CX 10.8 Update
Jul 2021

aruba
a Hewlett Packard Enterprise company

VXLAN Enhancements

Presenters
- Daryl Wan
- Justin Noonan

Agenda

1. VNBT - VXLAN GBP
2. VNBT - VXLAN Stub VTEP
3. IPv4 DHCP Relay over VXLAN
4. IPv4/IPv6 Ping/Traceroute over VXLAN
5. IPv4/IPv6 Radius/RadSec over VXLAN
6. IPv4 Multicast VXLAN
7. IPv4 VXLAN DCI Use Case Validation

VNBT - VXLAN GBP

VXLAN Group Based Policy (GBP) Overview

- This feature enhances the Campus VXLAN - Virtual Network Based Tunneling (VNBT) solution
- 10.8 adds support for VXLAN GBP in VXLAN overlay networks

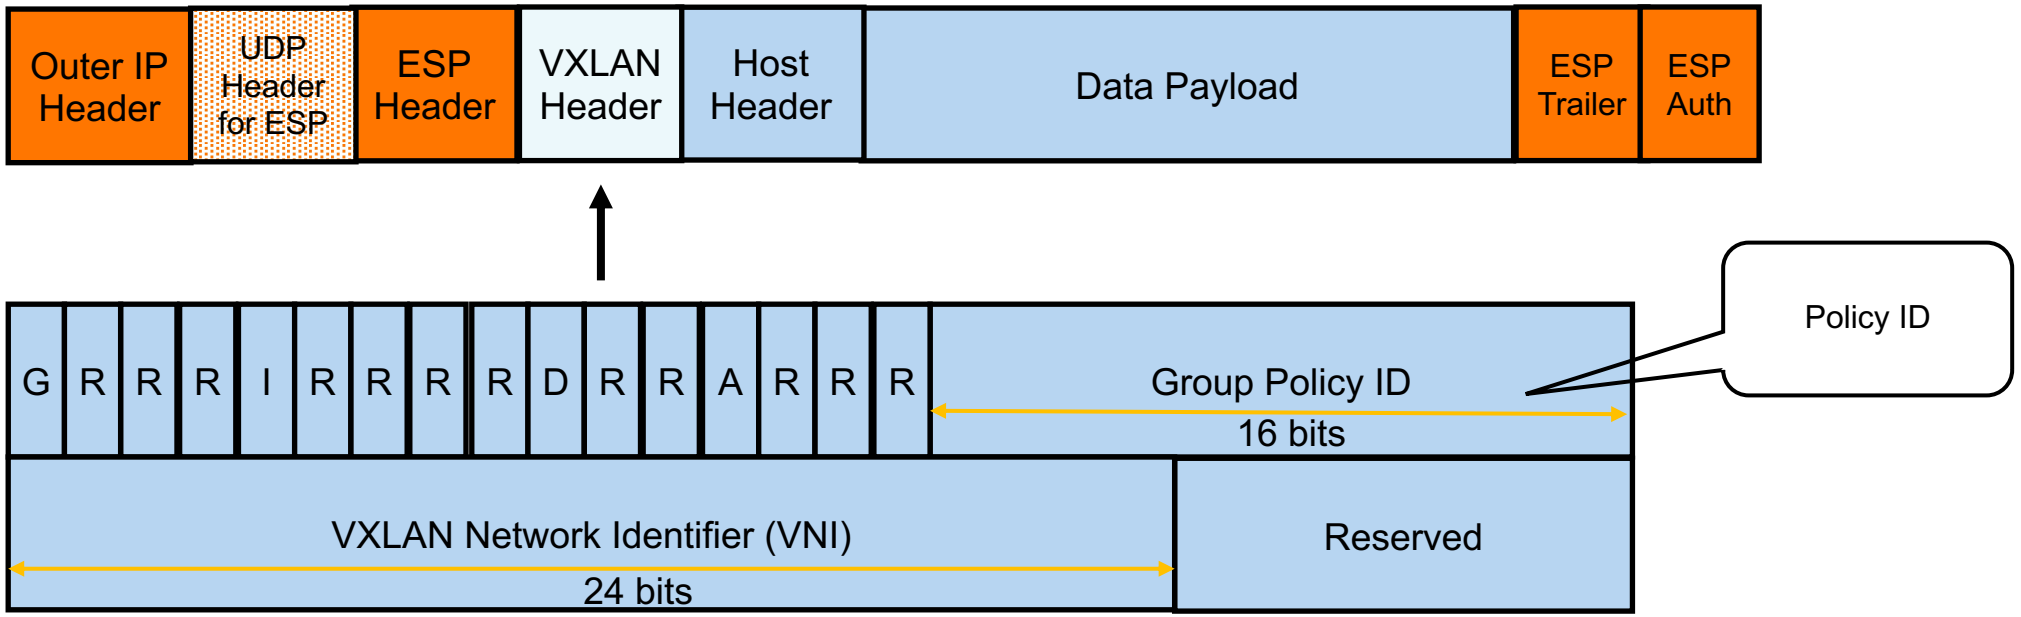

- Supported on 6300/6400/8360
- Enables role based policies
  - Role based policies are no longer tied to IP addresses
  - Source based roles (e.g. employee or guest) are assigned to a user/device on ingress VTEP and would remain effective even if it authenticates at a different location, or is assigned to a different IP subnet
- Refer to VXLAN GBP session for more details

VNBT - VXLAN Stub VTEP

VXLAN Stub VTEP Overview

- This feature enhances the Campus VXLAN - Virtual Network Based Tunneling (VNBT) solution
- Provides VXLAN overlay network connectivity between Gateway clients and VNBT clients
- Supported platforms: 6300, 6400, 8360
- Recommended platforms: 6400, 8360
- Refer to VXLAN Stub VTEP session for more details

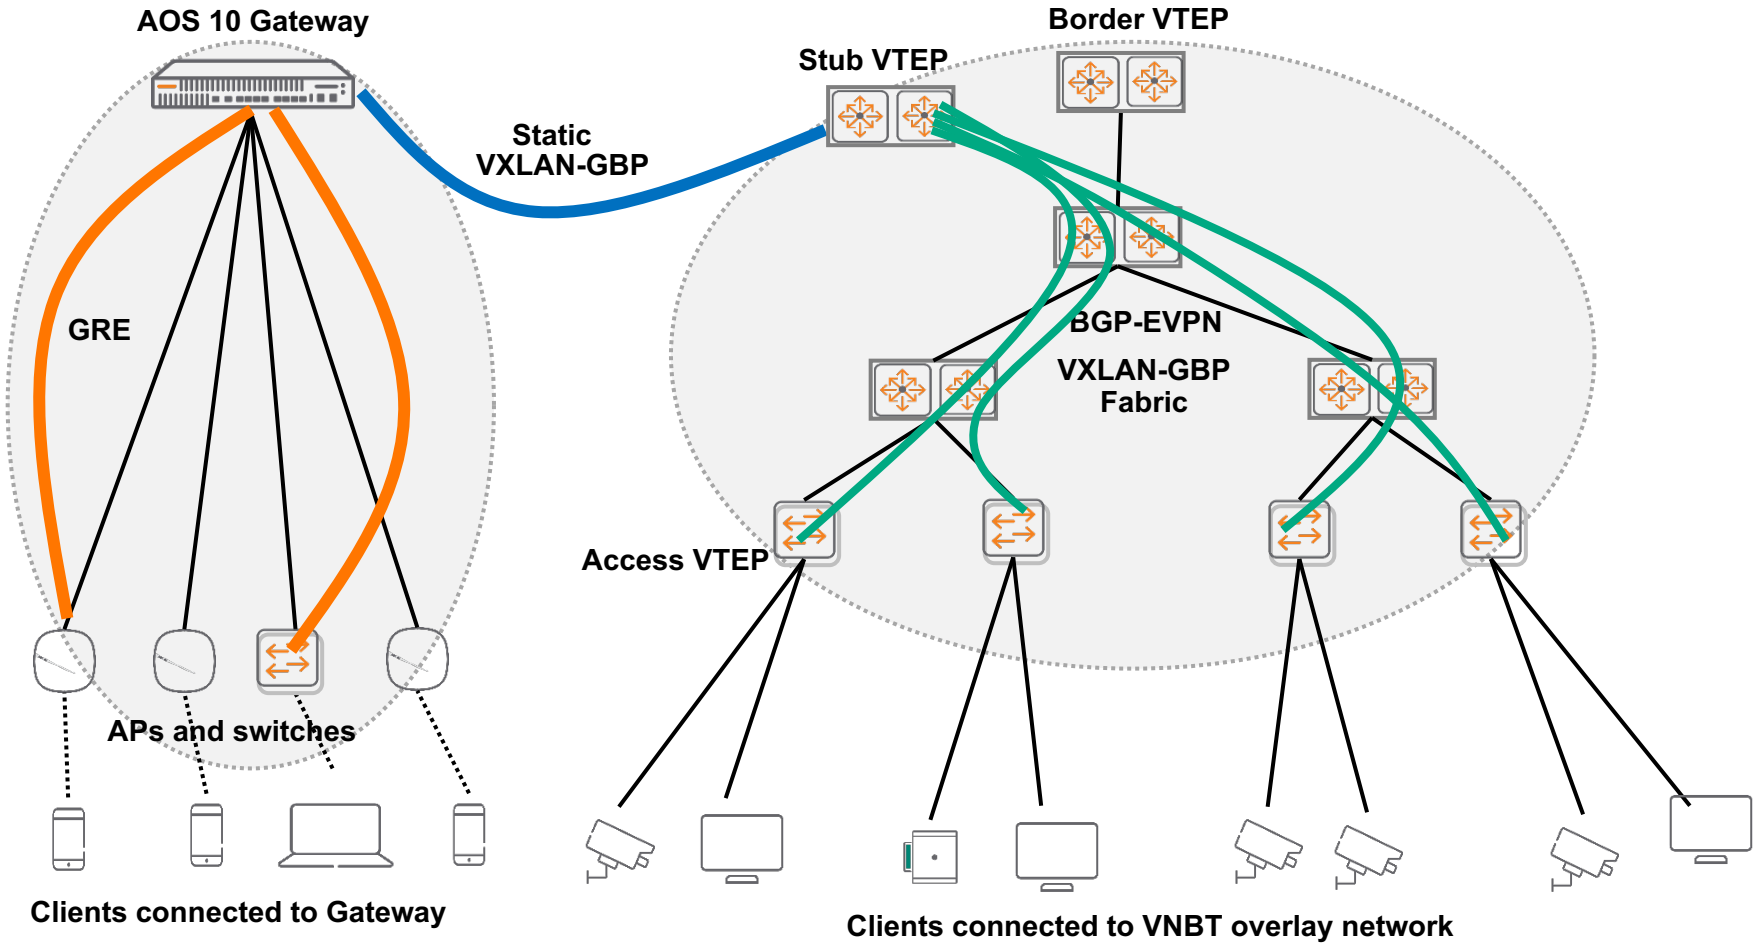

IPv4 DHCP Relay over VXLAN

IPv4 DHCP Relay in VXLAN Overlay Overview

- 10.7 added support for IPv4 DHCP relay with DHCP servers in the EVPN-VXLAN overlay network:
  - On 6300/6400/8360

- 10.8 adds feature parity for 8325/8400
  - Intra VRF – DHCP servers reachable from the same VRF by DHCP clients
  - Inter VRF – DHCP servers reachable from different VRF(s) by DHCP clients
  - Refer to 10.7 session on IPv4 DHCP Relay in VXLAN for more details

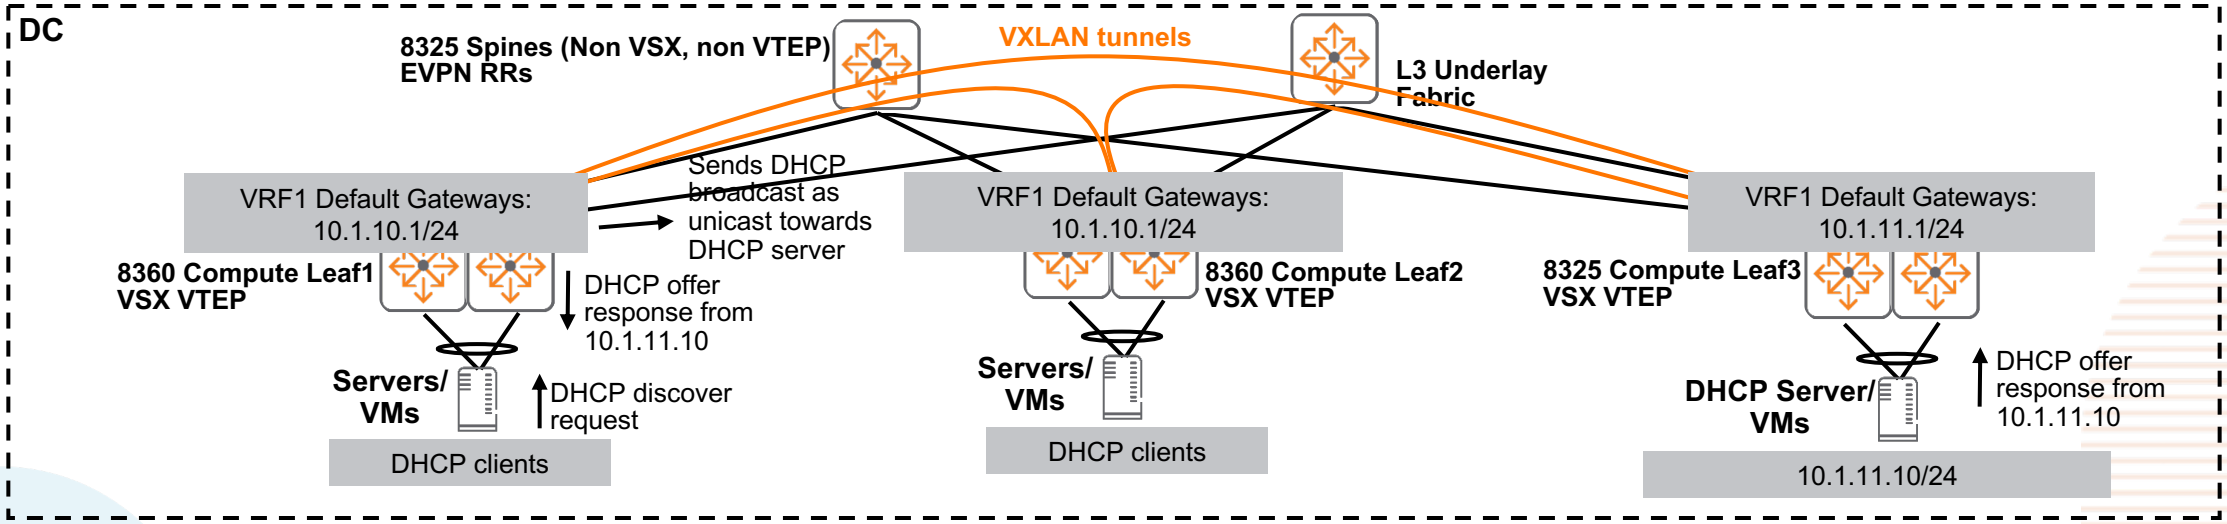

IPv4/IPv6 Ping/Traceroute over VXLAN

IPv4/IPv6 Ping/Traceroute over VXLAN

- Ping/traceroute from VTEP to VTEP, VTEP to host, host to VTEP over L2 VNI/L3 VNI are now supported
- Supported platforms: 6300, 6400, 8325, 8360, 8400

- Unique IP on VTEP should be used as ping/traceroute source/destination
- Both source/destination VTEPs require 10.8 for this feature to work

- Example of “DHCP relay over VXLAN” with unique IP (Lo13) on 8360 VSX switches with that will benefit from “ping/traceroute over VXLAN” connectivity verification

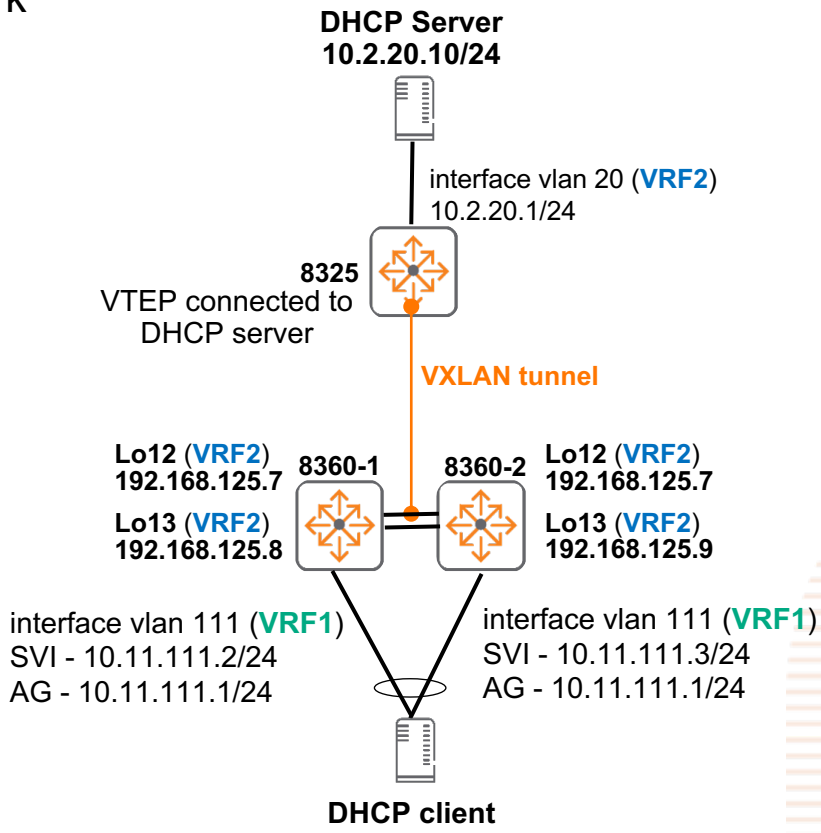

Demo – ping/traceroute over VXLAN

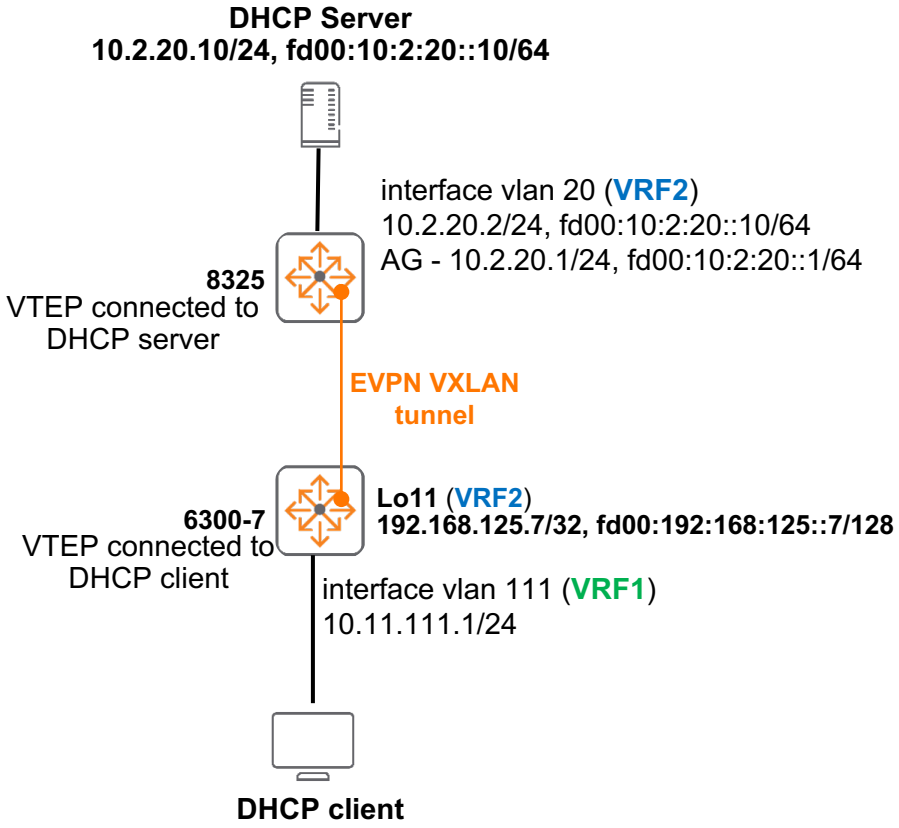

- IPv4/IPv6 ping/traceroute from Lo11 (**VRF2**) on DHCP client connected VTEP (6300-1) to DHCP server in **VRF2**

IPv4/IPv6 Radius/RadSec over VXLAN

IPv4/IPv6 Radius/RadSec over VXLAN

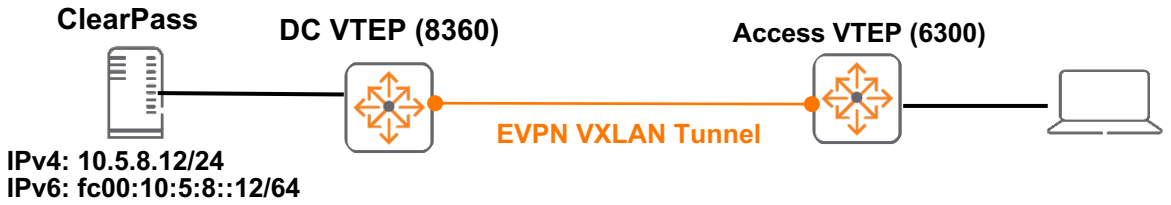

- Provides support for IPv4/IPv6 Radius/RadSec over VXLAN overlay network
- Supported platforms:

  Network users (supported on 6300/6400)
  Management users (supported on 6300/6400/8325/8360/8400)

- No additional configurations required on AOS-CX switches
- RADIUS or RADsec connection status can be verified using “show radius-server detail” command after enabling tracking (for RADsec without tracking connection status, also use “show radius-server detail”)

Caveats:

- MTU size (both interface and IP MTU) on all interfaces from RADsec client (AOS-CX switch) to RADsec server must be set to higher value based on the certificate size
- If IDEVID is used as RADsec client certificate, MTU size should be set above 1550 bytes to establish RADsec connection

Configuration

Configure connectivity over VXLAN tunnel

```
router bgp 65100
    bgp router-id 192.168.0.3
    no bgp fast-external-fallover
    neighbor 192.168.0.1 remote-as 65100
    neighbor 192.168.0.1 update-source loopback 0
    address-family l2vpn evpn
        neighbor 192.168.0.1 activate
        neighbor 192.168.0.1 send-community extended
    exit-address-family
!
    vrf VRF1
        address-family ipv4 unicast
            redistribute connected
            redistribute local loopback
        exit-address-family
        address-family ipv6 unicast
            redistribute local loopback
        exit-address-family
interface vxlan 1
    source ip 192.168.0.3
    no shutdown
    vni 58
        vlan 58
    vni 200
    vni 202
        vlan 202
    vni 100001
        vrf VRF1
        routing
interface loopback 10
    vrf attach VRF1
    ip address 192.168.10.1/32
    ipv6 address fd00:192:168:10::1/128
    ip pim-sparse enable

ip source-interface radius interface loopback10 vrf VRF1
ipv6 source-interface radius interface loopback10 vrf VRF1
```

IPv4 RADIUS Configuration (with RadSec)

```
radius-server host 10.5.8.12 tls vrf VRF1
```

IPv6 RADIUS Configuration (without RadSec)

```
radius-server host fc00:10:5:8::12 key ciphertext <key> vrf VRF1
```

Validate connectivity over VXLAN tunnel

```
VNBT-Access(config)# show ip route 10.5.8.12 vrf VRF1

VRF: VRF1

   Prefix           : 10.5.8.12/32                          VRF(egress)        : -
   Nexthop          : 192.168.0.1                           Interface          : -
   Origin           : bgp                                   Type               : bgp_evpn
   Distance         : 200                                   Metric             : 0
   Age              : 07h:29m:42s                           Tag                : 0
   Encap Type       : vxlan                                 Encap Details      : l3vni 100001



VNBT-Access(config)# ping 10.5.8.12 source loopback10 vrf VRF1
PING 10.5.8.12 (10.5.8.12) from 192.168.10.1 : 100(128) bytes of data.
108 bytes from 10.5.8.12: icmp_seq=1 ttl=62 time=0.527 ms
108 bytes from 10.5.8.12: icmp_seq=2 ttl=62 time=0.416 ms
108 bytes from 10.5.8.12: icmp_seq=3 ttl=62 time=0.403 ms
108 bytes from 10.5.8.12: icmp_seq=4 ttl=62 time=0.592 ms
108 bytes from 10.5.8.12: icmp_seq=5 ttl=62 time=0.542 ms

--- 10.5.8.12 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4079ms
rtt min/avg/max/mdev = 0.403/0.496/0.592/0.073 ms
```

RadSec Certificate Configuration

– RadSec Certificate Installation

1. Create a new certificate on the switch

```
VNBT-Access(config)# crypto pki certificate radsec
VNBT-Access(config-cert-radsec)# subject
  common-name  Specify common name.
  country      Specify the two letter ISO 3166-1 country code
  locality     Specify locality
  org          Specify organization
  org-unit     Specify organization unit
  state        Specify state
  <cr>
VNBT-Access(config-cert-radsec)# subject
Do you want to use the switch serial number as the common name (y/n)?
y
Common Name: SG06KWN00J
Org Unit:
Org Name:
Locality:
State:
Country:
```

RadSec Certificate Configuration

– RadSec Certificate Installation

2. Create a certificate signing request (CSR)

```
VNBT-Access(config-cert-radsec)# enroll terminal
You are enrolling a certificate with the following attributes:
Subject: C=<empty>, ST=<empty>, L=<empty>, OU=<empty>, O=<empty>,
         CN=SG06KWN00J
Key Type: RSA (2048)

Continue (y/n)? y

-----BEGIN CERTIFICATE REQUEST-----
MIICWjCCAUICAQIwFTETMBEGA1UEAwwKU0cwNktXTjAwSjCCASIwDQYJKoZIhvcN
AQEBBQADggEPADCCAQoCggEBAJBUcH1jXrDI+dF9yXZg6SPacVhZY9zfZo7BJBRn
r+mcyjyrluizpZxKWdwaSnSke+Cn+KufnO+9GwFF3wd/V4tkibxUiAy4XKspqwfb
OSFnr8ASgjeBwC8A3MhcxN/3x0a9G83rQxV9rpHi0dEInKab2hHSDQkNtooGPQTO
wFBf3OqvgHvXplIQ9u7WwETNkxccjx2XJdy9JhGO0WwGv6Jr0aCVEJI0GiWJCWTZ
BTQFoHOaThUSfpZyWt0mdtAIE96Qg9IMGuu+KF9A5bcM9M4Ot8HHwK57eah2DKaO
0rQlb5E0Jq0suq+b+bcGWbRGBf+dw84CQ9pghkj6Kd3bHasCAwEAAaAAMA0GCSqG
SIb3DQEBCwUAA4IBAQBNEYCreKC7m2y5tS1TYQcFyZM/L1W7Lppnime1I+GgeMcm
nuFxM61ZmIDEtfLmGexNqf1EJjhSdBHfA75shWFzeUCa4KezvoWcXqMqt/X3Xnz+
T/xt2IYnZUpuShOOGJrF4gWZPsc1k2lQqJ5J2blKWfJUaiGWdyPhiZH3s3DxvSyX
KP640108ZZlRMIVK+eiChcGQm6AwzUkl4Rwb/xCKAkBY6mYAWCIrOMgsdNs5uVVf
EeaaI+LKzYF6NT1vVfkV32z2jE09N198sApjKZUWdhq6c/A4oBNtH4KzajsMAbsx
En8kM/AR2B8cUS1agxpqZxz4vJiwrJ0rs0bN2Sy0
-----END CERTIFICATE REQUEST-----
```

RadSec Certificate Configuration

– RadSec Certificate Installation

3. Sign certificate with external Certificate Authority and import signed certificate

```
VNBT-Access(config-cert-radsec)# import terminal ta-profile cppm
Paste the certificate in PEM format below, then hit enter and ctrl-D:
VNBT-Access(config-cert-import)# -----BEGIN CERTIFICATE-----
ADBEMRMwEQYKCZImiZPyLGQBGRYDbmV0MRYwFAYKCZImiZPyLGQBGRYGdG1lbGFi
ikQvgegS06OBvezzkqKMuKqTQvWZshAgxTxBC0nfCJid4abB4LE=

←----------Certificate Contents ---------------------→

VNBT-Access(config-cert-import)# -----END CERTIFICATE-----
VNBT-Access(config-cert-import)#
Leaf certificate is validated with cppm and imported successfully.
```

4. Associate certificate to be used for RadSec clients and validate that the certificate was applied

```
VNBT-Access(config)# crypto pki application radsec-client certificate radsec
VNBT-Access(config)# show crypto pki certificate

Certificate Name                 Cert Status    EST Status        Associated Applications
-------------------------------- -------------- ----------------- -------------------------------
radsec                           installed      n/a               radsec-client
local-cert                       installed      n/a               captive-portal, est-client, https-server,
syslog-client
device-identity                  installed      n/a               none
```

RadSec Certificate Configuration

– Verify Certificate and Connection

5. Validate RadSec connection

```
VNBT-Access(config)# show radius-server detail
******* Global RADIUS Configuration *******

Shared-Secret: None
Timeout: 5
Auth-Type: pap
Retries: 3
TLS Timeout: 5
Tracking Time Interval (seconds): 300
Tracking Retries: 3
Tracking User-name: radius-tracking-user
Tracking Password: None
Number of Servers: 2
****** RADIUS Server Information ******
Server-Name               : 10.5.8.12
Auth-Port                 : 2083
Accounting-Port           : 2083
VRF                       : VRF1
TLS Enabled               : Yes
TLS Connection Status     : tls_connection_established
Timeout                   : 5
Auth-Type                 : pap
Server-Group              : radsec
```

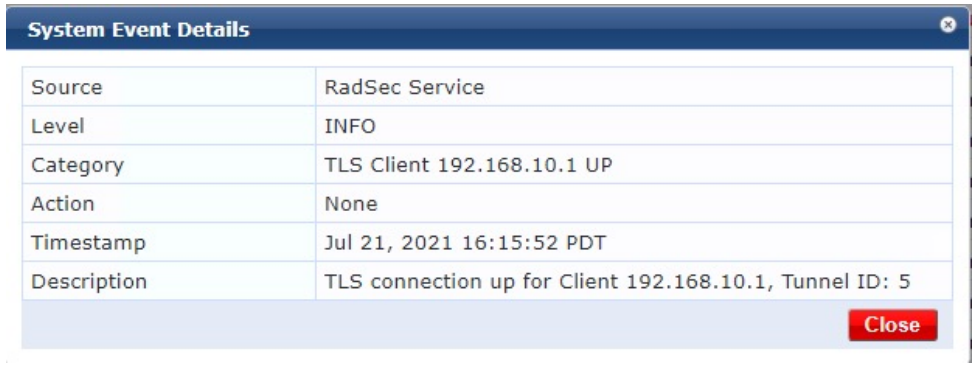

Troubleshooting RadSec

```
****** RADIUS Server Information ******
Server-Name                 : aoss-cppm.tmelab.net
Auth-Port                   : 2083
Accounting-Port             : 2083
VRF                         : default
TLS Enabled                 : Yes
TLS Connection Status       : tls_connection_failed
```

– debug radius

– debug pki

– SSL Errors (Debug Output)

```
2021-07-13:17:19:17.546168|port-accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_RADIUS|SSL_FUNCTION:
function: SSL_connect failed, error string: error:14094418:SSL routines:func(148):reason(1048),
error: 1
```

– https://boringssl.googlesource.com/boringssl/+/HEAD/include/openssl/ssl.h (starting on line number 5338)

  – Boring SSL fork of Open SSL

```
#define SSL_R_TLSV1_ALERT_UNKNOWN_CA 1048
```

  – Needed to generate CSR from Leaf Cert in switch and signed by CA

IPv4/IPv6 RadSec over VXLAN Demo

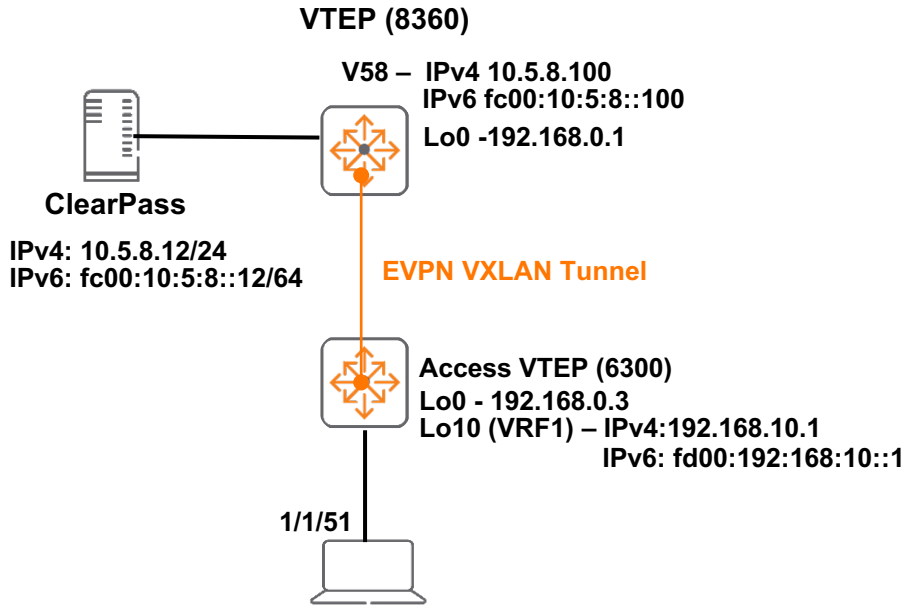

Use Case 1

- Demo IPv4 RadSec over VXLAN for Network Users

Use Case 2

- Demo IPv6 RADIUS over VXLAN for Management Users

IPv4 Multicast VXLAN

Overlay RP on VSX VTEPs

- Allows VSX VTEPs to function as overlay RPs
- Previous releases only supported overlay RP on standalone and VSF switches
- Supported platforms: 6400, 8360

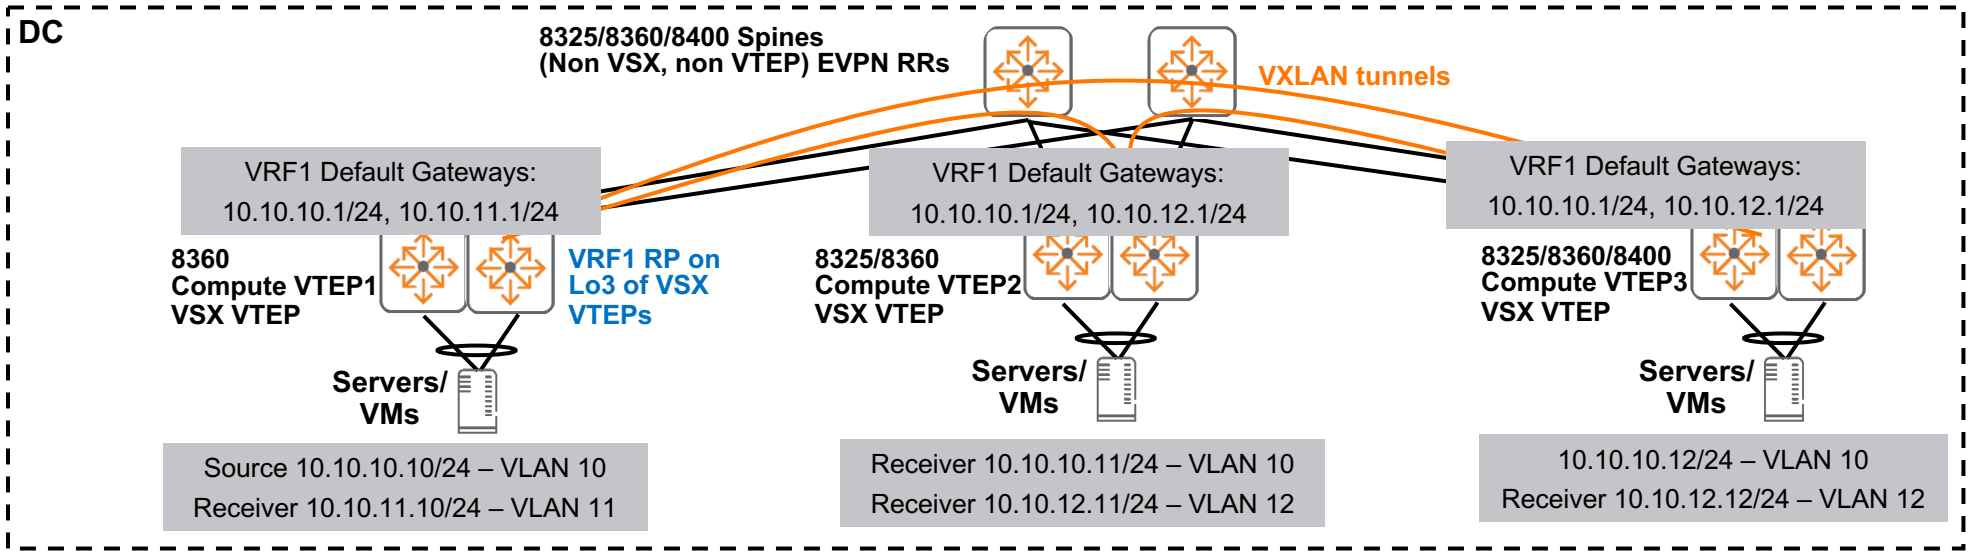

Overlay RP on VSX VTEPs (Sample Configs)

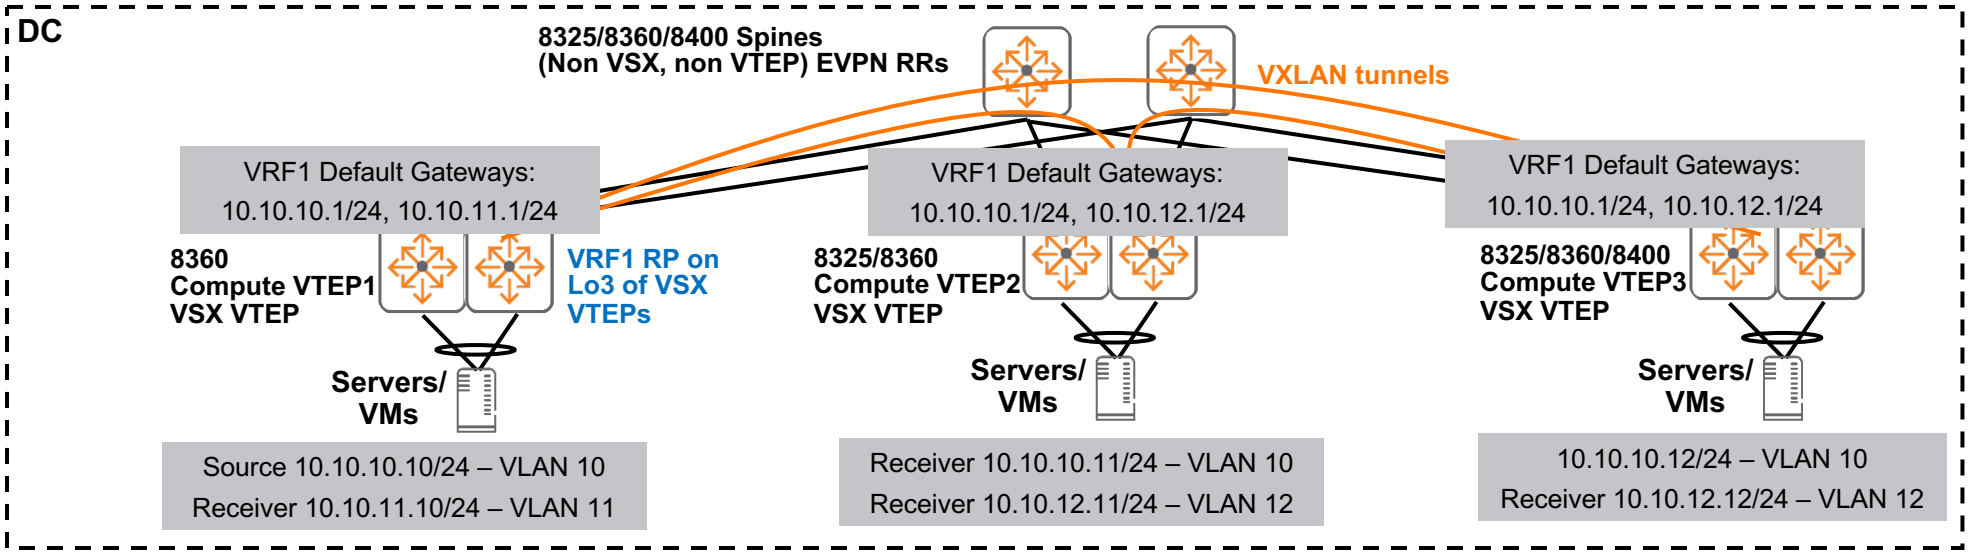

VSX VTEP (Primary) with preferred RP and BSR candidates

```
interface loopback 3
    vrf attach VRF1
    ip address 40.1.1.1/32
    ip ospf 1 area 0.0.0.0
    ip pim-sparse enable

router pim vrf VRF1
    enable
    rp-candidate source-ip-interface loopback3 group-prefix 224.0.0.0/4
    rp-candidate priority 200
    bsr-candidate source-ip-interface loopback3
    bsr-candidate priority 100
```

VSX VTEP (Secondary)

```
interface loopback 3
    vrf attach VRF1
    ip address 41.1.1.1/32
    ip ospf 1 area 0.0.0.0
    ip pim-sparse enable

router pim vrf VRF1
    enable
    rp-candidate source-ip-interface loopback3 group-prefix 224.0.0.0/4
    bsr-candidate source-ip-interface loopback3
```

VXLAN - IPv4 DCI Use Case Validation

VXLAN - IPv4 DCI Use Case Validation

- QA resources and time were allocated to validate VXLAN Data Center Interconnect (DCI) based on current features
- Supported platforms: 8325, 8360

Use Case 1 – Dedicated L2 VTEP (L2 DCI over VXLAN)

- Bridge L2 traffic/VLANs between DCs/Sites within a campus
- L2 VTEPs would connect to other network devices (e.g. L3 core switch via 802.1Q trunks)
- L3 routing functionality maintained by other devices, e.g. 3rd party L3 core switch

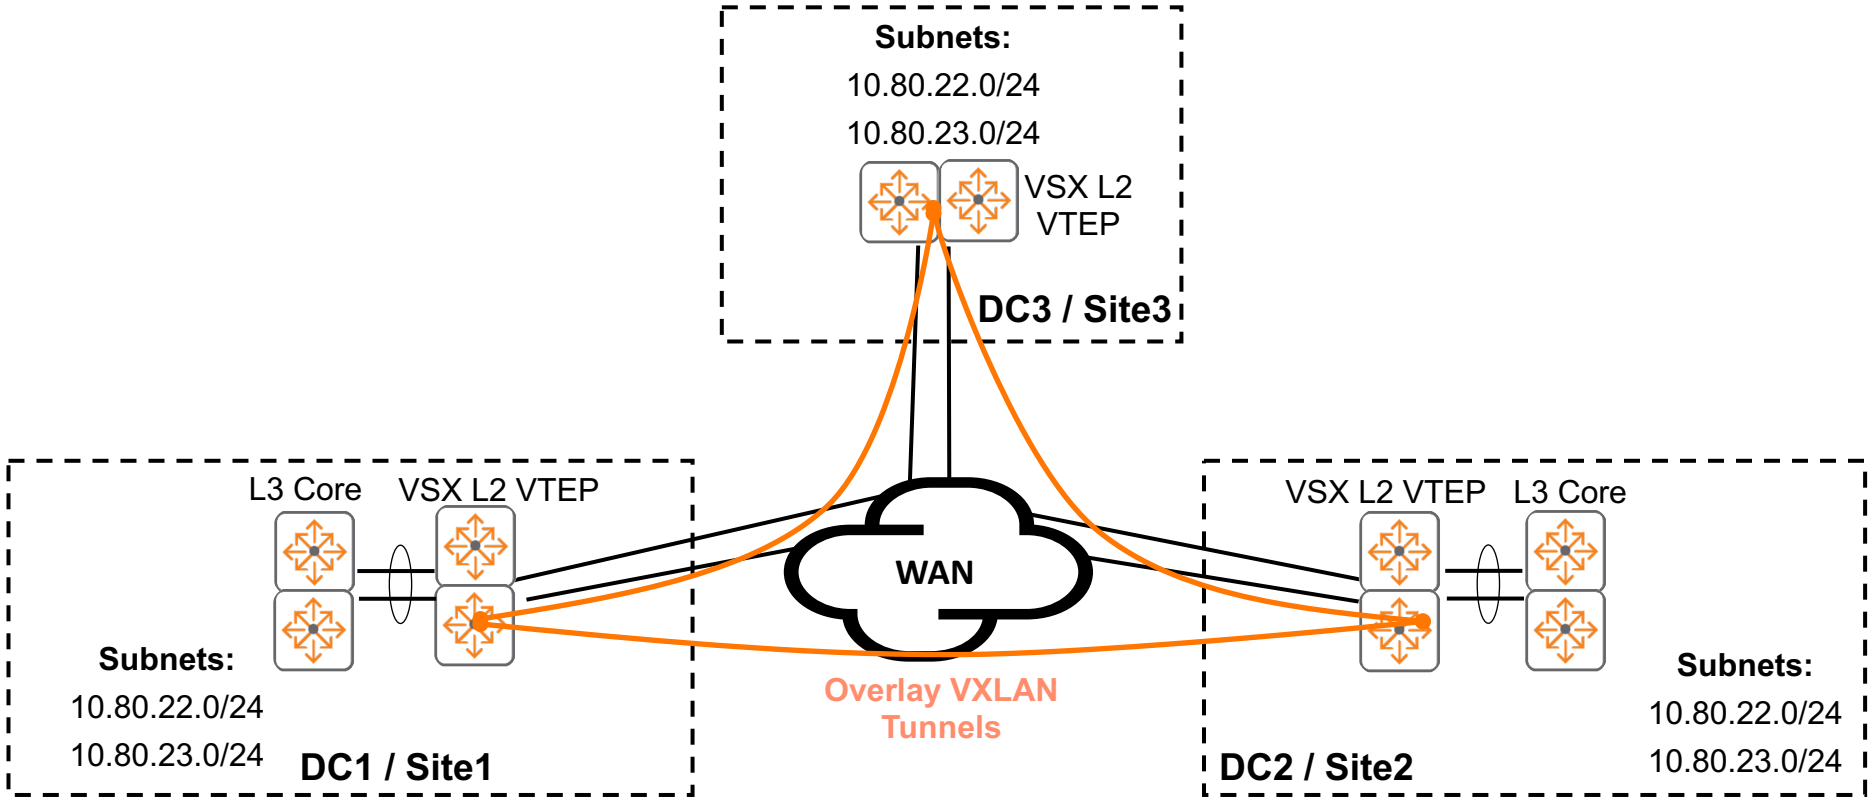

Use Case 2 – Distributed L3 Gateways (L2/L3 DCI over VXLAN)

- Small scale DCs
- VRRP not required for distributed default gateways (uses anycast distributed L3 gateways)
- ARP suppression available to minimize ARP broadcasts across WAN
- Supports inter-DC L2/L3 unicast and L2/3 multicast traffic

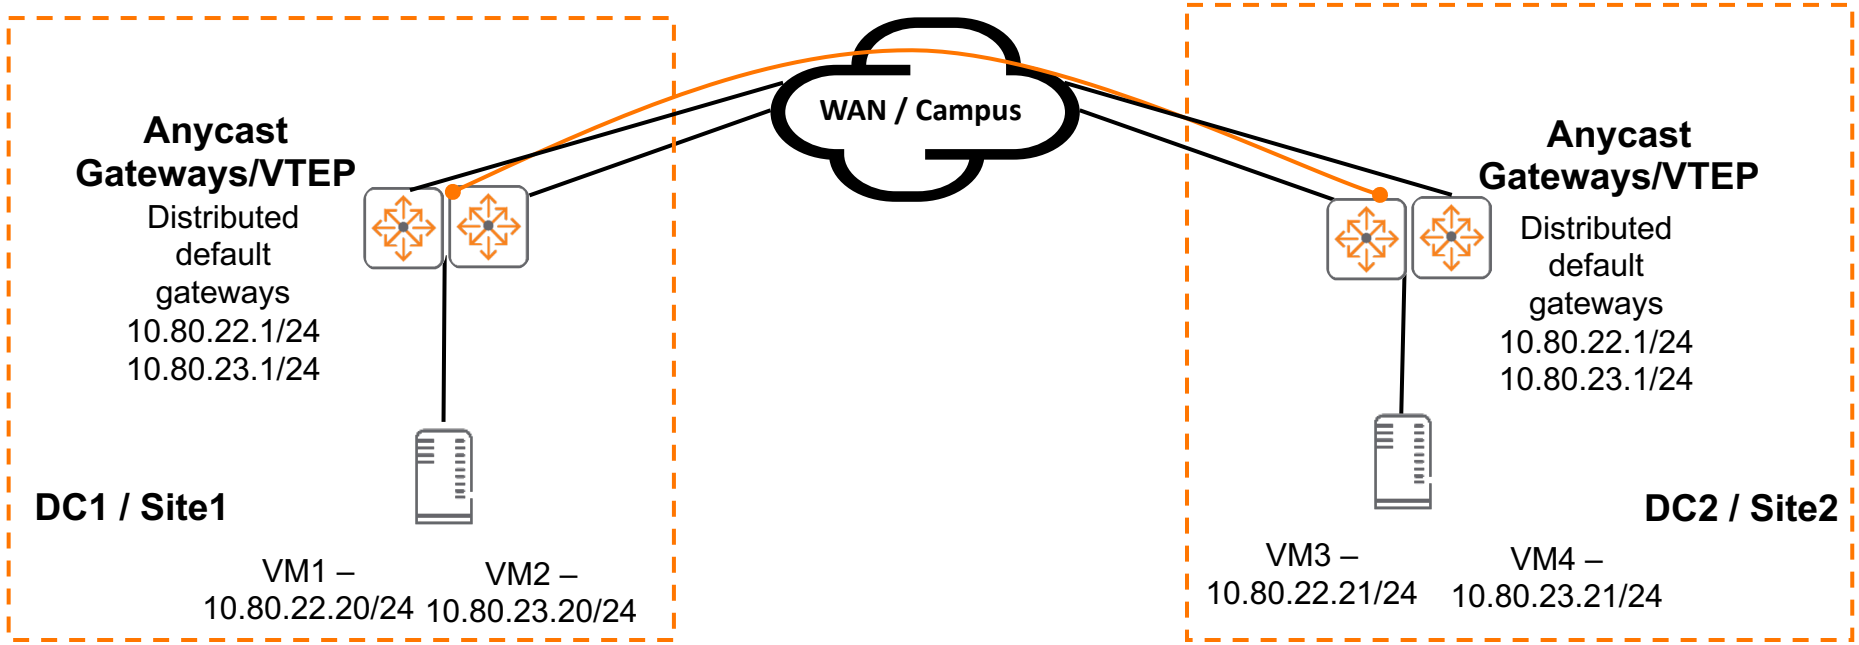

Use Case 2 – Distributed L3 Gateways (L2/L3 DCI over VXLAN)

- Expand on use case 2 – more racks/VTEPs, VTEPs are fully meshed
- VRRP not required for distributed default gateways (uses Anycast distributed L3 gateways)
- ARP suppression available to minimize ARP broadcasts across WAN
- Supports inter-DC L2/L3 unicast and L2/3 multicast traffic

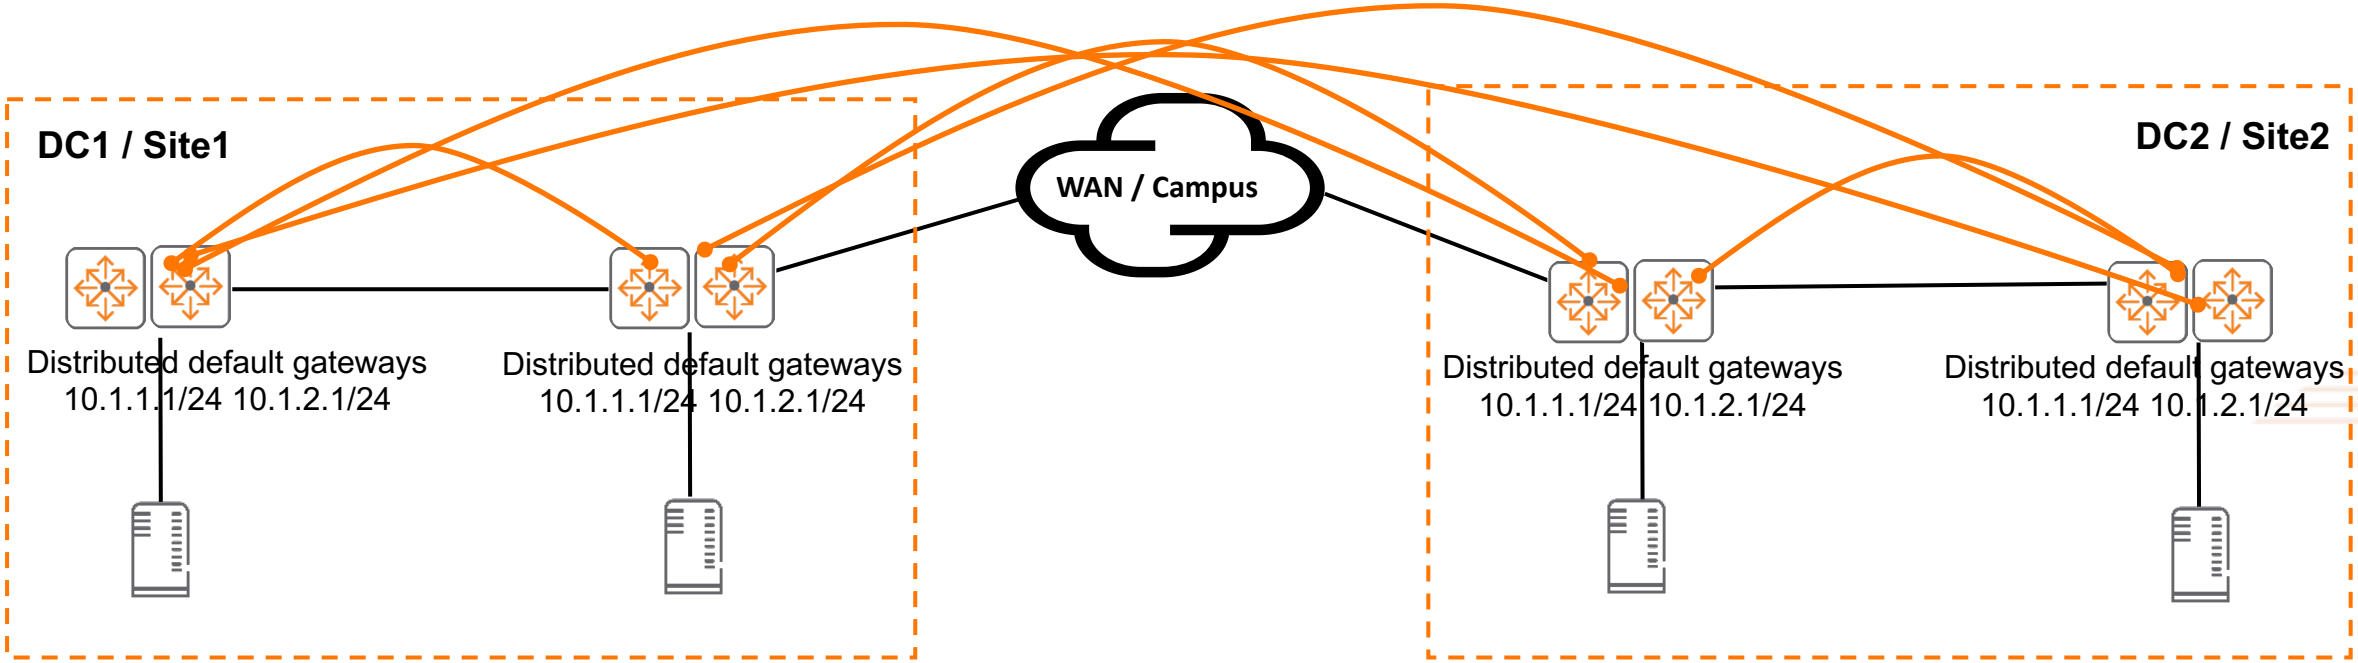

aruba
a Hewlett Packard Enterprise company

Thank you

daryl.wan@hpe.com
justin.noonan@hpe.com