



## VXLAN Underlay Infrastructure Routing

[« FabricPath vPC+ | VXLAN Flood and Learn »](#)

Last updated: January 19, 2017

### VXLAN Underlay Infrastructure Routing

#### Task

- Disable the East/West links between the 5Ks, and the East/West links between the 7Ks.
- Configure IGP routing between your VXLAN Leafs (N5Ks) and Spines (N7Ks) as follows:
  - Configure Loopback0 interfaces with the IP addresses 10.0.0.X/32, where X is the device number.
  - Configure the links from the Leafs to the Spines as IP unnumbered to the Loopback0 interface.
  - Advertise the Loopback0 interfaces and all links between the Leafs and Spines into OSPF area 0.
  - Both the Leafs and Spines should generate syslog messages for any OSPF adjacency changes.
  - Once complete, you should have full IP reachability between all of the Loopback0 interfaces, and redundancy between the Leafs if one of the Spines goes down.
- Configure Multicast routing with Phantom RP between your VXLAN Leafs (N5Ks) and Spines (N7Ks) as follows:
  - Configure a new Loopback1 on your first Spine with the IP address 10.0.0.254/31, and advertise it into OSPF.
  - Configure a new Loopback1 on your second Spine with the IP address 10.0.0.254/30, and advertise it into OSPF.
  - Enable PIM Sparse Mode on the Loopback0 interfaces and all links between the Leafs and Spines, and the new Loopback1 interfaces of the Spines.
  - Configure the Leafs and the Spines to statically use the Bidirectional RP address of 10.0.0.255.
  - Once complete, join a multicast group (e.g. 224.1.1.1) on the Loopback0 interface of your Leafs, and then ping this address while sourcing the packet from the interface going to the first Spine. If your setup is working, you should get an ICMP echo-reply from the remote Leaf.
  - Next, disable the Loopback1 of your first Spine (simulating a node failure of the Spine), and ping the multicast group you joined on the Leafs while sourcing the packet from the interface to the second Spine. If your setup is working, failover should have automatically occurred to the secondary RP, and you should get an ICMP echo-reply from the remote Leaf.
- Configure Bidirectional Forwarding Detection (BFD) between your VXLAN Leafs (N5Ks) and Spines (N7Ks) as follows:
  - Enable BFD echos on the links between the Leafs and the Spines.
  - Generate BFD packets every 250ms, and wait 1 second before declaring the neighbor down.
  - Register both OSPF and PIM to BFD.
  - To verify this configuration, ping from your first Leaf to your first Spine with an unlimited packet count. Next, disable IP unnumbered on the link from the Spine to that Leaf. If successful, you should see the BFD neighbor immediately go down between the Leaf and Spine, traffic automatically re-route, and little or no packet loss should occur for the ping test that you have running. Once complete, re-enabled ip unnumbered on the Spine's link to the Leaf, and the BFD session should return.
- Configure BGP routing between your VXLAN Leafs (N5Ks) and Spines (N7Ks) as follows:
  - Configure BGP AS 65001 on your Leafs and Spines.
  - Both the Leafs and Spines should generate syslog messages for any BGP neighbor changes.
  - Peer iBGP from your Leafs to the Spines using your Loopback0 interfaces.
  - Enable the Address-Families IPv4 Unicast and L2VPN EVPN for these peerings.
  - The spines should be Route Reflectors for both of these AFIs.
  - Ensure to send Extended Communities between the Leafs and Spines for the L2VPN EVPN AFIs.
  - Once complete, you should see the peerings established with the `show bgp ipv4 unicast summary` and `show bgp l2vpn evpn summary`.

#### Configuration [Click to collapse](#)

```
N5K1:
nv overlay evpn
feature ospf
feature bgp
feature pim
feature bfd
!
ip pim rp-address 10.0.0.255 group-list 224.0.0.0/4 bidir
ip pim bfd
!
interface Ethernet1/3
shutdown

interface Ethernet1/4
shutdown

interface Ethernet1/5
no switchport
medium p2p
bfd interval 250 min_rx 250 multiplier 4
ip unnumbered loopback0
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode

interface Ethernet1/6
no switchport
medium p2p
bfd interval 250 min_rx 250 multiplier 4
ip unnumbered loopback0
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode

interface loopback0
ip address 10.0.0.51/32
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode
ip igmp join-group 224.1.1.1

router ospf 1
bfd
log-adjacency-changes

router bgp 65001
log-neighbor-changes
address-family l2vpn evpn
neighbor 10.0.0.71
remote-as 65001
```

```

update-source loopback0
address-family ipv4 unicast
address-family l2vpn evpn
  send-community extended
neighbor 10.0.0.72
remote-as 65001
update-source loopback0
address-family ipv4 unicast
address-family l2vpn evpn
  send-community extended

N5K2:
nv overlay evpn
feature ospf
feature bgp
feature pim
feature bfd
!
ip pim rp-address 10.0.0.255 group-list 224.0.0.0/4 bidir
ip pim bfd
!
interface Ethernet1/3
  shutdown

interface Ethernet1/4
  shutdown

interface Ethernet1/5
  no switchport
  medium p2p
  bfd interval 250 min_rx 250 multiplier 4
  ip unnumbered loopback0
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode

interface Ethernet1/6
  no switchport
  medium p2p
  bfd interval 250 min_rx 250 multiplier 4
  ip unnumbered loopback0
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode

interface loopback0
  ip address 10.0.0.52/32
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode
  ip igmp join-group 224.1.1.1

router ospf 1
  bfd
  log-adjacency-changes

router bgp 65001
  log-neighbor-changes
  address-family l2vpn evpn
  neighbor 10.0.0.71
    remote-as 65001
  update-source loopback0
  address-family ipv4 unicast
  address-family l2vpn evpn
    send-community extended
  neighbor 10.0.0.72
    remote-as 65001
  update-source loopback0
  address-family ipv4 unicast
  address-family l2vpn evpn
    send-community extended

N7K1:
nv overlay evpn
feature ospf
feature bgp
feature pim
feature bfd
ip pim rp-address 10.0.0.255 group-list 224.0.0.0/4 bidir

ip pim bfd

interface Ethernet1/1
  shutdown

interface Ethernet1/2
  shutdown

interface Ethernet1/3
  medium p2p
  bfd interval 250 min_rx 250 multiplier 4
  ip unnumbered loopback0
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/4
  medium p2p
  bfd interval 250 min_rx 250 multiplier 4
  ip unnumbered loopback0
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode
  no shutdown

```

```

interface loopback0
 ip address 10.0.0.71/32
 ip router ospf 1 area 0.0.0.0
 ip pim sparse-mode

interface loopback1
 ip address 10.0.0.254/31
 ip ospf network point-to-point
 ip router ospf 1 area 0.0.0.0
 ip pim sparse-mode

router ospf 1
 bfd
 log-adjacency-changes

router bgp 65001
 log-neighbor-changes
 address-family 12vpn evpn
 neighbor 10.0.0.51
 remote-as 65001
 update-source loopback0
 address-family ipv4 unicast
 send-community extended
 route-reflector-client
 address-family 12vpn evpn
 send-community extended
 route-reflector-client
 neighbor 10.0.0.52
 remote-as 65001
 update-source loopback0
 address-family ipv4 unicast
 send-community extended
 route-reflector-client
 address-family 12vpn evpn
 send-community extended
 route-reflector-client

N7K2:
nv overlay evpn
feature ospf
feature bgp
feature pim
feature bfd

ip pim rp-address 10.0.0.255 group-list 224.0.0.0/4 bidir

ip pim bfd

interface Ethernet1/1
 shutdown

interface Ethernet1/2
 shutdown

interface Ethernet1/3
 medium p2p
 bfd interval 250 min_rx 250 multiplier 4
 ip unnumbered loopback0
 ip router ospf 1 area 0.0.0.0
 ip pim sparse-mode
 no shutdown

interface Ethernet1/4
 medium p2p
 bfd interval 250 min_rx 250 multiplier 4
 ip unnumbered loopback0
 ip router ospf 1 area 0.0.0.0
 ip pim sparse-mode
 no shutdown

interface loopback0
 ip address 10.0.0.72/32
 ip router ospf 1 area 0.0.0.0
 ip pim sparse-mode

interface loopback1
 ip address 10.0.0.254/30
 ip ospf network point-to-point
 ip router ospf 1 area 0.0.0.0
 ip pim sparse-mode

router ospf 1
 bfd
 log-adjacency-changes

router bgp 65001
 log-neighbor-changes
 address-family 12vpn evpn
 neighbor 10.0.0.51
 remote-as 65001
 update-source loopback0
 address-family ipv4 unicast
 route-reflector-client
 address-family 12vpn evpn
 send-community extended
 route-reflector-client
 neighbor 10.0.0.52
 remote-as 65001
 update-source loopback0
 address-family ipv4 unicast
 route-reflector-client

```

```
address-family ipv4 vvpn
send-community extended
route-reflector-client
```

## Verification

OSPF neighbors are up on the unnumbered link, and the link is running OSPF network point-to-point via the command `medium p2p`.

```
N5K1# show ip interface brief
IP Interface Status for VRF "default"(1)
Interface      IP Address      Interface Status
Lo0            10.0.0.51      protocol-up/link-up/admin-up
Eth1/5        unnumbered
              (loopback0)   protocol-up/link-up/admin-up
Eth1/6        unnumbered
              (loopback0)   protocol-up/link-up/admin-up

N5K1# show ip ospf neighbors
OSPF Process ID 1 VRF default
Total number of neighbors: 2
Neighbor ID    Pri State      Up Time  Address      Interface
10.0.0.71     1 FULL/-      01:05:50 10.0.0.71    Eth1/5
10.0.0.72     1 FULL/-      01:34:57 10.0.0.72    Eth1/6

N5K1# show ip ospf interface e1/5
Ethernet1/5 is up, line protocol is up
Unnumbered interface using IP address of loopback0 (10.0.0.51)
Process ID 1 VRF default, area 0.0.0.0
Enabled by interface configuration
State P2P, Network type P2P, cost 4
BFD is enabled
Index 1, Transmit delay 1 sec
1 Neighbors, flooding to 1, adjacent with 1
Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
Hello timer due in 00:00:03
No authentication
Number of opaque link LSAs: 0, checksum sum 0
```

Each of the Leafs have Equal Cost Multipath (ECMP) routes to each other's Loopback0 networks, which will be used for the VXLAN Tunnel Endpoints (VTEPs) in later tasks.

```
N5K1# show ip route ospf
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>

10.0.0.52/32, ubest/mbest: 2/0
 *via 10.0.0.71, Eth1/5, [110/9], 01:07:20, ospf-1, intra
 *via 10.0.0.72, Eth1/6, [110/9], 01:07:20, ospf-1, intra
10.0.0.71/32, ubest/mbest: 1/0
 *via 10.0.0.71, Eth1/5, [110/5], 01:07:20, ospf-1, intra
 via 10.0.0.71, Eth1/5, [250/0], 01:35:43, am
10.0.0.72/32, ubest/mbest: 1/0
 *via 10.0.0.72, Eth1/6, [110/5], 01:36:22, ospf-1, intra
 via 10.0.0.72, Eth1/6, [250/0], 01:35:43, am
10.0.0.252/30, ubest/mbest: 1/0
 *via 10.0.0.72, Eth1/6, [110/5], 01:22:41, ospf-1, intra
10.0.0.254/31, ubest/mbest: 1/0
 *via 10.0.0.71, Eth1/5, [110/5], 01:07:20, ospf-1, intra
```

PIM is enabled on all interfaces running IGP, including the Loopback0 interfaces, which will be the source of the VTEPs.

```
N5K1# show ip pim interface brief
PIM Interface Status for VRF "default"
Interface      IP Address      PIM DR Address  Neighbor  Border
Count          Interface
Ethernet1/5    10.0.0.51      10.0.0.71       1         no
Ethernet1/6    10.0.0.51      10.0.0.72       1         no
loopback0      10.0.0.51      10.0.0.51       0         no
```

The Leafs should be PIM adjacent to each Spine.

```
N5K1# show ip pim neighbor
PIM Neighbor Status for VRF "default"
Neighbor      Interface      Uptime    Expires    DR      Bidir-  BFD
Priority      Capable State
10.0.0.71     Ethernet1/5    03:06:01  00:01:32  1       yes     Up
10.0.0.72     Ethernet1/6    03:06:01  00:01:36  1       yes     Up
```

The below Leaf is configured to listen for the multicast group 224.1.1.1.

```
N5K1# show ip igmp groups
IGMP Connected Group Membership for VRF "default" - 1 total entries
Type: S - Static, D - Dynamic, L - Local, T - SSM Translated
Group Address  Type Interface      Uptime    Expires    Last Reporter
224.1.1.1     L   loopback0        03:06:16  never     10.0.0.51
```

The remote Leaf is using the RP address 10.0.0.255 (the Phantom RP address) for group 224.1.1.1.

<https://t.me/learningnets>

```
N5K2# show ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP disabled
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None
```

```
RP: 10.0.0.255, (1),
uptime: 03:29:48 priority: 0,
RP-source: (local),
group ranges:
224.0.0.0/4 (bidir)
```

The reverse route (RPF) to the RP is via the first Spine, connected on Eth1/5.

```
N5K2# show ip route rpf 10.0.0.255
IP Route Table for VRF "default", RPF for multicast source
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

10.0.0.254/31, ubest/mbest: 1/0
 *via 10.0.0.71, Eth1/5, [110/5], 03:27:01, ospf-1, intra
```

ICMP echos sent out this interface to 224.1.1.1 are received and replied to via the remote Leaf.

```
N5K2# ping multicast 224.1.1.1 interface Ethernet1/5
PING 224.1.1.1 (224.1.1.1): 56 data bytes
64 bytes from 10.0.0.51: icmp_seq=0 ttl=253 time=1.394 ms
64 bytes from 10.0.0.51: icmp_seq=1 ttl=253 time=1.254 ms
64 bytes from 10.0.0.51: icmp_seq=2 ttl=253 time=1.144 ms
64 bytes from 10.0.0.51: icmp_seq=3 ttl=253 time=1.126 ms
64 bytes from 10.0.0.51: icmp_seq=4 ttl=253 time=1.037 ms
```

The primary RP, N7K1 in this case, is the root of the tree for (\*, 224.1.1.1), while the secondary RP only has a shorter match for 224.0.0.0/4.

```
N7K1# show ip mroute 224.1.1.1
IP Multicast Routing Table for VRF "default"

(*, 224.1.1.1/32), bidir, uptime: 03:33:31, pim ip
Incoming interface: loopback1, RPF nbr: 10.0.0.255
Outgoing interface list: (count: 3)
 Ethernet1/3, uptime: 03:19:49, pim
 Ethernet1/4, uptime: 03:31:01, pim
 loopback1, uptime: 03:31:02, pim, (RPF)

N7K2# show ip mroute 224.1.1.1
IP Multicast Routing Table for VRF "default"

(*, 224.0.0.0/4), bidir, uptime: 03:34:39, pim ip
Incoming interface: loopback1, RPF nbr: 10.0.0.255
Outgoing interface list: (count: 1)
 loopback1, uptime: 03:34:39, pim, (RPF)
```

When N7K1 fails, N7K2 takes over as the RP due to IGP convergence.

```
N7K1# config t
Enter configuration commands, one per line. End with CNTL/Z.
N7K1(config)# int lo1
N7K1(config-if)# shut

N5K2# show ip route rpf 10.0.0.255
IP Route Table for VRF "default", RPF for multicast source
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

10.0.0.252/30, ubest/mbest: 1/0
 *via 10.0.0.72, Eth1/6, [110/5], 03:38:52, ospf-1, intra
```

Leaf N5K2 now falls back to the shorter match (/30 vs. the previous /31) to reach the RP address, specifically out port Eth1/6.

```
N5K2# ping multicast 224.1.1.1 interface Ethernet1/6
PING 224.1.1.1 (224.1.1.1): 56 data bytes
64 bytes from 10.0.0.51: icmp_seq=0 ttl=253 time=1.309 ms
64 bytes from 10.0.0.51: icmp_seq=1 ttl=253 time=1.02 ms
64 bytes from 10.0.0.51: icmp_seq=2 ttl=253 time=1.039 ms
64 bytes from 10.0.0.51: icmp_seq=3 ttl=253 time=1.045 ms
64 bytes from 10.0.0.51: icmp_seq=4 ttl=253 time=1.035 ms

--- 224.1.1.1 ping multicast statistics ---
5 packets transmitted,
From member 10.0.0.51: 5 packets received, 0.00% packet loss
--- in total, 1 group member responded ---
```

The secondary RP, N7K2, now takes over as the root of the tree for (\*, 224.1.1.1).

```
N7K2# show ip mroute 224.1.1.1
IP Multicast Routing Table for VRF "default"

(*, 224.1.1.1/32), bidir, uptime: 00:00:31, pim ip
Incoming interface: loopback1, RPF nbr: 10.0.0.255
Outgoing interface list: (count: 3)
Ethernet1/3, uptime: 00:00:31, pim
loopback1, uptime: 00:00:31, pim, (RPF)
Ethernet1/4, uptime: 00:00:31, pim
```

Each of the Leafs are BFD adjacent with both of the Spines, and both OSPF and PIM are registered to BFD.

```
N5K1# show bfd neighbors details

OurAddr      NeighAddr    LD/RD          RH/RS          Holdown(mult)  State      Int
Vr-f
10.0.0.51    10.0.0.72    1090519042/1107296257 Up              6801(4)     Up         Eth1/
6
default

Session state is Up and using echo function with 250 ms interval
Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None
MinTxInt: 250000 us, MinRxInt: 2000000 us, Multiplier: 4
Received MinRxInt: 2000000 us, Received Multiplier: 4
Holdown (hits): 8000 ms (0), Hello (hits): 2000 ms (8435)
Rx Count: 7523, Rx Interval (ms) min/max/avg: 0/1754/1704 last: 1198 ms ago
Tx Count: 8435, Tx Interval (ms) min/max/avg: 1515/1515/1515 last: 1484 ms ago
Registered protocols: ospf pim
Uptime: 14698 days 5 hrs 11 mins 40 secs
Last packet: Version: 1          - Diagnostic: 0
              State bit: Up      - Demand bit: 0
              Poll bit: 0        - Final bit: 0
              Multiplier: 4      - Length: 24
              My Discr.: 1107296257 - Your Discr.: 1090519042
              Min tx interval: 250000 - Min rx interval: 2000000
              Min Echo interval: 50000 - Authentication bit: 0
Down reason: None, Reason not-hosted: None

OurAddr      NeighAddr    LD/RD          RH/RS          Holdown(mult)  State      Int
Vr-f
10.0.0.51    10.0.0.71    1090519043/1107296259 Up              6847(4)     Up         Eth1/
5
default

Session state is Up and using echo function with 250 ms interval
Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None
MinTxInt: 250000 us, MinRxInt: 2000000 us, Multiplier: 4
Received MinRxInt: 2000000 us, Received Multiplier: 4
Holdown (hits): 8000 ms (0), Hello (hits): 2000 ms (8255)
Rx Count: 8251, Rx Interval (ms) min/max/avg: 0/1524/1520 last: 1152 ms ago
Tx Count: 8255, Tx Interval (ms) min/max/avg: 1515/1515/1515 last: 1484 ms ago
Registered protocols: pim ospf
Uptime: 14698 days 5 hrs 6 mins 59 secs
Last packet: Version: 1          - Diagnostic: 0
              State bit: Up      - Demand bit: 0
              Poll bit: 0        - Final bit: 0
              Multiplier: 4      - Length: 24
              My Discr.: 1107296259 - Your Discr.: 1090519043
              Min tx interval: 250000 - Min rx interval: 2000000
              Min Echo interval: 50000 - Authentication bit: 0
Down reason: None, Reason not-hosted: None
```

Reconvergence is now tested on a soft failure of the connection between N5K1 and N7K1. Note that a hard failure, e.g. a link down event, would immediately cause reconvergence as soon as the carrier delay / link debounce timer expires, which is typically only a few milliseconds. BFD in this case is protecting against an upper layer failure that does not update the layer 1 link status.

N5K1 prefers to route directly to N7K1 to reach its Loopback0.

```
N5K1# show ip route 10.0.0.71
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>

10.0.0.71/32, ubest/mbest: 1/0
*via 10.0.0.71, Eth1/5, [110/5], 03:32:03, ospf-1, intra
via 10.0.0.71, Eth1/5, [250/0], 04:00:26, am
```

Ping test starts.

```
N5K1# ping 10.0.0.71 count unlimited
PING 10.0.0.71 (10.0.0.71): 56 data bytes
64 bytes from 10.0.0.71: icmp_seq=0 ttl=254 time=1.26 ms
64 bytes from 10.0.0.71: icmp_seq=1 ttl=254 time=0.972 ms
64 bytes from 10.0.0.71: icmp_seq=2 ttl=254 time=0.973 ms
64 bytes from 10.0.0.71: icmp_seq=3 ttl=254 time=0.972 ms
64 bytes from 10.0.0.71: icmp_seq=4 ttl=254 time=0.96 ms
<snip>
```

The soft failure is simulated by removing IP processing on N7K1's link to N5K1.

```
N7K1# config t
```

```
Enter configuration commands, one per line. End with CNTL/Z.
N7K1(config)# interface e1/3
N7K1(config-if)# no ip unnumbered lo0
N7K1(config-if)#
```

N5K1 detects the failure, and one ping is lost.

```
64 bytes from 10.0.0.71: icmp_seq=1861 ttl=254 time=0.871 ms
64 bytes from 10.0.0.71: icmp_seq=1862 ttl=254 time=0.807 ms
64 bytes from 10.0.0.71: icmp_seq=1863 ttl=254 time=0.855 ms
64 bytes from 10.0.0.71: icmp_seq=1864 ttl=254 time=0.874 ms
64 bytes from 10.0.0.71: icmp_seq=1865 ttl=254 time=0.852 ms
64 bytes from 10.0.0.71: icmp_seq=1866 ttl=254 time=0.87 ms
64 bytes from 10.0.0.71: icmp_seq=1867 ttl=254 time=0.856 ms
64 bytes from 10.0.0.71: icmp_seq=1868 ttl=254 time=0.871 ms
64 bytes from 10.0.0.71: icmp_seq=1869 ttl=254 time=0.935 ms
2017 Jan 4 21:10:15 N5K1 %OSPF-5-ADJCHANGE: ospf-1 [6829] Nbr 10.0.0.71 on Ethernet1/5 went DOWN
2017 Jan 4 21:10:15 N5K1 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519043 to neighbor 10.0.0.71 on interface Eth1/5 has gone down. Reason: Administratively Down.
2017 Jan 4 21:10:15 N5K1 %PIM-5-DR_CHANGE: pim [9068] DR change from 10.0.0.71 to 10.0.0.51 on interface Ethernet1/5
2017 Jan 4 21:10:15 N5K1 %BFD-5-SESSION_REMOVED: BFD session to neighbor 10.0.0.71 on interface Eth1/5 has been removed
Request 1870 timed out
64 bytes from 10.0.0.71: icmp_seq=1871 ttl=252 time=1.279 ms
64 bytes from 10.0.0.71: icmp_seq=1872 ttl=252 time=0.787 ms
64 bytes from 10.0.0.71: icmp_seq=1873 ttl=252 time=0.828 ms
<snip>
```

N5K1 now prefers to route to the second Spine to reach the first Spine.

```
N5K1# show ip route 10.0.0.71
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>

10.0.0.71/32, ubest/mbest: 1/0
   *via 10.0.0.72, Eth1/6, [110/13], 00:02:23, ospf-1, intra
```

BGP peerings are established both for AFI IPv4 Unicast and AFI L2VPN EVPN. Extended communities are exchanged for L2VPN EVPN, which is required to encode the Route Target of the L2VPN route, which ultimately determines the VXLAN Network Identifier (VNI) membership.

```
N5K1# show bgp ipv4 unicast summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 10.0.0.51, local AS number 65001
BGP table version is 5, IPv4 Unicast config peers 2, capable peers 2
0 network entries and 0 paths using 0 bytes of memory
BGP attribute entries [0/0], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [0/0]

Neighbor      V   AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.71     4 65001    220    223     5    0   0 03:27:12 0
10.0.0.72     4 65001    216    220     5    0   0 03:27:10 0

N5K1# show bgp l2vpn evpn summary
BGP summary information for VRF default, address family L2VPN EVPN
BGP router identifier 10.0.0.51, local AS number 65001
BGP table version is 4, L2VPN EVPN config peers 2, capable peers 2
0 network entries and 0 paths using 0 bytes of memory
BGP attribute entries [0/0], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [0/0]

Neighbor      V   AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.71     4 65001    220    223     4    0   0 03:27:15 0
10.0.0.72     4 65001    216    220     4    0   0 03:27:13 0

N5K1# show bgp l2vpn evpn neighbors 10.0.0.71 | include community
Extended community attribute sent to this neighbor
```



[« FabricPath vPC+ | VXLAN Flood and Learn »](#)

[^ back to top](#)