

November 2020

VXLAN on CX 8400 and 8360 Switches

Daryl Wan
Satheesh Kumar
Pramod Kumar A S
Matthew Finlay

aruba
a Hewlett Packard
Enterprise company



Agenda: 8360/8400 VXLAN



1

Overview

2

Use Cases

3

Details/Caveats

4

Configuration and Best Practices

5

Troubleshooting

6

Resources

7

Demo

Overview

aruba

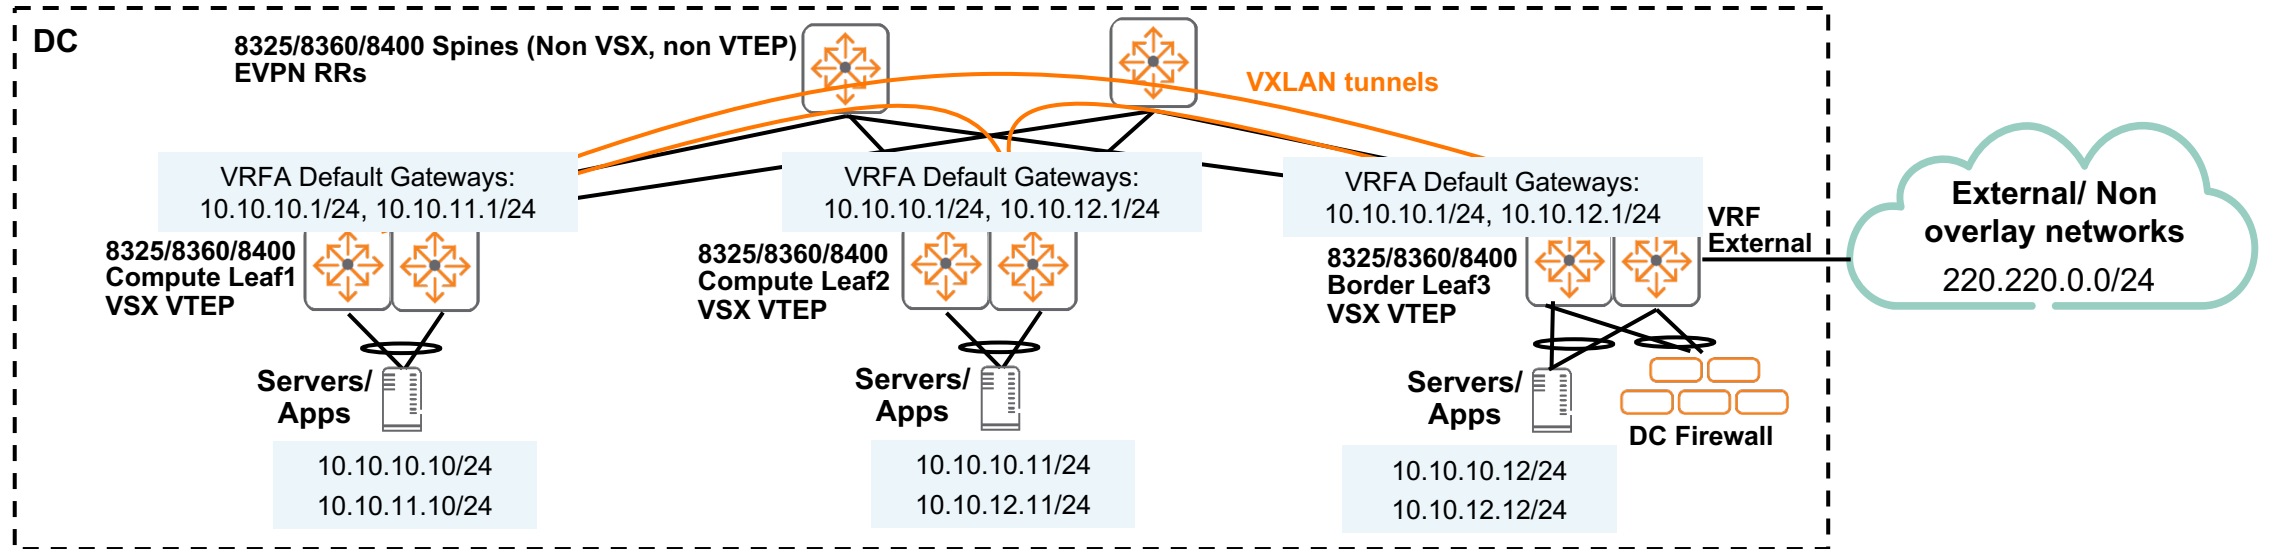
a Hewlett Packard
Enterprise company

8360/8400 VXLAN Overview

- 10.5 provides VXLAN EVPN support for 6300/6400/8325 and EVPN support for 8400
- 10.6 adds support for VXLAN EVPN on 8360/8400
- 8360 is targeted for core and aggregation layer of campuses, as well as Data Center (DC) leaf and compact spine requirements
- 8400 is targeted for campus/DC core or large scale border router requirements

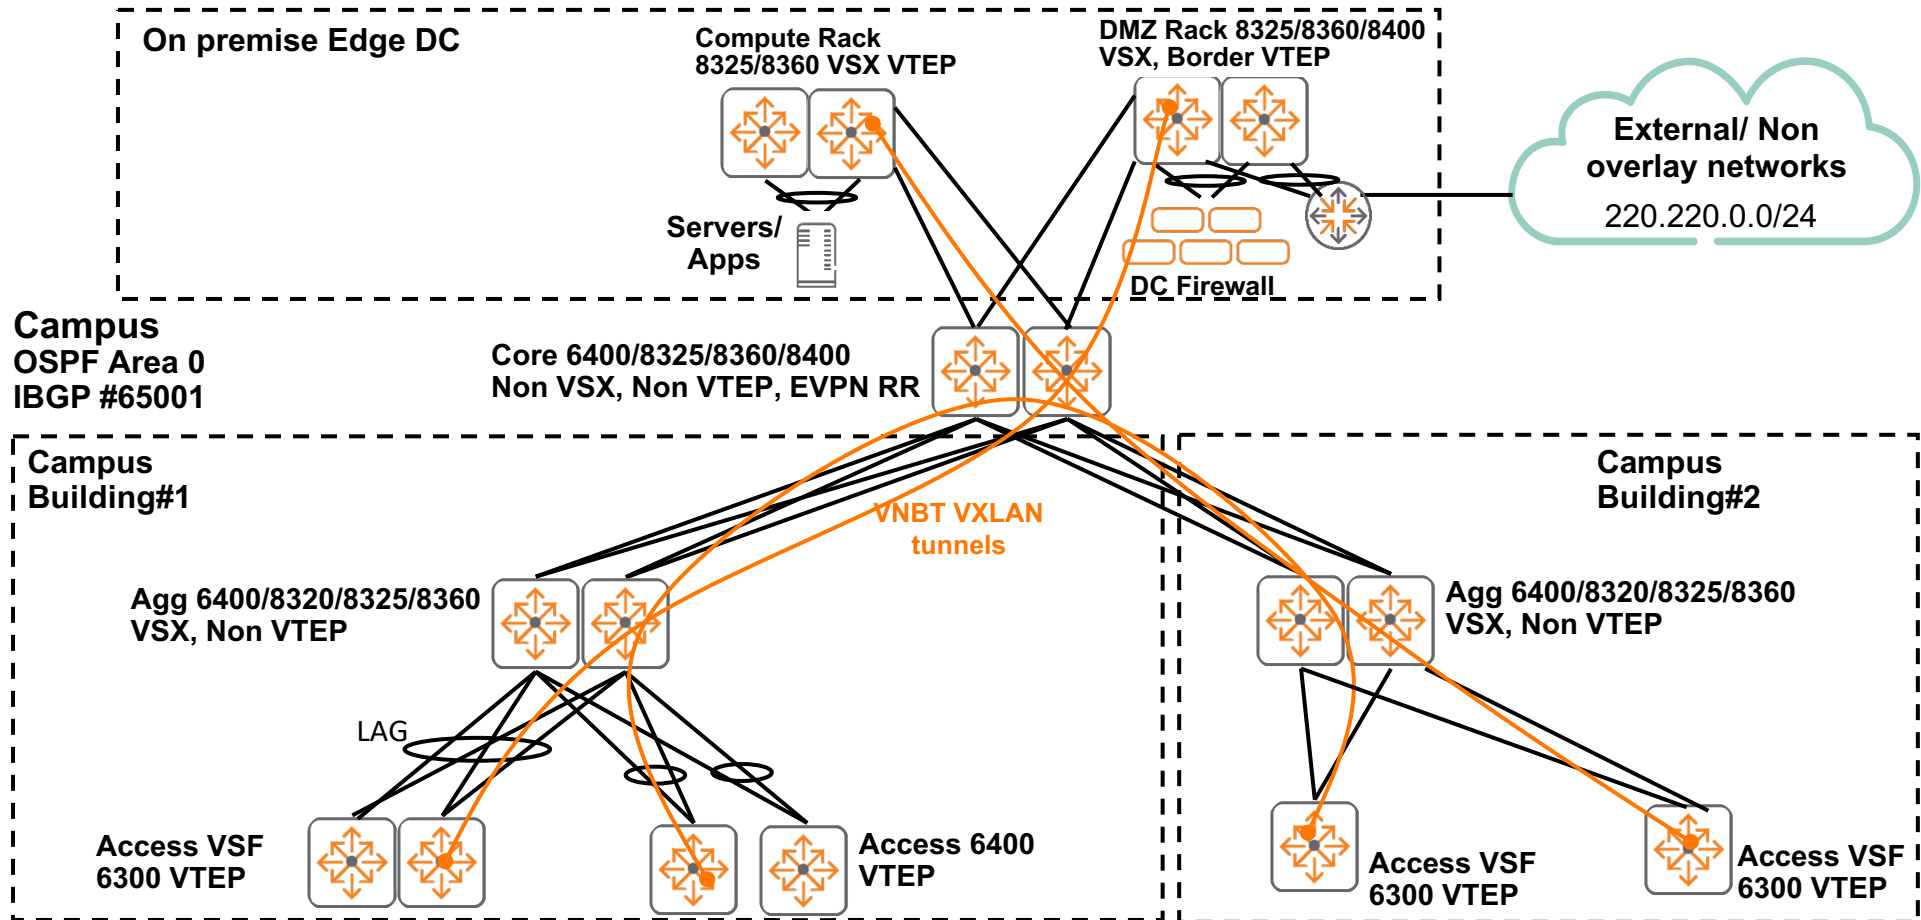
Use Cases

DC Use Case



- Distributed L3 gateways (symmetric IRB) recommended for DC
- Before 10.6
 - 8325 recommended as both compute & border leaf/VTEP
 - 8325/8400 recommended as EVPN RR spine
- From 10.6
 - 8360/8400 can be used as both compute & border leaf/VTEP
 - 8360 can be used as EVPN RR compact spine

Campus/Edge DC Use Case



- Distributed L3 gateways (symmetric IRB) recommended for Campus
- Before 10.6
 - 8325 recommended as both compute & border VTEP in edge DC
 - 8325/8400 recommended as EVPN RR campus core
- From 10.6
 - 8360 can be used as VSX compute VTEP in edge DC or EVPN RR (compact campus core)
 - 8360/8400 can be used as VSX border VTEP
 - 8360 also suitable as non VTEP Campus Agg switches or compact Core switch

AOS-CX 10.6 VXLAN EVPN Scale*

Switch Model	6300	6400	8325	8360	8400
MAC	32,768	32,768	98,304	114,688	768,000
IPv4 Host Table (ARP)	49,152	49,152	120,000	163,840	756,000
IPv6 Host Table (ND)	49,152	49,152	52,000	163,840	524,000
IPv4 Unicast Routes	61,000	61,000	131,072	1,048,576	1,011,712
BGP Peers	256	256	256	256	256
EVPN VXLAN VTEP Peers (Number of Tunnels)	256	256	1024	1,024	1,024
EVPN VXLAN L2 VNI (VLANs)	1,024	1,024	4,039 (VLAN 1 not supported)	4,093	4,093

* Refer to data sheet after 10.6 GA for final confirmed numbers

Details

aruba

a Hewlett Packard
Enterprise company

8360 VXLAN EVPN Details

8360 Scale

- VNI – 4093
 - VTEP – 1024
 - L3VNI – 255
 - Overlay Hosts – 110K
-
- Above values are per switch and not shown in datasheet
 - Refer to AOS-CX TFM for final confirmed numbers

8400 VXLAN EVPN Details

8400 Scale

- VNI – 4093
 - VTEP – 1024
 - L3VNI – 255
 - Overlay Hosts – 54K
-
- Above values are per switch and not shown in datasheet
 - Refer to AOS-CX TFM for final confirmed numbers

8400 VXLAN EVPN Caveats

Caveats	Comments
DSCP value is not copied from inner header to outer header when encapsulated over VXLAN <ul style="list-style-type: none">• 8400 as core/agg switch will not prioritize traffic in VXLAN tunnels correctly	Planned to be fixed in the future release.
GRE tunnels are not supported on JL687A (25G LCs) due to VXLAN enablement	Planned to be fixed in the future release.

** Refer to 10.6 release notes for final list of caveats*

Configuration

aruba

a Hewlett Packard
Enterprise company

VXLAN EVPN Distributed L3 Gateway Configuration

- Refer to <https://community.arubanetworks.com/t5/Data-Center-Networking-Solutions/AOS-CX-VXLAN-EVPN-Symmetric-IRB-Distributed-L3-Gateways/ta-p/664764>
- Above 10.5 configs documented also applicable to 8360/8400
- **In future 10.6 CPE**, the same SVI/AG IP will be supported for distributed L3 gateways with IPv4 unicast

```
interface vlan11
  vrf attach VRF1
  ip address 10.1.11.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.11.1
```

Best Practices

aruba

a Hewlett Packard
Enterprise company

VXLAN EVPN Distributed L3 Gateway Best Practices

- Refer to <https://community.arubanetworks.com/t5/Data-Center-Networking-Solutions/AOS-CX-VXLAN-EVPN-Symmetric-IRB-Distributed-L3-Gateways/ta-p/664764>
- Above 10.5 recommendations also applicable to 8360/8400
- 8400
 - Use route/subnet redistribution instead of host redistribution as applicable based on host scale

aruba

a Hewlett Packard
Enterprise company

Troubleshooting

VXLAN EVPN Distributed L3 Gateway Verification

- Refer to <https://community.arubanetworks.com/t5/Data-Center-Networking-Solutions/AOS-CX-VXLAN-EVPN-Symmetric-IRB-Distributed-L3-Gateways/ta-p/664764>
- Above 10.5 verification commands documented also applicable to 8360/8400
- Check Tunnels are operational - “show interface vxlan vteps”
 - If not, check OSPF/BGP peers are established – “show bgp all-vrf all summary”
 - EVPN is advertising and receiving VPN info from peers – “show bgp l2vpn evpn”
 - Check spine / remote VTEP-peer has the neighbor and EVPN configuration to share routes and l2vpn
- Check VNI is operational - “show interface vxlan vni”
 - If not, check EVPN is advertising and receiving VPN info – “show bgp l2vpn evpn”
- Check if the host/route is learnt on the VXLAN as applicable – “show ip route bgp all-vrfs”
 - If not, check route or host redistribute configuration on the EVPN configuration
 - Check EVPN is advertising and receiving VPN info for the VNI.
 - Check if the host or next hop over VXLAN are learnt – “show arp all-vrf” or “show ipv6 neighbors all-vrfs” as applicable.

8400 VXLAN EVPN Troubleshooting Specifics

- VXLAN info available through diagdump's and support files "**diag-dump vxlan basic**"
 - Diagdump MM – support file path: feature/vxlan/diagdump.txt
 - Shows VXLAN info in AOS-CX (Tunnels/Routes/VPN PI data)
 - Shows VXLAN 8400 platform info (overlay_*)
 - Shows counter info for the internal logic (OVERLAY: Non-Zero [Error/Warn/Info/Debug] Counters)
 - Diagdump LC – support file path: ./diag_dump_vxlan.txt ""**diag-dump vxlan basic module <LC>**""
 - Shows VXLAN info in LC CPU (section - CACHE_FULL_DUMP_BEGIN/..._END)
 - Shows VXLAN info programmed in the ASIC (section - ASIC_DUMP_FULL_BEGIN/..._END)
 - Shows counter info for the internal logic (OVERLAY: Non-Zero [Error/Warn/Info/Debug] Counters)
- Check L3PD programming for host and routes for overlay hosts and routes in diag dump output.
- Check VXLAN counters "show int vxlan 1 statistics"

8400 VXLAN EVPN Troubleshooting Specifics

Diagdump MM

```
MM:
overlay_vpn_s information
  Cache Entries: 3
VPN*      VNI      VRF-NAME      VRF      VLAN      MC-ID      IS-L3VNI      HA-STATE
-----
10         10010      VRF1           4         10         10         false         normal
11         10011      VRF1           4         11         11         false         normal
10240      100001     VRF1           4         10240      10240      true          normal
```

Diagdump LC

```
LC CPU:
VPN AGENT:
overlay_vpn_s information
  Cache Entries: 3
VPN*      VNI      VRF-NAME      VRF      VLAN      MC-ID      IS-L3VNI      HA-STATE
-----
10         10010      VRF1           4         10         10         false         normal
11         10011      VRF1           4         11         11         false         normal
10240      100001     VRF1           4         10240      10240      true          normal
```

```
LC ASIC:
HW VPNs:
VPN      VNI
-----
10        10010
11        10011
10240     100001
```

8400 VXLAN EVPN Troubleshooting Specifics

show int vxlan 1 statistics

Interface	RX Bytes		RX Packets		RX Drops		TX Bytes		TX Drops	
RX Multicast	TX Broadcast		TX Multicast		RX Pause		TX Pause		RX Broadcast	

vxlan1	1351626		9015		0		591570		4553	
0	0		0						0	

8360 VXLAN EVPN Troubleshooting Specifics

- Show commands for vxlan tunnels
 - **show interface vxlan**
 - **show interface vxlan vteps**
 - **show interface vxlan vni**
 - **show interface vxlan 1 statistics**
- PI Commands for VXLAN tunnels
 - Captured in the Wiki link
- Command to collect debug information for vxlan tunnels is “**diag-dump vxlan basic**”
 - Sample display captured in the Wiki link
 - Shows VXLAN info in AOS-CX (Tunnels/Routes/VPN PI hmap)
 - The 8360 Tunnel, VPN Platform information in the MM and LC are captured
 - The Plugin cache(MM), Agent Cache(LC) and the ASIC(LC) dumps are displayed under “**PLUGIN CACHE DUMP BEGIN**” , “**AGENT CACHE DUMP BEGIN**” and “**ASIC DUMP BEGIN**” in the **diag dump vxlan basic** command
- VXLAN info available through diagdump’s in feature/vxlan/diagdump.txt in support-files

8360 VXLAN EVPN Troubleshooting Specifics

Plugin Cache (on MM)

***** PLUGIN CACHE DUMP BEGIN *****

VPN ENTRY

Vni id	Rep Grp Id	Vlan	MM State	Status
-----	-----	----	-----	-----
20	34	200	Undef	False
100000	33	4101	Undef	False
10	32	100	Undef	False

TUNNEL/ECMP/Nexthop CACHE

Tunnel Index	Src IP	Dst IP	Internal Vlan	VRF	Ecmp Id	TTL	MTU
Udp Port	MM State	Tunnel Status	Vlan Conf				
Ecmp Conf							
-----	-----	-----	-----	---	-----	---	---
-----	-----	-----	-----				
1	1.1.1.1	3.3.3.3	7167	2	1	255	
9162	4789	Undef	Configured	False			
False							
ECMP Id	Vlan id	L3 intf id	Log Port	Num of Nhops	Member State	Nexthop_Mac	
Port Name	ISL Redirected						---
-----	-----	-----	-----	-----	-----	-----	
1	4123	5	16777508	1	3	90:20:C2:BA:B3:00	
1/1/12	NA						
.....							

8360 VXLAN EVPN Troubleshooting Specifics

Plugin Cache (on MM) Continued....

```
NETVP CACHE
*****
Vni id      Tunnel id    MM State
-----
100000      1             Undef
10          1             Undef
=====
***** PLUGIN CACHE DUMP END *****
```

8360 VXLAN EVPN Troubleshooting Specifics

Agent Cache (on LC)

```
-----  
[Start] Daemon switchd_agent0  
-----
```

```
LC:          NODE ID:  
-----
```

```
2            1  
=====
```

```
***** AGENT CACHE DUMP BEGIN *****
```

```
VPN ENTRY
```

```
*****
```

Vni id	Rep Grp Id	Vlan	MM State	Status
-----	-----	-----	-----	-----
20	34	200	Undef	False
100000	33	4101	Undef	True
10	32	100	Undef	True

8360 VXLAN EVPN Troubleshooting Specifics

Agent Cache (on LC) continued..

TUNNEL/ECMP/Nexthop CACHE

Tunnel Index	Src IP	Dst IP	Internal Vlan	VRF	Ecmp Id	TTL	MTU
Udp Port	MM State	Tunnel Status	Vlan Conf				
Ecmp Conf							
-----	-----	-----	-----	---	-----	---	---
-----	-----	-----	-----				

1	1.1.1.1	3.3.3.3	7167	2	1	255	
9162	4789	Undef	Configured	True			
True							
ECMP Id	Vlan id	L3 intf id	Log Port	Num of Nhops	Member State	Nexthop_Mac	
Port Name	ISL Redirected						
-----	-----	-----	-----	-----	-----	-----	---
-----	-----						
1	4123	5	16777508	1	1	90:20:C2:BA:B3:00	
1/1/12	NA						
.....							

NETVP CACHE

Vni id	Tunnel id	MM State
-----	-----	-----
ni id	Tunnel id	MM State
[10/1726]		
-----	-----	-----
100000	1	Undef
10	1	Undef
=====		

***** AGENT CACHE DUMP END *****

8360 VXLAN EVPN Troubleshooting Specifics

ASIC Dump (on LC)

***** ASIC DUMP BEGIN *****

VPN ENTRY

Vni Id	Rep Grp Id	Dst Vlan	Tunnel Index	Use_Pkt_Dst_Vlan
-----	-----	-----	-----	-----
100000	33	4101	1	False
10	32	100	1	False

TUNNEL ENTRY

Tunnel Index	Status	Src IP	Dst IP	ECMP Id	MTU	TTL
Internal Vlan	Encap pkts	Encap bytes				
-----	-----	-----	-----	-----	---	---
1	Enabled	1.1.1.1	3.3.3.3	1	9162	255
7167	1716	173436				

TUNNEL EXIT

Tunnel Index	Status	Src IP	Decap Pkts	Decap Bytes
-----	-----	-----	-----	-----
1	Enabled	3.3.3.3	1716	173436

TUNNEL ECMP MEMBERS

NFD ID	Local ASIC Port	Local Port Type	VLAN Number	Nexthop MAC
-----	-----	-----	-----	-----
5 36	PHY	4123	90:20:C2:BA:B3:00	

.....

8360 VXLAN EVPN Troubleshooting Specifics

ASIC Dump (on LC) continued

NETVP ENTRY

Vni_id	Vlan Num	Rep Grp Id	Ucast Override	Mcast Override	IP Mcast Override	Bcast Override
100000	4101	33	True	True	True	True
10	100	32	True	True	True	True

Available REP entries: 57342

***** ASIC DUMP END *****

[End] Daemon switchd_agent0

[End] Feature vxlan

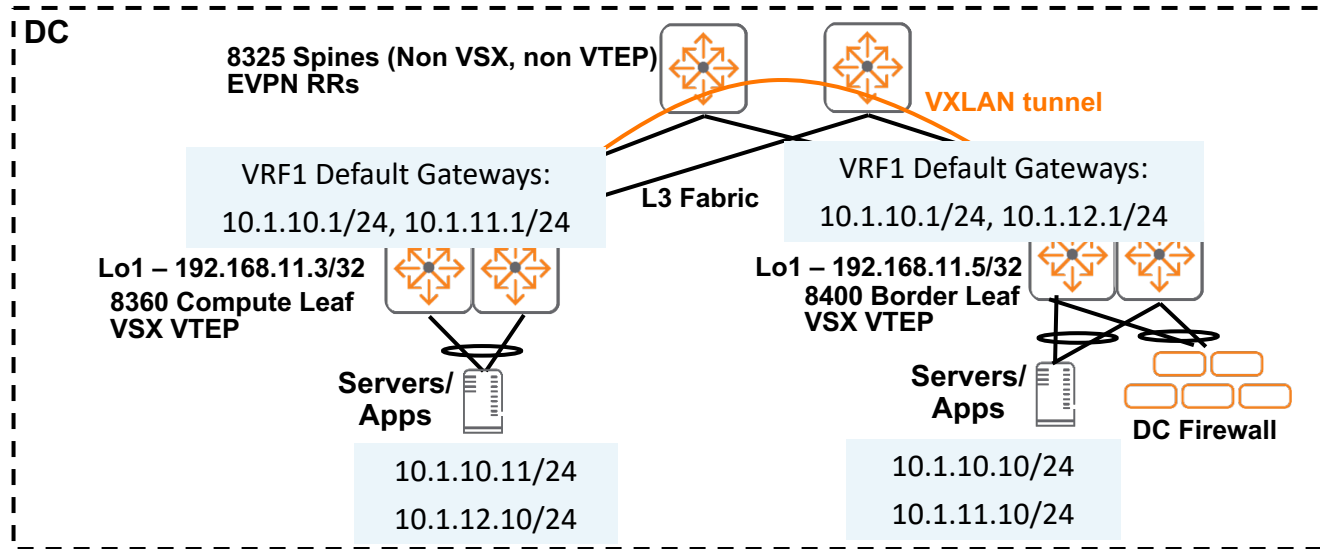
8360 VXLAN EVPN Troubleshooting Specifics

- Show commands for vxlan tunnels
 - **show interface vxlan**
 - **show interface vxlan vteps**
 - **show interface vxlan vni**
 - **show interface vxlan 1 statistics**
- Commands to collect debug information for VXLAN tunnels in "**diag-dump vxlan basic**"
 - Show VXLAN info in AOS-CX (Tunnels/Routes)
 - The 8360 Tunnel, VPN Platform information in the MM and LC are captured
 - The Plugin cache (MM), Agent cache (LC) and ASIC (LC) dumps are displayed under "**PLUGIN CACHE DUMP BEGIN**", "**AGENT CACHE DUMP BEGIN**" and "**ASIC DUMP BEGIN**" in this command's output
- VXLAN info available through diagdump's in feature/vxlan/diagdump.txt in support files

Demo



Demo – Distributed L3 Gateways with 8360/8400 VSX VTEPs



- Demo will show
 - VMs are able to communicate using both L2 VNI (same subnet), L3 VNI (different subnet)
 - Commands to help verify

Thank you