

Dépannage d'EVPN/VxLAN dans un environnement multisite

Table des matières

[Introduction](#)

[Conditions préalables](#)

[Exigences](#)

[Composants utilisés](#)

[Topologie](#)

[Vérification du plan de contrôle](#)

[Vérification de la table ARP dans Source-Leaf](#)

[Vérification de la route DC-1 Leaf1 vers Spine1](#)

[Vérification de la route Leaf1 vers Spine2 de DC-1](#)

[DC-1 BGW-1](#)

[DC-2 BGW-2](#)

[Annonce de route spine vers leaf DC-2](#)

[Vérification de DC-2 Leaf-1](#)

[Dépannage du plan de données](#)

[Vérification De Base Sur Source-Host4](#)

[Vérification de la feuille de destination 1](#)

[Configurer la session de surveillance](#)

[Contrôles de base sur le premier saut DC-2, Leaf-1](#)

[Surveiller la session sur le leaf source](#)

[Surveiller la session sur le leaf de destination](#)

[Capture sur le leaf source](#)

[Contrôles de base sur le dernier saut DC-1, Leaf-2](#)

Introduction

Ce document décrit comment dépanner Ethernet VPN/Virtual Extensible LAN (EVPN/VxLAN) dans un environnement multisite.

Conditions préalables

Exigences

Cisco vous recommande de prendre connaissance des rubriques suivantes :

- VPN de couche 3 à commutation multiprotocole par étiquette (MPLS)
- Protocole MP-BGP (Multiprotocol-Border Gateway Protocol)
- EVPN


Composants utilisés

Les informations contenues dans ce document sont basées sur les versions de matériel et de logiciel suivantes :

leaf1#	N5K-C5672UP-16G-SUP	système : version 7.3(0)N1(1)
leaf2#	N9K-C92160YC-X	NXOS : version 9.2(3)
spine1#	N9K-C9396PX	NXOS : version 9.2(3)
spine2#	N9K-C9396PX	NXOS : version 9.2(3)
MultisiteBG1#	N9K-C93108TC-EX	NXOS : version 9.2(3)
MultisiteBG2#	N9K-C93108TC-FX	NXOS : version 9.3(1)
multisitespine2#	N9K-C9372TX-E	NXOS : version 9.2(3)
Multistespine1#	N9K-C92160YC-X	NXOS : version 9.2(3)
Feuille1 multisite#	N9K-C93108TC-EX	NXOS : version 7.0(3)I7(5)

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. Si votre réseau est en ligne, assurez-vous de bien comprendre l'incidence possible des commandes.

Topologie



Ce document décrit l'origine du trafic (DC-1, Hôte1/2 - 172.20.2.10/11), puis accompagne nos paquets jusqu'à la destination DC-2, Hôte4 (172.20.2.100).

Vérification du plan de contrôle



Flux de trafic :
 Recherchez l'adresse IP et le VLAN appropriés sur le noeud source.

Étape 1.

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#

Vous avez maintenant l'accessibilité du noeud source à sa SVI-GW de 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#

Étape 2. Passez au premier saut, **leaf1**, et vérifiez les éléments au niveau des couches 2 et 3.

Voyez maintenant si **leaf1** apprend le mac, arp du noeud source ou non.

Oui, l'adresse MAC (0035.1ac1.37c2), IP 172.20.2.10 du noeud source est apprise sur eth1/23 via Vlan1000.

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#

```

Vérification de la table ARP dans Source-Leaf

```
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#
```

Ceci montre comment leaf1 obtient le message pour faire une entrée IP/mac du noeud source.

```

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
admin_dist: 0 seq_num: 0 rt_flags: L soo: 0 dg_count: 0 res: 0 esi: (F) nh_count: 1

[04/24/20 13:10:42.651 UTC 1c 4173] (1000,0035.1ac1.37c2):Mobility check for new rte from prod:
3

[04/24/20 13:10:42.651 UTC 1d 4173] (1000,0035.1ac1.37c2):Current non-del-pending route
local:yes, remote:no, linked mac-ip count:1

[04/24/20 13:10:42.651 UTC 1e 4173] (1000,0035.1ac1.37c2):Mobility type: local-to-local; New
route SOO: 0, Seq num: 0; Existing route SOO: 0, Seq num: 0

[04/24/20 13:10:42.651 UTC 1f 4173] (1000,0035.1ac1.37c2):Local Update, Add to DB

[04/24/20 13:10:42.651 UTC 20 4173] (1000,0035.1ac1.37c2,3):Using seq number from Recv-based
route

<.....snipped for brevity.....>

```

leaf1#

Une fois que vous avez vérifié comment l'entrée mac/IP a été faite dans leaf1, c'est l'un des points d'observation importants où les informations mac/mac-ip sont liées avec BGP et sont annoncées comme des routes l2vpn/evpn.

```
leaf1#show bgp l2vpn evpn 0035.1ac1.37c2
```

```
Route Distinguisher: 192.168.1.3:33767 (L2VNI 2001002)
```

```
BGP routing table entry for [2]:[0]:[0]:[48]:[0035.1ac1.37c2]:[0]:[0.0.0.0]/216, version 99
```

```
Paths: (1 available, best #1)
```

```
Flags: (0x00010a) on xmit-list, is not in l2rib/evpn
```

```
Advertised path-id 1
```

```
Path type: local, path is valid, is best path
```

```
AS-Path: NONE, path locally originated
```

```
192.168.1.3 (metric 0) from 0.0.0.0 (192.168.1.3)
```

```
Origin IGP, MED not set, localpref 100, weight 32768
```

```
Received label 2001002
```

```
Extcommunity:
```


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#

Vérification de la route Leaf1 vers Spine2 de DC-1

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-redirect, 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#
```

Comme il y a plusieurs spines, vérifiez différentes commandes sur différents spines pour confirmer (garder les mises à jour de route intactes).

Dans DC-1, SPINE-1 et SPINE-2 ont un voisinage EVPN avec Leaf1, Leaf2 et BGW-1

```
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#
```

Jusqu'à présent, vous avez atteint la couche SPINE, voyez maintenant si elle est transmise à la passerelle frontière (DC-1, BGW-1) ou non.

```
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-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
```

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#

Vérifiez au niveau de la couche BORDER de DC-1, à partir de laquelle les routes sont propagées vers un site différent.

Vous pouvez alors voir avec qui, quel type de routes et combien de routes sont échangées dans BGW-1

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-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 (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#
```

Une fois que les routes sont reçues au niveau de DC-2/BGW-2, vous voyez l'annonce de route vers son SPINE-1 connecté (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#
```

Les routes sont également annoncées à DC-2, Leaf-1 où la destination est connectée.

Annonce de route spine vers leaf DC-2

```
Multistespinel#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	

```
Multistespinel#
```

```
Multistespinel#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

Multistespinel#

Vous voyez maintenant que la topologie de couche 2 a été étendue de DC-1, Leaf-1 à DC-2, Leaf-1.

Vérification de DC-2 Leaf-1

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
36 network entries and 36 paths using 6120 bytes of memory
BGP attribute entries [16/2560], BGP AS path entries [1/6]
BGP community entries [0/0], BGP clusterlist entries [1/4]

Table with columns: Neighbor, V, AS, MsgRcvd, MsgSent, TblVer, InQ, OutQ, Up/Down, State/PfxRcd. Rows show neighbor details for 192.168.2.1 and 192.168.2.2.

MultisteLeaf1#

MultisteLeaf1#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,
(T) - True, (F) - False, C - ControlPlane MAC, ~ - vsan

Table with columns: VLAN, MAC Address, Type, age, Secure, NTFY, Ports. Rows show MAC address table entries for VLANs 1, 101, and 1000.

```

C 1000    e0d1.73a4.1d41    dynamic  0          F        F        nve1(10.111.111.1)
C 1001    0035.1ac1.37c3    dynamic  0          F        F        nve1(10.111.111.1)
* 1001    00c8.8bf9.5f42    dynamic  0          F        F        Eth1/16
C 1001    e0d1.73a4.1d42    dynamic  0          F        F        nve1(10.111.111.1)
G   -     0000.2222.3333    static   -          F        F        sup-eth1(R)
G   -     00be.75f4.544d    static   -          F        F        sup-eth1(R)
G  101    00be.75f4.544d    static   -          F        F        sup-eth1(R)
G 1000    00be.75f4.544d    static   -          F        F        sup-eth1(R)
G 1001    00be.75f4.544d    static   -          F        F        sup-eth1(R)

```

MultisteLeaf1#

DC-2, Leaf-1 a appris le Mac de l'hôte 4 via le VLAN1000.

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#

Dépannage du plan de données

La vérification du plan de données est testée sur plusieurs périphériques afin de comprendre



avant sur N9K.

Vous pouvez commencer à partir de l'hôte 4 et envoyer une simple requête ping à l'hôte 2. Bien qu'il s'agisse d'un trafic de transit pour les commutateurs LEAF, vérifiez toujours si vous êtes en mesure de traiter ces paquets correctement ou non. Confirmez à DC-1/Leaf-2.

Afin de confirmer, surveillez l'interface eth-1/23, et envoyez ces paquets au CPU et exécutez ethanalyzer/dmirror pour le même.

Configurer la session de surveillance

```
leaf2#configure terminal
leaf2(config)#monitor session 1
leaf2(config-monitor)#source interface ethernet 1/23
leaf2(config-monitor)#destination interface sup-eth 0
leaf2(config-monitor)#no shut
leaf2(config-monitor)#exit
```

```
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#
```

Dans le cas où vous voulez envoyer des données de la source et capturer sur le leaf de destination, voici la procédure :

1. Lancez une requête ping de l'hôte 4 vers l'hôte 2.
2. SOURCE : 172.20.2.100 // Adresse MAC source : 00:C8:8B:F9:5F:41

3. DESTINATION : 172.20.2.11 // Adresse MAC DST : E0:D1:73:A4:1D:41
4. Taille du paquet : 777
5. Ne pas fragmenter : OUI
6. Chaîne de données : beef

Dans cet exemple, vous disposez d'un nombre de paquets suffisant pour capturer les paquets sur les feuilles source et de destination.

```
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#
```

Contrôles de base sur le premier saut DC-2, Leaf-1



```
MultisteLeaf1#show module
Mod Ports Module-Type Model Status
-----
1 54 48x10GT + 6x40G/100G Ethernet Module N9K-C93108TC-EX active *

Mod Sw Hw Slot
-----
1 7.0(3)I7(5) 1.3 NA

Mod MAC-Address(es) Serial-Num
-----
1 00-be-75-f4-54-46 to 00-be-75-f4-54-95 FDO220225UX

Mod Online Diag Status
-----
1 Pass

* this terminal session
MultisteLeaf1#
```

MultisteLeaf1#show cdp neighbors

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

```
Device-ID Local Intrfce Hldtme Capability Platform Port ID
MX066-H-03-SW.cisco.com
mgmt0 141 S I WS-C2960X-48T Gig1/0/31
Multistespinel(FD022150SJZ)
Eth1/6 142 R S s N9K-C92160YC- Eth1/47
toMultisiteLeaf1 Eth1/16 128 R S I WS-C3750X-24S Gig2/0/1 <<<<<<<<<<<< Towards Host-4
```

Total entries displayed: 3
MultisteLeaf1#

MultisteLeaf1#show hardware internal tah interface et1/16

#####

IfIndex: 0x1a001e00
DstIndex: 6084
IfType: 26
Asic: 0

Asic: 0 <<<<<<<<<<<<<<<<<<<<<<<<<

AsicPort: 11

SrcId: 22 <<<<<<<<<<<<<<<<<<<<<<<<<

Slice: 0 <<<<<<<<<<<<<<<<<<<<<<<<<

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#

Vérifiez au niveau matériel que vous avez réellement reçu les paquets de l'hôte 4.

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

.1q 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#
```

Si vous voulez voir en tant que pcap, laissez-le envoyer une requête ping en continu, la session de surveillance est configurée sur le leaf source et de destination.

Surveiller la session sur le leaf source

```
MultisteLeaf1#show run | section monitor
```

```
monitor session 1
```

```
source interface Ethernet1/16 both
```

```
destination interface sup-eth0
```

```
no shut
```

```
MultisteLeaf1#
```

Surveiller la session sur le leaf de destination

```
leaf2(config)#show run | section monitor
```

```
monitor session 1
```

```
source interface Ethernet1/23 both
```

```
destination interface sup-eth0
```

```
no shut
```

```
leaf2(config)#
```

Capture sur le leaf source

```
MultisteLeaf1#ethalyzer local interface inband display-filter "ip.addr==172.20.2.100 &&  
ip.addr==172.20.2.11 && icmp" limit-captured-frames 0 detail
```

```
Capturing on inband
```

```
Frame 1 (791 bytes on wire, 791 bytes captured)
```

```
Arrival Time: May 31, 2020 15:44:46.767411000
```

```
[Time delta from previous captured frame: 0.000000000 seconds]
```


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)
.... 00000000 = IG bit: Individual address (unicast)
.... 00000000 = 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)
.... 00000000 = IG bit: Individual address (unicast)
.... 00000000 = LG bit: Globally unique address (factory default)
Type: 802.1Q Virtual LAN (0x8100)
802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 1000
00000000 = Priority: 0
...00000000 = CFI: 0
.... 001111101000 = 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)
000000.. = Differentiated Services Codepoint: Default (0x00)
.... 0000 = ECN-Capable Transport (ECT): 0
.... 0000 = 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 efg_.....
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#

Contrôles de base sur le dernier saut DC-1, Leaf-2



```
leaf2#show module
```

```
Mod Ports Module-Type Model Status
```

```
-----
```

```
1 54 48x10G + 4x40G + 2x100G Ethernet Modu N9K-C92160YC-X active *
```

```
Mod Sw Hw Slot
```

```
-----
```

```
1 9.2(3) 1.3 NA
```

```
Mod MAC-Address(es) Serial-Num
```

```
-----
```

```
1 70-79-b3-3e-81-1c to 70-79-b3-3e-81-69 FD022111H2V
```

```
Mod Online Diag Status
```

```
-----
```

```
1 Pass
```

```
* this terminal session
```

```
leaf2#show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute
```

```
Device-ID Local Intrfce Hldtme Capability Platform Port ID
```

```
Toleaf2 Eth1/23 142 S I WS-C3750X-24S Gig1/0/1 <<<<<<<<<<<<<<<< Towards Host-2
```

```
switch(SAL2024RRYF)
```

```
Eth1/47 175 R S I s N9K-C9372PX-E Eth1/4
```

```
spine2(SAL1949UELD)
```


0290 be ef 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 be ef
02b0 be ef 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 be ef
02d0 be ef 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: 00000000147F4ADBEEFBEEFBEEFBEEFBEEFBEEFBEEFBEEFBEEFBEEF...
[Length: 749]

Frame 2 (791 bytes on wire, 791 bytes captured)

Arrival Time: May 31, 2020 15:45:55.694904000

[Time delta from previous captured frame: 0.020346000 seconds]

[Time delta from previous displayed frame: 0.020346000 seconds]

[Time since reference or first frame: 0.020346000 seconds]

Frame Number: 2

Frame Length: 791 bytes

Capture Length: 791 bytes

[Frame is marked: False]

[Protocols in frame: eth: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: 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

À propos de cette traduction

Cisco a traduit ce document en traduction automatisée vérifiée par une personne dans le cadre d'un service mondial permettant à nos utilisateurs d'obtenir le contenu d'assistance dans leur propre langue.

Il convient cependant de noter que même la meilleure traduction automatisée ne sera pas aussi précise que celle fournie par un traducteur professionnel.