ef be ef be ef be ef be ef be ef be ef be ef .....
00f0 be ef be ef be ef be ef be ef be ef be ef be ef .....
0100 be ef be ef be ef be ef be ef be ef be ef be ef .....
0110 be ef be ef be ef be ef be ef be ef be ef be ef .....
0120 be ef be ef be ef be ef be ef be ef be ef be ef .....
0130 be ef be ef be ef be ef be ef be ef be ef be ef .....
0140 be ef be ef be ef be ef be ef be ef be ef be ef .....
0150 be ef be ef be ef be ef be ef be ef be ef be ef .....
0160 be ef be ef be ef be ef be ef be ef be ef be ef .....
0170 be ef be ef be ef be ef be ef be ef be ef be ef .....
0180 be ef be ef be ef be ef be ef be ef be ef be ef .....
0190 be ef be ef be ef be ef be ef be ef be ef be ef .....
01a0 be ef be ef be ef be ef be ef be ef be ef be ef .....
01b0 be ef be ef be ef be ef be ef be ef be ef be ef .....
01c0 be ef be ef be ef be ef be ef be ef be ef be ef .....
01d0 be ef be ef be ef be ef be ef be ef be ef be ef .....
MultisteLeaf1#
```

对最后一跳DC-1、枝叶-2的基本检查

Address: 00:c8:8b:f9:5f:41 (00:c8:8b:f9:5f:41)
.... ..0 = IG bit: Individual address (unicast)
.... ..0. = LG bit: Globally unique address (factory default)
Type: IP (0x0800)
Internet Protocol, Src: 172.20.2.100 (172.20.2.100), Dst: 172.20.2.11 (172.20.2.11)
Version: 4
Header length: 20 bytes
Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
0000 00.. = Differentiated Services Codepoint: Default (0x00)
.... ..0. = ECN-Capable Transport (ECT): 0
.... ..0 = ECN-CE: 0
Total Length: 777
Identification: 0x8237 (33335)
Flags: 0x02 (Don't Fragment)
0.. = Reserved bit: Not Set
.1. = Don't fragment: Set
..0 = More fragments: Not Set
Fragment offset: 0
Time to live: 255
Protocol: ICMP (0x01)
Header checksum: 0xea49 [correct]
[Good: True]
[Bad : False]
Source: 172.20.2.100 (172.20.2.100)
Destination: 172.20.2.11 (172.20.2.11)
Internet Control Message Protocol
Type: 8 (Echo (ping) request)
Code: 0 ()
Checksum: 0x980f [correct]
Identifier: 0x001f
Sequence number: 2515 (0x09d3)
Data (749 bytes)

0000 00 00 00 00 14 7f 4a fd be ef be ef be ef be efJ.....
0010 be ef be ef be ef be ef be ef be ef be ef be ef
0020 be ef be ef be ef be ef be ef be ef be ef be ef
0030 be ef be ef be ef be ef be ef be ef be ef be ef
0040 be ef be ef be ef be ef be ef be ef be ef be ef
0050 be ef be ef be ef be ef be ef be ef be ef be ef
0060 be ef be ef be ef be ef be ef be ef be ef be ef
0070 be ef be ef be ef be ef be ef be ef be ef be ef
0080 be ef be ef be ef be ef be ef be ef be ef be ef
0090 be ef be ef be ef be ef be ef be ef be ef be ef
00a0 be ef be ef be ef be ef be ef be ef be ef be ef
00b0 be ef be ef be ef be ef be ef be ef be ef be ef
00c0 be ef be ef be ef be ef be ef be ef be ef be ef
00d0 be ef be ef be ef be ef be ef be ef be ef be ef
00e0 be ef be ef be ef be ef be ef be ef be ef be ef
00f0 be ef be ef be ef be ef be ef be ef be ef be ef
0100 be ef be ef be ef be ef be ef be ef be ef be ef
0110 be ef be ef be ef be ef be ef be ef be ef be ef
0120 be ef be ef be ef be ef be ef be ef be ef be ef
0130 be ef be ef be ef be ef be ef be ef be ef be ef
0140 be ef be ef be ef be ef be ef be ef be ef be ef
0150 be ef be ef be ef be ef be ef be ef be ef be ef
0160 be ef be ef be ef be ef be ef be ef be ef be ef
0170 be ef be ef be ef be ef be ef be ef be ef be ef
0180 be ef be ef be ef be ef be ef be ef be ef be ef
0190 be ef be ef be ef be ef be ef be ef be ef be ef
01a0 be ef be ef be ef be ef be ef be ef be ef be ef
01b0 be ef be ef be ef be ef be ef be ef be ef be ef
01c0 be ef be ef be ef be ef be ef be ef be ef be ef
01d0 be ef be ef be ef be ef be ef be ef be ef be ef
01e0 be ef be ef be ef be ef be ef be ef be ef be ef

01f0 be ef be ef be ef be ef be ef be ef be ef be ef
0200 be ef be ef be ef be ef be ef be ef be ef be ef
0210 be ef be ef be ef be ef be ef be ef be ef be ef
0220 be ef be ef be ef be ef be ef be ef be ef be ef
0230 be ef be ef be ef be ef be ef be ef be ef be ef
0240 be ef be ef be ef be ef be ef be ef be ef be ef
0250 be ef be ef be ef be ef be ef be ef be ef be ef
0260 be ef be ef be ef be ef be ef be ef be ef be ef
0270 be ef be ef be ef be ef be ef be ef be ef be ef
0280 be ef be ef be ef be ef be ef be ef be ef be ef
0290 be ef be ef be ef be ef be ef be ef be ef be ef
02a0 be ef be ef be ef be ef be ef be ef be ef be ef
02b0 be ef be ef be ef be ef be ef be ef be ef be ef
02c0 be ef be ef be ef be ef be ef be ef be ef be ef
02d0 be ef be ef be ef be ef be ef be ef be ef be ef
02e0 be ef be ef be ef be ef be ef be ef be ef 00

Data: 0000000147F4AFDBEEFBEEFBEEFBEEFBEEFBEEFBEEFBEEFBEEF...

[Length: 749]

关于此翻译

思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。