Hewlett Packard Enterprise

EVI 2.0 EVPN Data Center Interconnect (EVPN DCI) deployment guide

Contents

Deployment guide

Introduction

This deployment guide provides advice on deploying Comware switches with EVI (Ethernet Virtual Interconnect) 2.0 EVPN Data Center Interconnect (EVPN DCI). This solution supports active/active Data Centers (DC) by allowing L2/L3 networks to be extended across multiple DCs over any Wide Area Network (WAN) with IP connectivity while isolating spanning tree within each DC. As shown in the example in Figure 1, EVPN DCI will allow and provide network connectivity for servers, Virtual Machines (VMs) and applications on the 10.80.22.0/24 and 10.80.23.0/24 subnets across multiple DCs for high availability. EVPN DCI utilizes VXLAN as the data plane protocol and EVPN BGP as the control plane protocol.

Figure 1. EVI 2.0 EVPN DCI solution overview

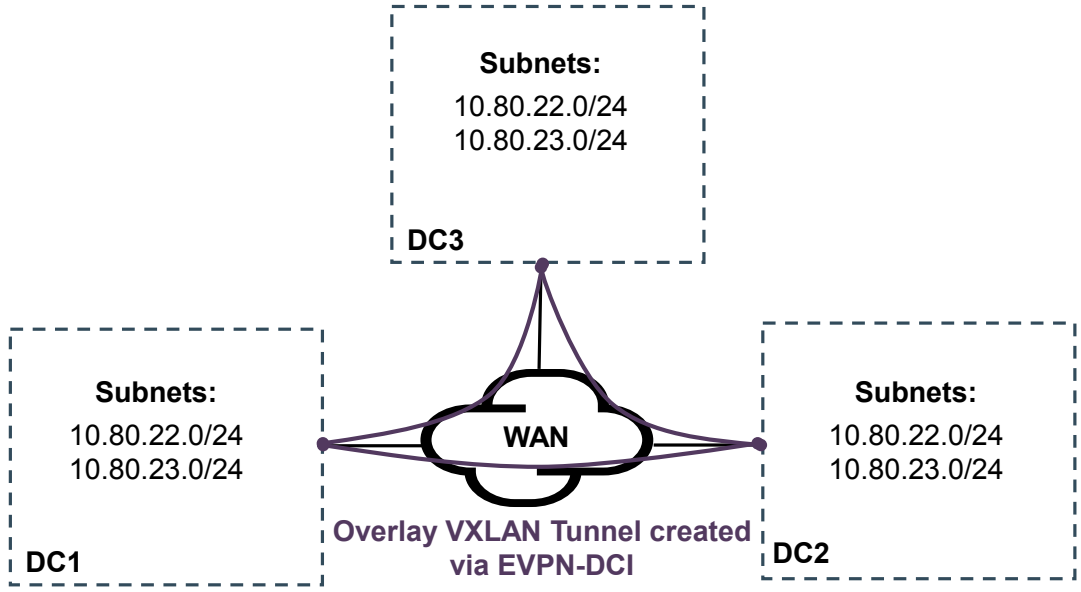

The sample configurations documented are based on FlexFabric 5940 “R2509P02” software but should also be applicable on other platforms and software that support EVPN-DCI.

3 use cases will be described in this document:

Use case #1 - Separate gateway and Edge Device (ED)

Use case #2 - Combined gateway and ED

Use case #3 - Intra-DC EVPN + EVPN DCI

In all use cases, it is recommended that EDs in each DC utilize two chassis Intelligent Resilient Framework (IRF) for multipathing, chassis high availability, single device configuration and single IP address benefits as shown in figure 2. When IRF EDs are deployed, Bridge Aggregation (BAGG) interfaces are typically used towards non ED devices within each DC, interfaces towards the WAN could either utilize BAGG, Route Aggregation (RAGG) interfaces or multiple IPs if required.

Figure 2. EVPN-DCI IRF EDs

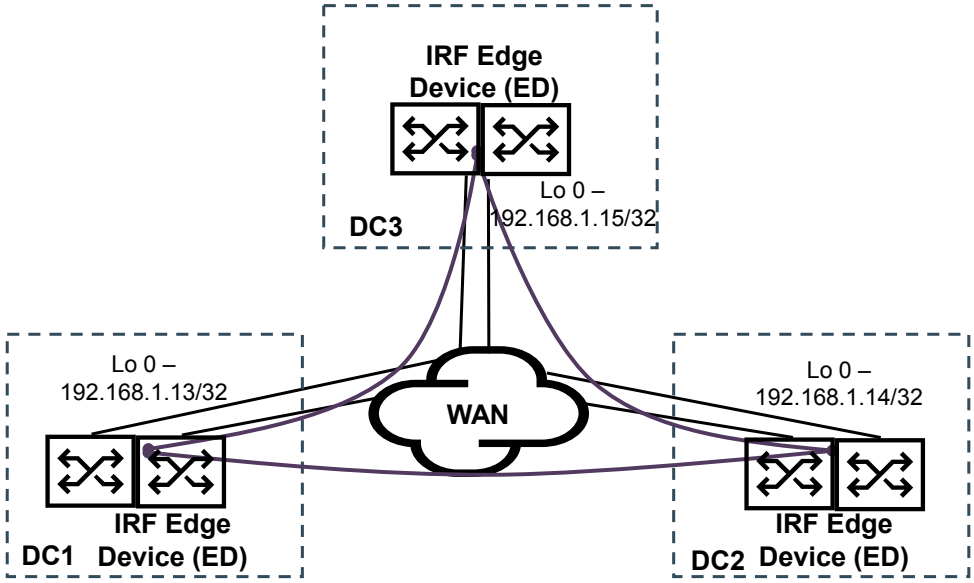

Use case #1 - separate gateway and ED

In this use case, we assume there is an existing DC1, however there now is a requirement to extend the existing L2/L3 networks to a new DC2.

As shown in figure 3, EDs in each DC would function as VXLAN Tunnel End Points (VTEPs) and connect to switches that function as the VRRP default gateway for servers and VMs. EDs would only bridge L2 traffic for the 10.80.22.0/24 and 10.80.23.0/24 subnets, there is no L3 routing functionality required for these two subnets on the EDs. To keep things simple, EDs in both DCs are placed in the same WAN BGP AS #100. Sample configurations are provided next to help understand this use-case. Additional DCs can be added using the same concepts.

Figure 3. Separate gateway and ED

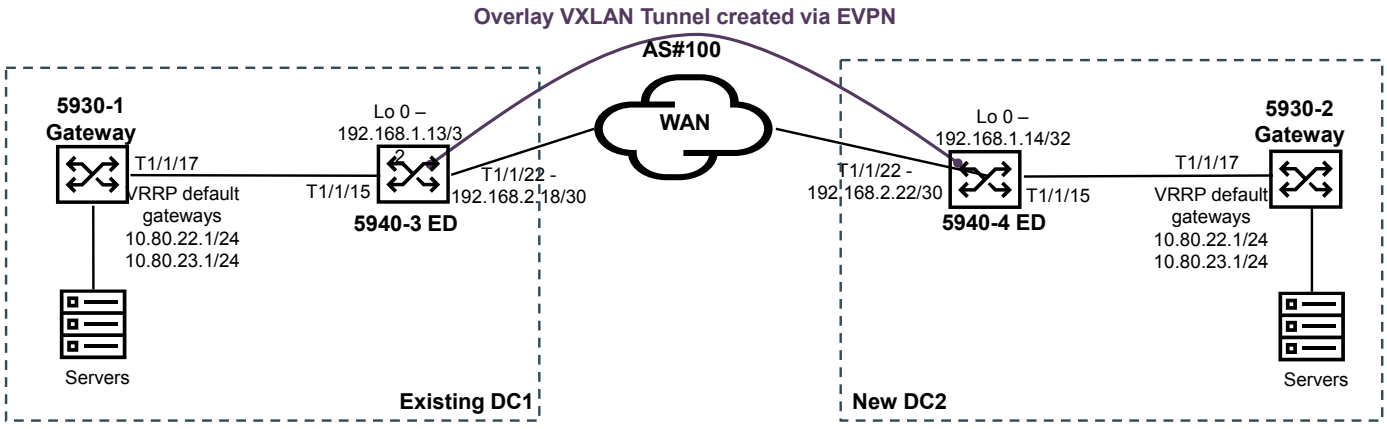

Sample 5930-1 gateway configuration

```
sysname 5930-1
#
vlan 22 to 23
#
interface Vlan-interface22
 ip address 10.80.22.2 255.255.255.0
 vrrp vrid 22 virtual-ip 10.80.22.1
#
interface Vlan-interface23
 ip address 10.80.23.2 255.255.255.0
 vrrp vrid 23 virtual-ip 10.80.23.1
#
interface Ten-GigabitEthernet1/1/17
 description To ED
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 22 to 23
#
interface Ten-GigabitEthernet1/1/18
 description To Server
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 22 to 23
```

After EVPN-DCI is configured on EDs, Master VRRP state can be confirmed using

```
<5930-1>dis vrrp
IPv4 Virtual Router Information:
 Running mode : Standard
 Total number of virtual routers : 2
 Interface          VRID  State        Running Adver   Auth     Virtual
                                        Pri     Timer   Type        IP
 ------------------------------------------------------------------------
 Vlan22             22    Master       100     100     None     10.80.22.1
 Vlan23             23    Master       100     100     None     10.80.23.1
```

After EVPN-DCI is configured on EDs, connectivity to 5930-2 should also be available.

```
<5930-1>ping 10.80.22.3
Ping 10.80.22.3 (10.80.22.3): 56 data bytes, press CTRL_C to break
56 bytes from 10.80.22.3: icmp_seq=0 ttl=255 time=1.491 ms
56 bytes from 10.80.22.3: icmp_seq=1 ttl=255 time=0.957 ms
56 bytes from 10.80.22.3: icmp_seq=2 ttl=255 time=1.098 ms
56 bytes from 10.80.22.3: icmp_seq=3 ttl=255 time=1.268 ms
56 bytes from 10.80.22.3: icmp_seq=4 ttl=255 time=0.986 ms
```

```
<5930-1>ping 10.80.23.3
Ping 10.80.23.3 (10.80.23.3): 56 data bytes, press CTRL_C to break
56 bytes from 10.80.23.3: icmp_seq=0 ttl=255 time=1.390 ms
56 bytes from 10.80.23.3: icmp_seq=1 ttl=255 time=1.338 ms
56 bytes from 10.80.23.3: icmp_seq=2 ttl=255 time=1.195 ms
56 bytes from 10.80.23.3: icmp_seq=3 ttl=255 time=1.131 ms
```

Sample 5940-3 ED configuration

```
sysname 5940-3
#
 vtep enable
 l2vpn enable
 stp global enable
#
# Loopback IP address is used as VXLAN tunnel source and destination address (from opposite DC)
 tunnel global source-address 192.168.1.13
#
 vxlan tunnel mac-learning disable
#
 router id 192.168.1.13
#
# OSPF is used for underlay connectivity, replace with other routing protocols if required
ospf 1
 area 0.0.0.0
  network 192.168.0.0 0.0.255.255
#
 lldp global enable
#
 system-working-mode standard
#
vlan 22 to 23
#
vsi vpna-22
 vxlan 22
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
#
vsi vpna-23
 vxlan 23
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
```

```
#
interface LoopBack0
 ip address 192.168.1.13 255.255.255.255
#
interface Ten-GigabitEthernet1/1/22
 port link-mode route
 description WAN
 mtu 9000
 ip address 192.168.2.18 255.255.255.252
 dci enable
#
# ACL is used to block VRRP traffic between DCs
acl mac 4000
 rule 10 deny type 0800 ffff dest-mac 0100-5e00-0012 ffff-ffff-ffff
 rule 100 permit
#
interface Ten-GigabitEthernet1/1/15
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 22 to 23
 packet-filter mac 4000 inbound
 packet-filter mac 4000 outbound
 service-instance 22
  encapsulation s-vid 22
  xconnect vsi vpna-22
 service-instance 23
  encapsulation s-vid 23
  xconnect vsi vpna-23
#
interface Vsi-interface1000
 l3-vni 1000
#
# EVPN peering to ED in other DC
bgp 100
 peer 192.168.1.14 as-number 100
 peer 192.168.1.14 connect-interface LoopBack0
 #
 address-family l2vpn evpn
  peer 192.168.1.14 enable
  peer 192.168.1.14 router-mac-local
```

After the ED in DC2 is configured with EVPN-DCI, remote MACs should be seen in the “l2vpn mac” and “bgp l2vpn evpn” tables.

```
[5940-3]dis l2vpn mac
MAC Address    State    VSI Name                        Link ID/Name    Aging
0000-5e00-0116 EVPN     vpna-22                         Tunnel0         NotAging
2c23-3a5a-b789 EVPN     vpna-22                         Tunnel0         NotAging
2c23-3ae7-da61 Dynamic  vpna-22                         XGE1/1/15       Aging
0000-5e00-0117 EVPN     vpna-23                         Tunnel0         NotAging
2c23-3a5a-b78a EVPN     vpna-23                         Tunnel0         NotAging
2c23-3ae7-da62 Dynamic  vpna-23                         XGE1/1/15       Aging
--- 6 mac address(es) found  ---
```

This command is used in the DC1 ED to locate a server in DC2 with MAC address of 2c23-3a5a-b789, it shows this MAC address is advertised by VTEP 192.168.1.14 (DC2 ED).

```
[5940-3]dis bgp l2vpn evpn | b 2c23-3a5a-b789
* >i [2][0][48][2c23-3a5a-b789][0][0.0.0.0]/104
                        192.168.1.14    0          100        0       i
* >  [2][0][48][2c23-3ae7-da61][0][0.0.0.0]/104
                        0.0.0.0         0          100        32768   i
* >  [3][0][32][192.168.1.13]/80
                        0.0.0.0         0          100        32768   i
* >i [3][0][32][192.168.1.14]/80
                        192.168.1.14    0          100        0       i
```

Sample 5940-4 ED configuration

```
sysname 5940-4
#
 vtep enable
 l2vpn enable
 stp global enable
#
 tunnel global source-address 192.168.1.14
#
 vxlan tunnel mac-learning disable
#
 router id 192.168.1.14
#
ospf 1
 area 0.0.0.0
  network 192.168.0.0 0.0.255.255
#
 lldp global enable
#
 system-working-mode standard
#
```

```
vlan 22 to 23
#
vsi vpna-22
 vxlan 22
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
#
vsi vpna-23
 vxlan 23
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
#
interface LoopBack0
 ip address 192.168.1.14 255.255.255.255
#
interface Ten-GigabitEthernet1/1/22
 port link-mode route
 description WAN
 mtu 9000
 ip address 192.168.2.22 255.255.255.252
 dci enable
#
acl mac 4000
 rule 10 deny type 0800 ffff dest-mac 0100-5e00-0012 ffff-ffff-ffff
 rule 100 permit
#
interface Ten-GigabitEthernet1/1/15
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 22 to 23
 packet-filter mac 4000 inbound
 packet-filter mac 4000 outbound
 service-instance 22
  encapsulation s-vid 22
  xconnect vsi vpna-22
 service-instance 23
  encapsulation s-vid 23
  xconnect vsi vpna-23
#
interface Vsi-interface1000
 l3-vni 1000
#
```

```
bgp 100
 peer 192.168.1.13 as-number 100
 peer 192.168.1.13 connect-interface LoopBack0
 #
 address-family l2vpn evpn
  peer 192.168.1.13 enable
  peer 192.168.1.13 router-mac-local
```

After the ED in DC1 is configured with EVPN-DCI, remote MACs should be seen in the "l2vpn mac" and "bgp l2vpn evpn" tables.

```
<5940-4>dis l2vpn mac
MAC Address    State    VSI Name                        Link ID/Name    Aging
0000-5e00-0116 Dynamic  vpna-22                         XGE1/1/15       Aging
2c23-3a5a-b789 Dynamic  vpna-22                         XGE1/1/15       Aging
2c23-3ae7-da61 EVPN     vpna-22                         Tunnel0         NotAging
0000-5e00-0117 Dynamic  vpna-23                         XGE1/1/15       Aging
2c23-3a5a-b78a Dynamic  vpna-23                         XGE1/1/15       Aging
2c23-3ae7-da62 EVPN     vpna-23                         Tunnel0         NotAging
--- 6 mac address(es) found  ---


<5940-4>dis bgp l2vpn evpn | b 2c23-3ae7-da61
* >i [2][0][48][2c23-3ae7-da61][0][0.0.0.0]/104
                         192.168.1.13    0          100        0       i
* >i [3][0][32][192.168.1.13]/80
                         192.168.1.13    0          100        0       i
* >  [3][0][32][192.168.1.14]/80
                         0.0.0.0         0          100        32768   i
```

Sample 5930-2 gateway configuration

```
sysname 5930-2
#
vlan 22 to 23
#
interface Vlan-interface22
 ip address 10.80.22.3 255.255.255.0
 vrrp vrid 22 virtual-ip 10.80.22.1
#
interface Vlan-interface23
 ip address 10.80.23.3 255.255.255.0
 vrrp vrid 23 virtual-ip 10.80.23.1
#
interface Ten-GigabitEthernet1/1/17
 description To ED
 port link-mode bridge
 port link-type trunk
```

```
 port trunk permit vlan 1 22 to 23
#
interface Ten-GigabitEthernet1/1/18
 description To Server
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 22 to 23
```

After the 5940s are configured, Master VRRP state can be confirmed using

```
<5930-2>dis vrrp
IPv4 Virtual Router Information:
 Running mode : Standard
 Total number of virtual routers : 2
 Interface          VRID  State        Running Adver   Auth     Virtual
                                        Pri     Timer   Type        IP
 ------------------------------------------------------------------------
 Vlan22             22    Master       100     100     None     10.80.22.1
 Vlan23             23    Master       100     100     None     10.80.23.1
```

Alternative configuration for use case #1 – single VSI

In the previous use case, we used different VSIs for different VLANs, each VSI typically represents a “tenant”, different VSIs allow for multi-tenancy. It is also possible to place different VLANs into the same VSI if desired.

The majority of the previous configs shown in use case #1 EDs would still apply; the only changes required on both 5940-3 and 5940-4 would be these.

```
vsi vpna
 vxlan 21
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
#
interface Ten-GigabitEthernet1/1/15
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 22 to 23
 packet-filter mac 4000 inbound
 packet-filter mac 4000 outbound
 service-instance 21
  encapsulation s-vid 22 to 23
  xconnect vsi vpna
```

After the ED in DC2 is configured with EVPN-DCI, remote MACs should be seen in the “l2vpn mac” and “bgp l2vpn evpn” tables.

```
[5940-3]dis l2vpn mac
MAC Address    State    VSI Name                        Link ID/Name    Aging
```

```
0000-5e00-0116 Dynamic  vpna                                XGE1/1/15       Aging
0000-5e00-0117 Dynamic  vpna                                XGE1/1/15       Aging
2c23-3a5a-b789 EVPN     vpna                                Tunnel0         NotAging
2c23-3a5a-b78a EVPN     vpna                                Tunnel0         NotAging
2c23-3ae7-da61 Dynamic  vpna                                XGE1/1/15       Aging
2c23-3ae7-da62 Dynamic  vpna                                XGE1/1/15       Aging
--- 6 mac address(es) found  ---

[5940-3]dis bgp l2vp evpn | b 2c23-3a5a-b789
* >i [2][0][48][2c23-3a5a-b789][0][0.0.0.0]/104
                        192.168.1.14    0          100        0       i
* >i [2][0][48][2c23-3a5a-b78a][0][0.0.0.0]/104
                        192.168.1.14    0          100        0       i
```

Use case #2 - combined gateway and ED

In this use case, we assume two new small scale DCs are being deployed.

As shown in figure 4, EDs in each DC would also function as the default gateway for servers and VMs. EDs would bridge L2 and route L3 traffic for the 10.80.22.0/24 and 10.80.23.0/24 subnets. Each DC has a unique BGP AS number to cater for expansion to use case #3. The same default gateway IPs are active in both DCs without VRRP. ARP suppression is also enabled in this use case to minimize ARP broadcasts across the WAN. Sample configurations are provided next to help understand this use-case. Additional DCs can be added using the same concepts.

Figure 4. Combined gateway and ED

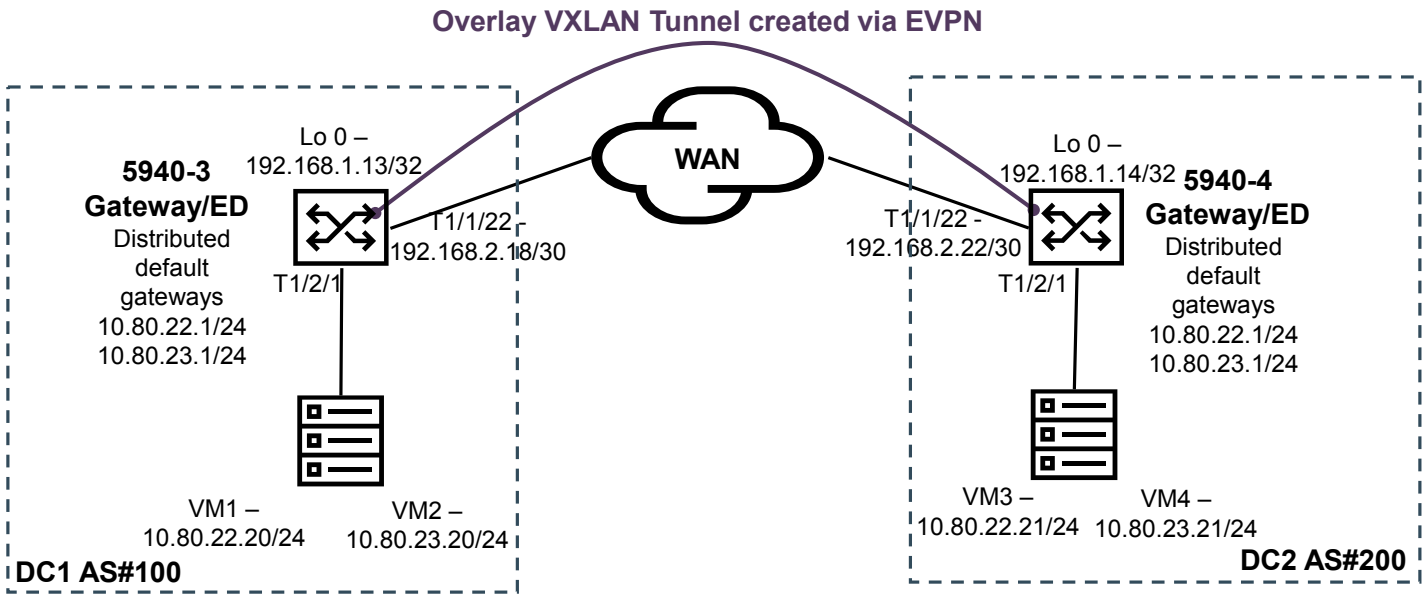

Sample 5940-3 ED configuration

```
sysname 5940-3
#
l2vpn enable
vxlan tunnel mac-learning disable
vxlan tunnel arp-learning disable
#
hardware-resource vxlan border24k
#
interface LoopBack0
 ip address 192.168.1.13 255.255.255.255
#
tunnel global source-address 192.168.1.13
#
vsi vpna-22
 arp suppression enable
 vxlan 22
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target 1:22 both
# Manual vpn-target is required as DCs are in different AS#s
#
vsi vpna-23
 arp suppression enable
 vxlan 23
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target 1:23 both
# Manual vpn-target is required as DCs are in different AS#s
#
bgp 100
 peer 192.168.1.14 as-number 200
 peer 192.168.1.14 connect-interface loopback 0
 peer 192.168.1.14 ebgp-max-hop 64
 address-family l2vpn evpn
  peer 192.168.1.14 enable
  peer 192.168.1.14 router-mac-local
#
interface ten-gigabitethernet 1/2/1
 service-instance 22
  encapsulation s-vid 22
  xconnect vsi vpna-22
 service-instance 23
  encapsulation s-vid 23
  xconnect vsi vpna-23
#
```

```
ip vpn-instance vpna
 route-distinguisher 1:1
 address-family ipv4
 vpn-target 2:2
 address-family evpn
 vpn-target 1:1
#
interface vsi-interface 22
 ip binding vpn-instance vpna
 ip address 10.80.22.1 255.255.255.0
 mac-address 1-22-1
 distributed-gateway local
#
interface vsi-interface 23
 ip binding vpn-instance vpna
 ip address 10.80.23.1 255.255.255.0
 mac-address 1-23-1
 distributed-gateway local
#
# Assign L3 gateway interface into VSI
vsi vpna-22
 gateway vsi-interface 22
#
vsi vpna-23
 gateway vsi-interface 23
#
interface vsi-interface 1000
 ip binding vpn-instance vpna
 l3-vni 1000
#
interface ten-gigabitethernet 1/1/22
 port link-mode route
 description WAN
 mtu 9000
 ip address 192.168.2.18 255.255.255.252
 dci enable
#
ospf 1
 area 0.0.0.0
  network 192.168.0.0 0.0.255.255
# OSPF is used for underlay connectivity, replace with other routing protocols if required
```

After the ED in DC2 is configured with EVPN-DCI, host routes of the remote servers/VMs should be seen in the “vpn” table.

```
[5940-3]dis ip ro vpn vpna

Destinations : 18       Routes : 18

Destination/Mask   Proto   Pre Cost        NextHop         Interface
0.0.0.0/32         Direct  0   0           127.0.0.1       InLoop0
10.80.22.0/24      Direct  0   0           10.80.22.1      Vsi22
10.80.22.0/32      Direct  0   0           10.80.22.1      Vsi22
10.80.22.1/32      Direct  0   0           127.0.0.1       InLoop0
10.80.22.21/32     BGP     255 0           192.168.1.14    Vsi1000
10.80.22.255/32    Direct  0   0           10.80.22.1      Vsi22
10.80.23.0/24      Direct  0   0           10.80.23.1      Vsi23
10.80.23.0/32      Direct  0   0           10.80.23.1      Vsi23
10.80.23.1/32      Direct  0   0           127.0.0.1       InLoop0
10.80.23.21/32     BGP     255 0           192.168.1.14    Vsi1000
```

The “VPN ARP” table should show locally connected servers

```
[5940-3]display arp vpn-instance vpna
  Type: S-Static   D-Dynamic   O-Openflow   R-Rule   M-Multiport  I-Invalid
IP address      MAC address    VID        Interface/Link ID        Aging Type
10.80.22.20     0050-569a-5d30 1          0x0                      19    D
10.80.23.20     0050-569a-3888 2          0x0                      19    D
192.168.1.14    d894-0322-1ec9 3          Tunnel0                  N/A   R
```

Local and remote MACs should be seen in the “l2vpn mac” and “bgp l2vpn evpn” tables.

```
<5940-3>dis l2vpn mac
MAC Address    State    VSI Name                        Link ID/Name    Aging
0050-569a-5d30 Dynamic  vpna-22                         XGE1/2/1        Aging
0050-569a-673a EVPN     vpna-22                         Tunnel0         NotAging
0050-569a-3888 Dynamic  vpna-23                         XGE1/2/1        Aging
0050-569a-6b5e EVPN     vpna-23                         Tunnel0         NotAging
--- 4 mac address(es) found  ---

<5940-3>dis bgp l2vpn evpn | b 0050-569a-673a
* >e [2][0][48][0050-569a-673a][32][10.80.22.21]/136
                        192.168.1.14    0                    0       200i
* >e [2][0][48][0050-569a-6b5e][32][10.80.23.21]/136
192.168.1.14    0                    0       200i
```

ARP suppression can be verified by using this command

```
[5940-3]dis arp suppression vsi
IP address      MAC address    Vsi Name                        Link ID    Aging
10.80.22.21     0050-569a-673a vpna-22                         0x5000000  N/A
```

```
10.80.22.20     0050-569a-5d30 vpna-22                          0x0        25
10.80.23.21     0050-569a-6b5e vpna-23                          0x5000000  N/A
10.80.23.20     0050-569a-3888 vpna-23                          0x0        25
```

Sample 5940-4 ED configuration

```
sysname 5940-4
#
l2vpn enable
vxlan tunnel mac-learning disable
vxlan tunnel arp-learning disable
#
hardware-resource vxlan border24k
#
interface LoopBack0
 ip address 192.168.1.14 255.255.255.255
#
tunnel global source-address 192.168.1.14
#
vsi vpna-22
 arp suppression enable
 vxlan 22
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target 1:22 both
# Manual vpn-target is required as DCs are in different AS#s
#
vsi vpna-23
 arp suppression enable
 vxlan 23
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target 1:23 both
# Manual vpn-target is required as DCs are in different AS#s
#
bgp 200
 peer 192.168.1.13 as-number 100
 peer 192.168.1.13 connect-interface loopback 0
 peer 192.168.1.13 ebgp-max-hop 64
 address-family l2vpn evpn
  peer 192.168.1.13 enable
  peer 192.168.1.13 router-mac-local
#
interface ten-gigabitethernet 1/2/1
 service-instance 22
  encapsulation s-vid 22
  xconnect vsi vpna-22
```

```
 service-instance 23
  encapsulation s-vid 23
  xconnect vsi vpna-23
#
ip vpn-instance vpna
 route-distinguisher 1:1
 address-family ipv4
 vpn-target 2:2
 address-family evpn
 vpn-target 1:1
#
interface vsi-interface 22
 ip binding vpn-instance vpna
 ip address 10.80.22.1 255.255.255.0
 mac-address 1-22-1
 distributed-gateway local
#
interface vsi-interface 23
 ip binding vpn-instance vpna
 ip address 10.80.23.1 255.255.255.0
 mac-address 1-23-1
 distributed-gateway local
#
vsi vpna-22
 gateway vsi-interface 22
#
vsi vpna-23
 gateway vsi-interface 23
#
interface vsi-interface 1000
 ip binding vpn-instance vpna
 l3-vni 1000
#
interface ten-gigabitethernet 1/1/22
 port link-mode route
 description WAN
 mtu 9000
 ip address 192.168.2.22 255.255.255.252
 dci enable
#
ospf 1
 area 0.0.0.0
  network 192.168.0.0 0.0.255.255
```

After the ED in DC2 is configured with EVPN-DCI, host routes of the remote servers/VMs should be seen in the “vpn” table.

```
[5940-4]dis ip ro vpn vpna

Destinations : 18       Routes : 18

Destination/Mask   Proto   Pre Cost        NextHop         Interface
10.80.22.0/24      Direct  0   0           10.80.22.1      Vsi22
10.80.22.0/32      Direct  0   0           10.80.22.1      Vsi22
10.80.22.1/32      Direct  0   0           127.0.0.1       InLoop0
10.80.22.20/32     BGP     255 0           192.168.1.13    Vsi1000
10.80.23.0/24      Direct  0   0           10.80.23.1      Vsi23
10.80.23.0/32      Direct  0   0           10.80.23.1      Vsi23
10.80.23.1/32      Direct  0   0           127.0.0.1       InLoop0
10.80.23.20/32     BGP     255 0           192.168.1.13    Vsi1000
```

The “VPN ARP” table should show locally connected servers

```
[5940-4]display arp vpn-instance vpna
  Type: S-Static   D-Dynamic   O-Openflow   R-Rule   M-Multiport  I-Invalid
IP address      MAC address    VID        Interface/Link ID        Aging Type
10.80.23.21     0050-569a-6b5e 1          0x0                      11    D
10.80.22.21     0050-569a-673a 0          0x0                      19    D
192.168.1.13    d894-0322-1e58 2          Tunnel0                  N/A   R
```

Local and remote MACs should be seen in the “l2vpn mac” and “bgp l2vpn evpn” tables.

```
<5940-4>dis l2vpn mac
MAC Address    State    VSI Name                        Link ID/Name    Aging
0050-569a-5d30 EVPN     vpna-22                         Tunnel0         NotAging
0050-569a-673a Dynamic  vpna-22                         XGE1/2/1        Aging
0050-569a-3888 EVPN     vpna-23                         Tunnel0         NotAging
0050-569a-6b5e Dynamic  vpna-23                         XGE1/2/1        Aging
--- 4 mac address(es) found  ---

<5940-4>dis bgp l2vpn evpn | b 0050-569a-5d30
* >e [2][0][48][0050-569a-5d30][32][10.80.22.20]/136
                       192.168.1.13    0                    0       100i
```

ARP suppression can be verified by using this command

```
[5940-4]dis arp suppression vsi
IP address      MAC address    Vsi Name                        Link ID    Aging
10.80.22.20     0050-569a-5d30 vpna-22                         0x5000000  N/A
10.80.22.21     0050-569a-673a vpna-22                         0x0        25
10.80.23.20     0050-569a-3888 vpna-23                         0x5000000  N/A
10.80.23.21     0050-569a-6b5e vpna-23                         0x0        25
```

Use case #3 - Intra-DC EVPN + EVPN DCI

In this use case, as shown in figure 5, we assume use case # 2 needs to scale up to accommodate more switches. IP un-numbered interfaces are used between switches in the same DC, BGP Route Reflector (RRs) are added in each DC to scale and ease intra-DC EVPN deployment. Two RRs are recommended in each DC. The default gateway IPs are now distributed across multiple leaf switches in both DCs while EDs in each DC are dedicated for DCI. The IGP protocol used within each DC is isolated and not extended across the WAN to create separate failure domains, this also means Leaf VTEPs in each DC will not need to know about VTEP loopback IPs in other DCs to save on route table resources. Sample configurations are provided next to help understand this use-case. Additional DCs can be added using the same concepts.

Figure 5. Intra-DC EVPN with RRs in each DC + EVPN DCI

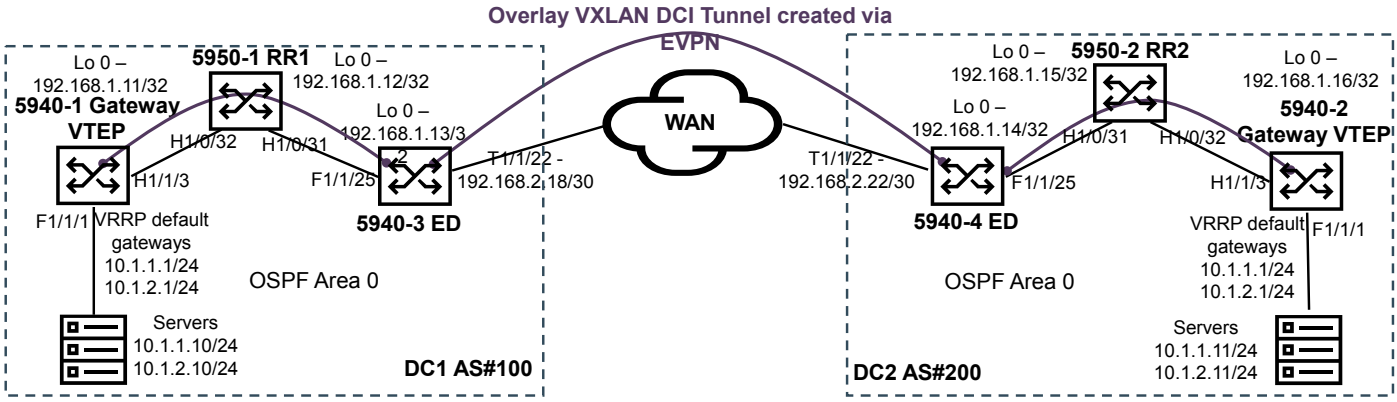

Sample 5940-1 gateway configuration

```
sysname 5940-1
#
l2vpn enable
vxlan tunnel arp-learning disable
vxlan tunnel mac-learning disable
#
hardware-resource vxlan l3gw24k
#
# 2 VSIs are created as 1 gateway interface can only be attached to 1 VSI
vsi vpna-1
 arp suppression enable
 vxlan 10
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
#
vsi vpna-2
 arp suppression enable
 vxlan 20
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
#
```

```
ospf 1
 area 0.0.0.0
  network 192.168.1.0 0.0.0.255
#
interface LoopBack0
 ip address 192.168.1.11 255.255.255.255
#
tunnel global source-address 192.168.1.11
#
interface FortyGigE1/1/1
 port link-mode bridge
 service-instance 1
  encapsulation s-vid 10
  xconnect vsi vpna-1
 service-instance 2
  encapsulation s-vid 11
  xconnect vsi vpna-2
#
interface HundredGigE1/1/3
 port link-mode route
 ip address unnumbered interface LoopBack0
 ospf network-type p2p
 lldp management-address arp-learning
 lldp tlv-enable basic-tlv management-address-tlv interface LoopBack0
#
ip vpn-instance vpna
 route-distinguisher 1:1
 #
 address-family evpn
  vpn-target 1:1 import-extcommunity
  vpn-target 1:1 export-extcommunity
#
interface Vsi-interface1
 ip binding vpn-instance vpna
 ip address 10.1.1.1 255.255.255.0
 mac-address 0001-0001-0001
 local-proxy-arp enable
 distributed-gateway local
# local-proxy-arp is used to respond back on all ARP requests, the EVPN table will be used
# to determine if the remote MAC addresses are known
#
interface Vsi-interface2
 ip binding vpn-instance vpna
 ip address 10.1.2.1 255.255.255.0
 mac-address 0002-0002-0002
 local-proxy-arp enable
```

```
 distributed-gateway local
#
vsi vpna-1
 gateway vsi-interface 1
#
vsi vpna-2
 gateway vsi-interface 2
#
interface Vsi-interface1000
 ip binding vpn-instance vpna
 l3-vni 1000
#
bgp 100
 peer 192.168.1.12 as-number 100
 peer 192.168.1.12 connect-interface LoopBack0
 #
 address-family l2vpn evpn
  peer 192.168.1.12 enable
# EVPN peering only to RR
```

After all other switches are configured correctly, the locally connected MAC address can be seen using

```
[5940-1]dis l2vpn mac
MAC Address    State    VSI Name                        Link ID/Name    Aging
7848-59ed-0ac8 Dynamic  vpna-1                          FGE1/1/1        Aging
7848-59ed-0ac9 Dynamic  vpna-2                          FGE1/1/1        Aging
--- 2 mac address(es) found  ---
```

Remote server IPs can be seen this way

```
<5940-1>dis ip ro vpn vpna

Destinations : 18       Routes : 18

Destination/Mask   Proto   Pre Cost        NextHop         Interface
10.1.1.0/24        Direct  0   0           10.1.1.1        Vsi1
10.1.1.0/32        Direct  0   0           10.1.1.1        Vsi1
10.1.1.1/32        Direct  0   0           127.0.0.1       InLoop0
10.1.1.11/32       BGP     255 0           192.168.1.13    Vsi1000
10.1.2.0/24        Direct  0   0           10.1.2.1        Vsi2
10.1.2.0/32        Direct  0   0           10.1.2.1        Vsi2
10.1.2.1/32        Direct  0   0           127.0.0.1       InLoop0
10.1.2.11/32       BGP     255 0           192.168.1.13    Vsi1000
```

While remote server MAC/IPs can be identified via

```
<5940-1>dis bgp l2vpn evpn

 BGP local router ID is 192.168.1.11
 Status codes: * - valid, > - best, d - dampened, h - history,
               s - suppressed, S - stale, i - internal, e - external
               Origin: i - IGP, e - EGP, ? - incomplete

 Total number of routes from all PEs: 8

 Route distinguisher: 1:1(vpna)
 Total number of routes: 4

     Network            NextHop         MED        LocPrf     PrefVal Path/Ogn

* >i [2][0][48][7848-59ed-0ad2][32][10.1.1.11]/136
                        192.168.1.13               100        0       200i
* >i [2][0][48][7848-59ed-0ad3][32][10.1.2.11]/136
                        192.168.1.13               100        0       200i
```

You will notice the only underlay OSPF routes learnt are from switch loopback IPs in the same DC, there are no routes required from the other DC.

```
<5940-1>dis ip ro

Destinations : 16       Routes : 16

Destination/Mask   Proto   Pre Cost        NextHop         Interface
0.0.0.0/32         Direct  0   0           127.0.0.1       InLoop0
10.10.10.0/24      Direct  0   0           10.10.10.173    MGE0/0/0
10.10.10.0/32      Direct  0   0           10.10.10.173    MGE0/0/0
10.10.10.173/32    Direct  0   0           127.0.0.1       InLoop0
10.10.10.255/32    Direct  0   0           10.10.10.173    MGE0/0/0
127.0.0.0/8        Direct  0   0           127.0.0.1       InLoop0
127.0.0.0/32       Direct  0   0           127.0.0.1       InLoop0
127.0.0.1/32       Direct  0   0           127.0.0.1       InLoop0
127.255.255.255/32 Direct  0   0           127.0.0.1       InLoop0
192.168.1.11/32    Direct  0   0           127.0.0.1       InLoop0
192.168.1.12/32    O_INTRA 10  1           192.168.1.12    HGE1/1/3
192.168.1.13/32    O_INTRA 10  2           192.168.1.12    HGE1/1/3
224.0.0.0/4        Direct  0   0           0.0.0.0         NULL0
224.0.0.0/24       Direct  0   0           0.0.0.0         NULL0
255.255.255.255/32 Direct  0   0           127.0.0.1       InLoop0
```

Sample 5950-1 RR1 configuration

```
sysname 5950-1
#
ospf 1
 area 0.0.0.0
  network 192.168.1.0 0.0.0.255
#
interface LoopBack0
 ip address 192.168.1.12 255.255.255.255
#
interface HundredGigE1/0/31
 port link-mode route
 ip address unnumbered interface LoopBack0
 ospf network-type p2p
 lldp management-address arp-learning
 lldp tlv-enable basic-tlv management-address-tlv interface LoopBack0
#
interface HundredGigE1/0/32
 port link-mode route
 ip address unnumbered interface LoopBack0
 ospf network-type p2p
 lldp management-address arp-learning
 lldp tlv-enable basic-tlv management-address-tlv interface LoopBack0
#
bgp 100
 peer 192.168.1.11 as-number 100
 peer 192.168.1.11 connect-interface LoopBack0
 peer 192.168.1.13 as-number 100
 peer 192.168.1.13 connect-interface LoopBack0
 #
 address-family l2vpn evpn
  undo policy vpn-target
  peer 192.168.1.11 enable
  peer 192.168.1.11 reflect-client
  peer 192.168.1.13 enable
  peer 192.168.1.13 reflect-client
```

After all other switches are configured correctly, you can check if EVPN peers within the DC are established using

```
<5950-1>dis bgp peer l2vp evpn

 BGP local router ID: 192.168.1.12
 Local AS number: 100
 Total number of peers: 2                 Peers in established state: 2

  * - Dynamically created peer
  Peer                    AS  MsgRcvd  MsgSent OutQ PrefRcv Up/Down  State
```

```
  192.168.1.11           100      321       353    0       6 03:41:04 Established
  192.168.1.13           100      314       329    0       4 03:41:00 Established
```

Remote server MAC/IPs learnt can be checked using

```
<5950-1>dis bgp l2vp ev

 BGP local router ID is 192.168.1.12
 Status codes: * - valid, > - best, d - dampened, h - history,
               s - suppressed, S - stale, i - internal, e - external
               Origin: i - IGP, e - EGP, ? - incomplete

 Total number of routes from all PEs: 10

 Route distinguisher: 1:10
 Total number of routes: 4

     Network            NextHop         MED        LocPrf     PrefVal Path/Ogn

* >i [2][0][48][7848-59ed-0ac8][0][0.0.0.0]/104
                        192.168.1.11    0          100        0       i
* >i [2][0][48][7848-59ed-0ac8][32][10.1.1.10]/136
                        192.168.1.11    0          100        0       i
* >i [2][0][48][7848-59ed-0ad2][32][10.1.1.11]/136
                        192.168.1.13               100        0       200i
* >i [3][0][32][192.168.1.11]/80
                        192.168.1.11    0          100        0       i

 Route distinguisher: 1:20
 Total number of routes: 4

     Network            NextHop         MED        LocPrf     PrefVal Path/Ogn

* >i [2][0][48][7848-59ed-0ac9][0][0.0.0.0]/104
                        192.168.1.11    0          100        0       i
* >i [2][0][48][7848-59ed-0ac9][32][10.1.2.10]/136
                        192.168.1.11    0          100        0       i
* >i [2][0][48][7848-59ed-0ad3][32][10.1.2.11]/136
                        192.168.1.13               100        0       200i
```

Sample 5940-3 ED configuration

```
sysname 5940-3
#
l2vpn enable
vxlan tunnel arp-learning disable
vxlan tunnel mac-learning disable
#
hardware-resource vxlan border24k
#
ip vpn-instance vpna
 route-distinguisher 1:2
 #
 address-family evpn
  vpn-target 1:1 import-extcommunity
  vpn-target 1:1 export-extcommunity
#
tunnel global source-address 192.168.1.13
# ISIS is used as the WAN routing protocol
isis 1
 is-level level-1
 network-entity 10.0000.0000.0003.00
# OSPF is used as the routing protocol within the DC
ospf 1
 area 0.0.0.0
  network 192.168.1.0 0.0.0.255
#
interface LoopBack0
 ip address 192.168.1.13 255.255.255.255
 isis enable 1
#
interface FortyGigE1/1/25
 port link-mode route
 ip address unnumbered interface LoopBack0
 ospf network-type p2p
 lldp management-address arp-learning
 lldp tlv-enable basic-tlv management-address-tlv interface LoopBack0
#
interface Ten-GigabitEthernet1/1/22
 port link-mode route
 description WAN
 mtu 9000
 ip address 192.168.2.18 255.255.255.252
 isis enable 1
 dci enable
#
interface Vsi-interface1000
```

```
 ip binding vpn-instance vpna
 l3-vni 1000
#
bgp 100
 peer 192.168.1.12 as-number 100
 peer 192.168.1.12 connect-interface LoopBack0
 peer 192.168.1.14 as-number 200
 peer 192.168.1.14 connect-interface LoopBack0
 peer 192.168.1.14 ebgp-max-hop 64
 #
 address-family l2vpn evpn
  peer 192.168.1.12 enable
  peer 192.168.1.12 next-hop-local
  peer 192.168.1.14 enable
  peer 192.168.1.14 router-mac-local
# next-hop-local is used when peering to RR and router-mac-local is used when peering to other
# EDs
```

After all other switches are configured correctly, you can check on the EVPN routes learnt using

```
<5940-3>dis ip ro vp vpna

Destinations : 16       Routes : 16

Destination/Mask   Proto   Pre Cost        NextHop         Interface
0.0.0.0/32         Direct  0   0           127.0.0.1       InLoop0
10.1.1.0/24        BGP     255 0           192.168.1.11    Vsi1000
10.1.1.10/32       BGP     255 0           192.168.1.11    Vsi1000
10.1.1.11/32       BGP     255 0           192.168.1.14    Vsi1000
10.1.2.0/24        BGP     255 0           192.168.1.11    Vsi1000
10.1.2.10/32       BGP     255 0           192.168.1.11    Vsi1000
10.1.2.11/32       BGP     255 0           192.168.1.14    Vsi1000
10.80.22.0/24      BGP     255 0           192.168.1.14    Vsi1000
10.80.23.0/24      BGP     255 0           192.168.1.14    Vsi1000
```

Remote server MAC/IPs learnt can be checked using

```
<5940-3>dis bgp l2vpn evpn

 BGP local router ID is 192.168.1.13
 Status codes: * - valid, > - best, d - dampened, h - history,
               s - suppressed, S - stale, i - internal, e - external
               Origin: i - IGP, e - EGP, ? - incomplete

 Total number of routes from all PEs: 8

 Route distinguisher: 1:2(vpna)
```

```
 Total number of routes: 4

     Network            NextHop         MED        LocPrf     PrefVal Path/Ogn

* >i [2][0][48][7848-59ed-0ac8][32][10.1.1.10]/136
                        192.168.1.11    0          100        0       i
* >i [2][0][48][7848-59ed-0ac9][32][10.1.2.10]/136
                        192.168.1.11    0          100        0       i
* >e [2][0][48][7848-59ed-0ad2][32][10.1.1.11]/136
                        192.168.1.14                          0       200i
* >e [2][0][48][7848-59ed-0ad3][32][10.1.2.11]/136
                        192.168.1.14                          0       200i
```

You will notice the ED has underlay OSPF routes for switch loopback IPs in the same DC and ISIS learns about the loopback IP from the ED in the other DC.

```
<5940-3>dis ip ro

Destinations : 27       Routes : 27

Destination/Mask   Proto   Pre Cost        NextHop         Interface
0.0.0.0/32         Direct  0   0           127.0.0.1       InLoop0
10.10.10.0/24      Direct  0   0           10.10.10.39     MGE0/0/0
10.10.10.0/32      Direct  0   0           10.10.10.39     MGE0/0/0
10.10.10.39/32     Direct  0   0           127.0.0.1       InLoop0
10.10.10.255/32    Direct  0   0           10.10.10.39     MGE0/0/0
15.0.0.0/8         Static  60  0           10.10.10.254    MGE0/0/0
127.0.0.0/8        Direct  0   0           127.0.0.1       InLoop0
127.0.0.0/32       Direct  0   0           127.0.0.1       InLoop0
127.0.0.1/32       Direct  0   0           127.0.0.1       InLoop0
127.255.255.255/32 Direct  0   0           127.0.0.1       InLoop0
172.16.200.0/24    Direct  0   0           172.16.200.2    Vlan200
172.16.200.0/32    Direct  0   0           172.16.200.2    Vlan200
172.16.200.2/32    Direct  0   0           127.0.0.1       InLoop0
172.16.200.255/32  Direct  0   0           172.16.200.2    Vlan200
172.16.201.0/24    Static  60  0           172.16.200.254  Vlan200
192.168.1.11/32    O_INTRA 10  2           192.168.1.12    FGE1/1/25
192.168.1.12/32    O_INTRA 10  1           192.168.1.12    FGE1/1/25
192.168.1.13/32    Direct  0   0           127.0.0.1       InLoop0
192.168.1.14/32    IS_L1   15  20          192.168.2.17    XGE1/1/22
192.168.2.16/30    Direct  0   0           192.168.2.18    XGE1/1/22
192.168.2.16/32    Direct  0   0           192.168.2.18    XGE1/1/22
192.168.2.18/32    Direct  0   0           127.0.0.1       InLoop0
192.168.2.19/32    Direct  0   0           192.168.2.18    XGE1/1/22
192.168.2.20/30    IS_L1   15  20          192.168.2.17    XGE1/1/22
```

Configs for switches in DC2 will be similar to switches in DC1, the same verification commands used in DC1 can also be used.

Sample 5940-4 ED configuration

```
sysname 5940-4
#
l2vpn enable
vxlan tunnel arp-learning disable
vxlan tunnel mac-learning disable
#
hardware-resource vxlan border24k
#
ip vpn-instance vpna
 route-distinguisher 1:3
 #
 address-family evpn
  vpn-target 1:1 import-extcommunity
  vpn-target 1:1 export-extcommunity
#
tunnel global source-address 192.168.1.14
# ISIS is used as the WAN routing protocol
isis 1
 is-level level-1
 network-entity 10.0000.0000.0003.00
# OSPF is used as the routing protocol within the DC
ospf 1
 area 0.0.0.0
  network 192.168.1.0 0.0.0.255
#
interface LoopBack0
 ip address 192.168.1.14 255.255.255.255
 isis enable 1
#
interface FortyGigE1/1/25
 port link-mode route
 ip address unnumbered interface LoopBack0
 ospf network-type p2p
 lldp management-address arp-learning
 lldp tlv-enable basic-tlv management-address-tlv interface LoopBack0
#
interface Ten-GigabitEthernet1/1/22
 port link-mode route
 description WAN
 mtu 9000
 ip address 192.168.2.22 255.255.255.252
 isis enable 1
 dci enable
#
```

```
interface Vsi-interface1000
 ip binding vpn-instance vpna
 l3-vni 1000
#
bgp 200
 peer 192.168.1.15 as-number 200
 peer 192.168.1.15 connect-interface LoopBack0
 peer 192.168.1.13 as-number 100
 peer 192.168.1.13 connect-interface LoopBack0
 peer 192.168.1.13 ebgp-max-hop 64
 #
 address-family l2vpn evpn
  peer 192.168.1.15 enable
  peer 192.168.1.15 next-hop-local
  peer 192.168.1.13 enable
  peer 192.168.1.13 router-mac-local
# next-hop-local is used when peering to RR and router-mac-local is used when peering to other
# EDs
```

Sample 5950-2 RR2 configuration

```
sysname 5950-2
#
ospf 1
 area 0.0.0.0
  network 192.168.1.0 0.0.0.255
#
interface LoopBack0
 ip address 192.168.1.15 255.255.255.255
#
interface HundredGigE1/0/31
 port link-mode route
 ip address unnumbered interface LoopBack0
 ospf network-type p2p
 lldp management-address arp-learning
 lldp tlv-enable basic-tlv management-address-tlv interface LoopBack0
#
interface HundredGigE1/0/32
 port link-mode route
 ip address unnumbered interface LoopBack0
 ospf network-type p2p
 lldp management-address arp-learning
 lldp tlv-enable basic-tlv management-address-tlv interface LoopBack0
#
bgp 200
 peer 192.168.1.14 as-number 200
 peer 192.168.1.14 connect-interface LoopBack0
```

```
 peer 192.168.1.16 as-number 200
 peer 192.168.1.16 connect-interface LoopBack0
 #
 address-family l2vpn evpn
  undo policy vpn-target
  peer 192.168.1.14 enable
  peer 192.168.1.14 reflect-client
  peer 192.168.1.16 enable
  peer 192.168.1.16 reflect-client
```

Sample 5940-2 gateway configuration

```
sysname 5940-2
#
l2vpn enable
vxlan tunnel arp-learning disable
vxlan tunnel mac-learning disable
#
hardware-resource vxlan l3gw24k
#
# 2 VSIs are created as 1 gateway interface can only be attached to 1 VSI
vsi vpna-1
 arp suppression enable
 vxlan 10
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
#
vsi vpna-2
 arp suppression enable
 vxlan 20
 evpn encapsulation vxlan
  route-distinguisher auto
  vpn-target auto export-extcommunity
  vpn-target auto import-extcommunity
#
ospf 1
 area 0.0.0.0
  network 192.168.1.0 0.0.0.255
#
interface LoopBack0
 ip address 192.168.1.16 255.255.255.255
#
tunnel global source-address 192.168.1.16
#
```

```
interface FortyGigE1/1/1
 port link-mode bridge
 service-instance 1
  encapsulation s-vid 10
  xconnect vsi vpna-1
 service-instance 2
  encapsulation s-vid 11
  xconnect vsi vpna-2
#
interface HundredGigE1/1/3
 port link-mode route
 ip address unnumbered interface LoopBack0
 ospf network-type p2p
 lldp management-address arp-learning
 lldp tlv-enable basic-tlv management-address-tlv interface LoopBack0
#
ip vpn-instance vpna
 route-distinguisher 1:4
 #
 address-family evpn
  vpn-target 1:1 import-extcommunity
  vpn-target 1:1 export-extcommunity
#
interface Vsi-interface1
 ip binding vpn-instance vpna
 ip address 10.1.1.1 255.255.255.0
 mac-address 0001-0001-0001
 local-proxy-arp enable
 distributed-gateway local
# local-proxy-arp is used to respond back on all ARP requests, the EVPN table will be used
# to determine if the remote MAC addresses are known
#
interface Vsi-interface2
 ip binding vpn-instance vpna
 ip address 10.1.2.1 255.255.255.0
 mac-address 0002-0002-0002
 local-proxy-arp enable
 distributed-gateway local
#
vsi vpna-1
 gateway vsi-interface 1
#
vsi vpna-2
 gateway vsi-interface 2
#
interface Vsi-interface1000
```

```
 ip binding vpn-instance vpna
 l3-vni 1000
#
bgp 100
 peer 192.168.1.15 as-number 100
 peer 192.168.1.15 connect-interface LoopBack0
 #
 address-family l2vpn evpn
  peer 192.168.1.15 enable
# EVPN peering only to RR
```

```
[5940-2-FortyGigE1/1/1]dis l2vpn mac
MAC Address    State    VSI Name                        Link ID/Name    Aging
7848-59ed-0ad2 Dynamic  vpna-1                          FGE1/1/1        Aging
7848-59ed-0ad3 Dynamic  vpna-2                          FGE1/1/1        Aging
--- 2 mac address(es) found  ---
```

```
[5940-2]dis ip ro vp vpna

Destinations : 20       Routes : 20

Destination/Mask   Proto   Pre Cost        NextHop         Interface
0.0.0.0/32         Direct  0   0           127.0.0.1       InLoop0
10.1.1.0/24        Direct  0   0           10.1.1.1        Vsi1
10.1.1.0/32        Direct  0   0           10.1.1.1        Vsi1
10.1.1.1/32        Direct  0   0           127.0.0.1       InLoop0
10.1.1.10/32       BGP     255 0           192.168.1.14    Vsi1000
10.1.1.255/32      Direct  0   0           10.1.1.1        Vsi1
10.1.2.0/24        Direct  0   0           10.1.2.1        Vsi2
10.1.2.0/32        Direct  0   0           10.1.2.1        Vsi2
10.1.2.1/32        Direct  0   0           127.0.0.1       InLoop0
10.1.2.10/32       BGP     255 0           192.168.1.14    Vsi1000
```

Resources, contacts, or additional links

R250x-HPE FlexFabric 5940 Switch Series EVPN Configuration Guide

Learn more at

hpe.com

Sign up for updates

A00000256ENW, January 2017