EVPN Campus Design and Implementation

Sergey Nasonov
Solutions Engineer
snasonov@cisco.com

CCIE R&S 62572

Session ID: BRKENS-2041

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until 14 November 2025.

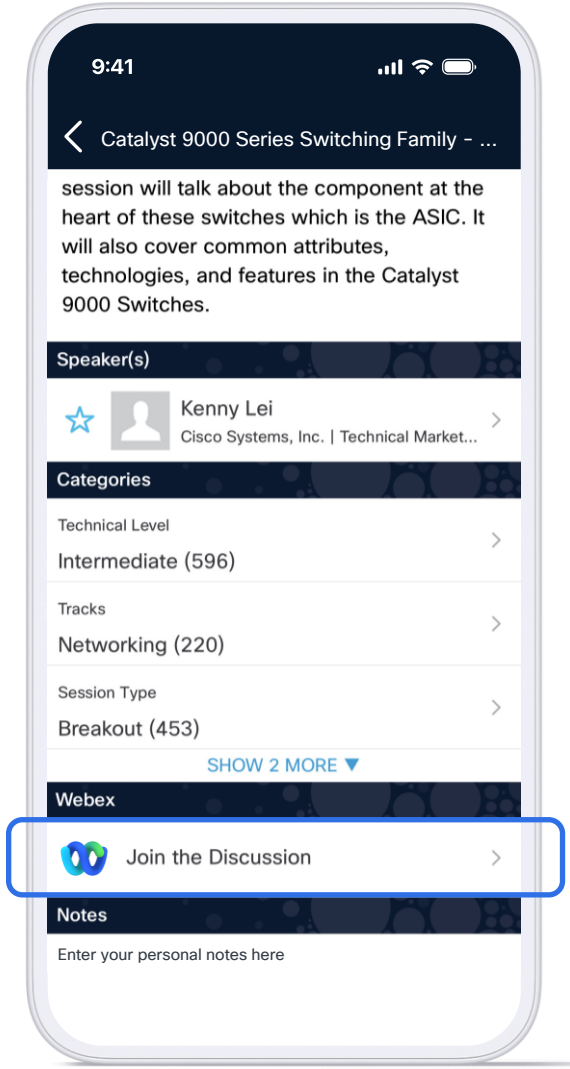

https://ciscolive.ciscoevents.com/
ciscolivebot/**#BRKENS-2041**

Agenda

01 How BGP EVPN Works

02 Campus EVPN Architecture

03 Design Considerations

04 Implementation Details

05 Multicast in EVPN

06 Closing

How EVPN Works

EVPN Enables Fabric Overlay

- Ethernet VPN (EVPN) is an address family in BGP, intended to carry Layer 2 (MAC Address) and Layer 3 (IP address) in the overlay over transport network.
- An overlay is a logical topology used to virtually connect devices, built on top of physical underlay topology.
- “We can solve any problem by introducing an extra level of indirection.” © David Wheeler.

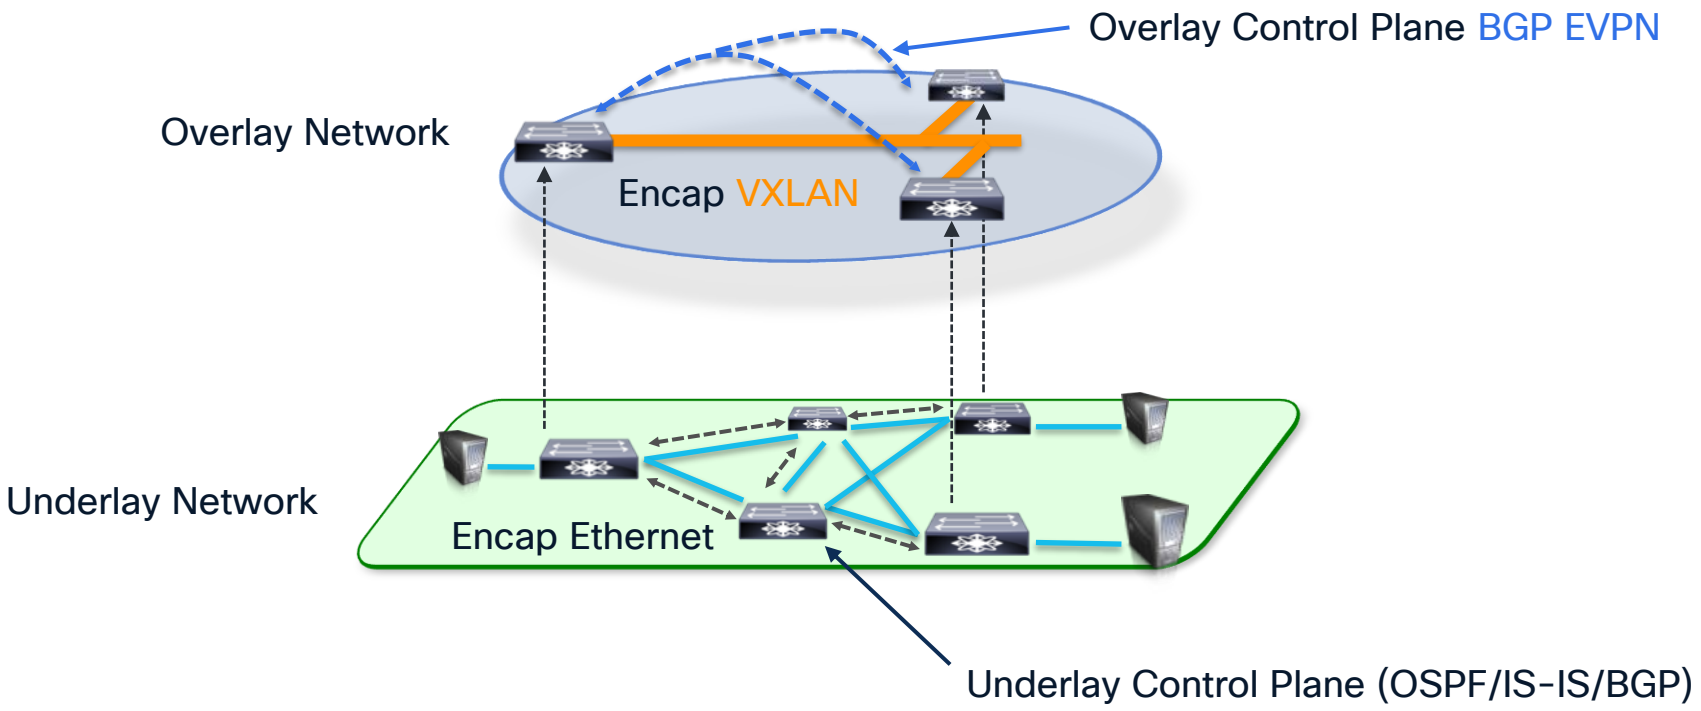

Overlay Data Plane - VXLAN

Also known as MAC-in-UDP

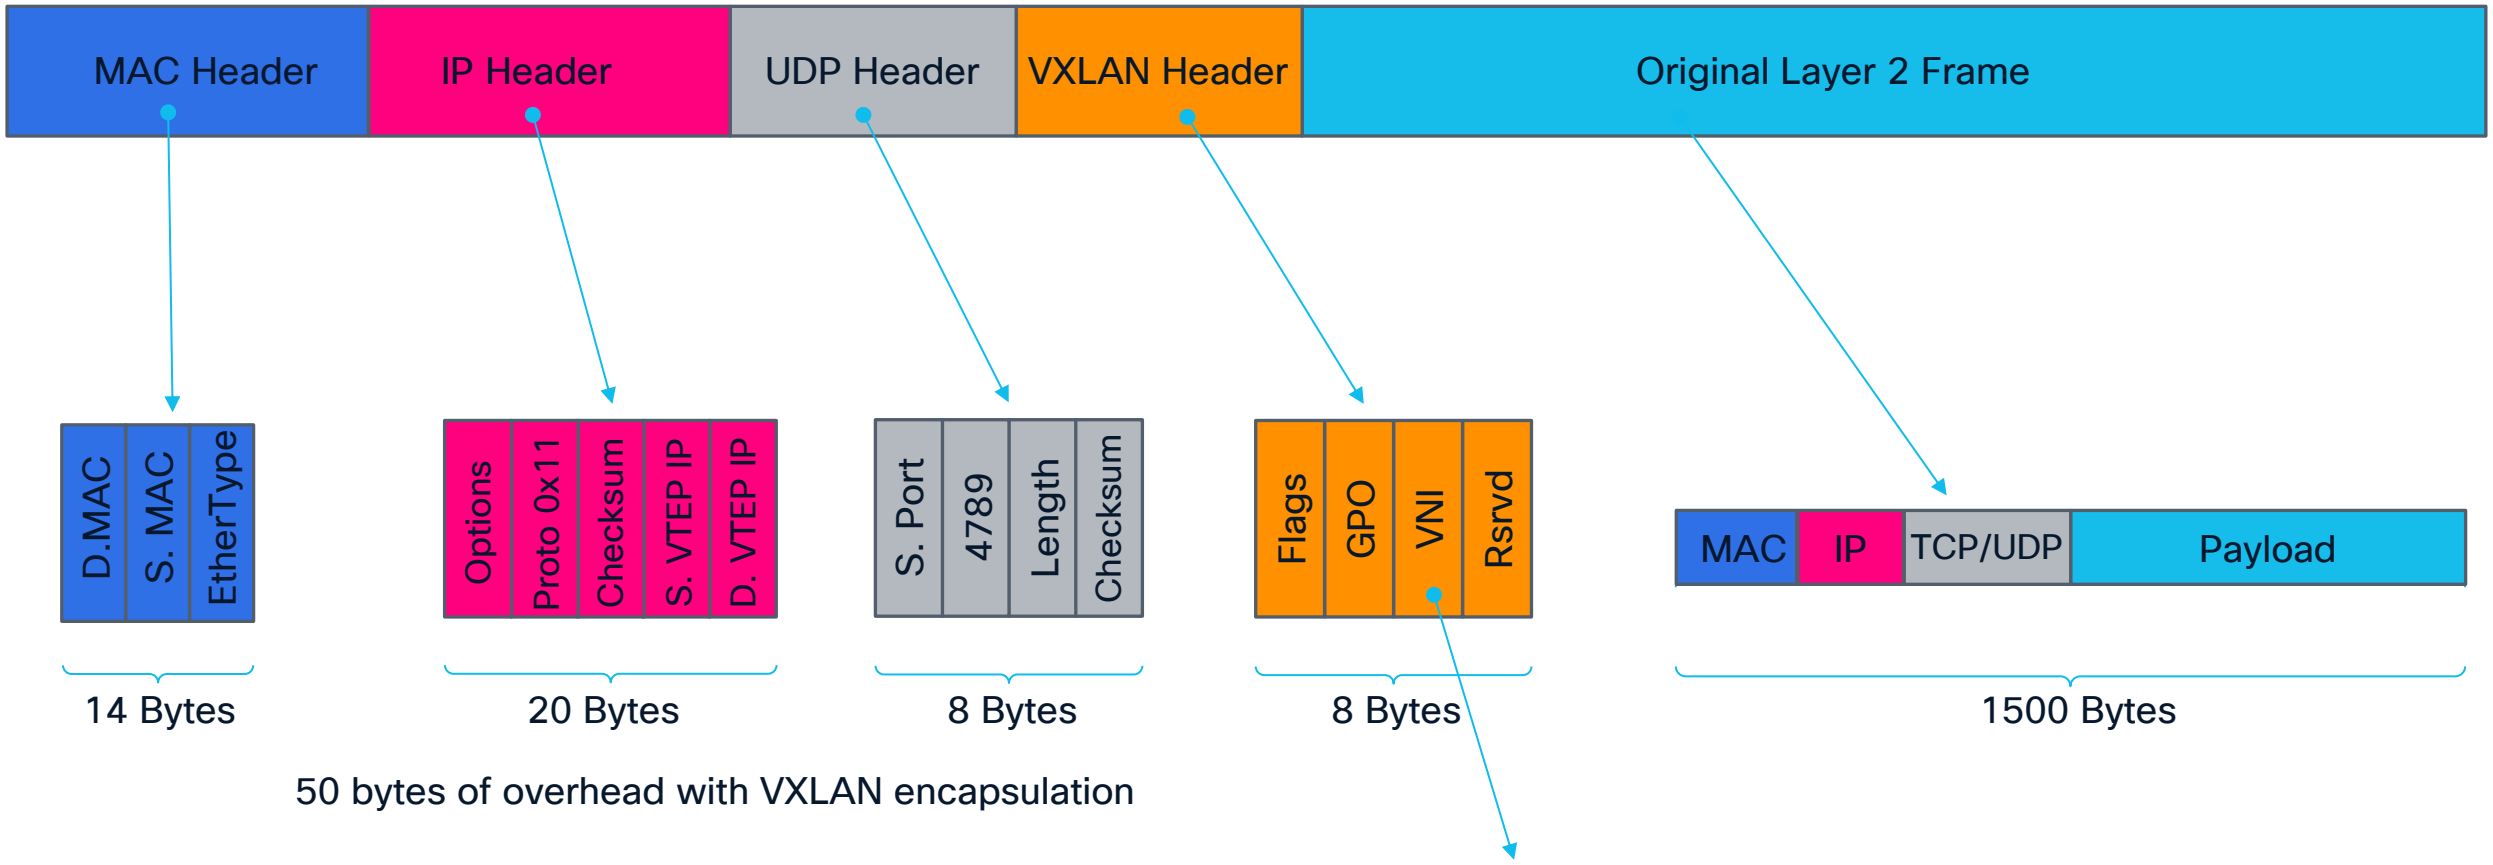

50 bytes of overhead with VXLAN encapsulation

VNI – 24-bit field (16 million segments). First 4k are not assignable.
C9300 – 1k VNI (L2+L3)
C9500 – 1k VNI (L2+L3) / C9500X – 4k

cisco

Overlay Control Plane – BGP EVPN

Layer 2 Stretch with Distributed Anycast Gateway (DAG) – Starting Topology

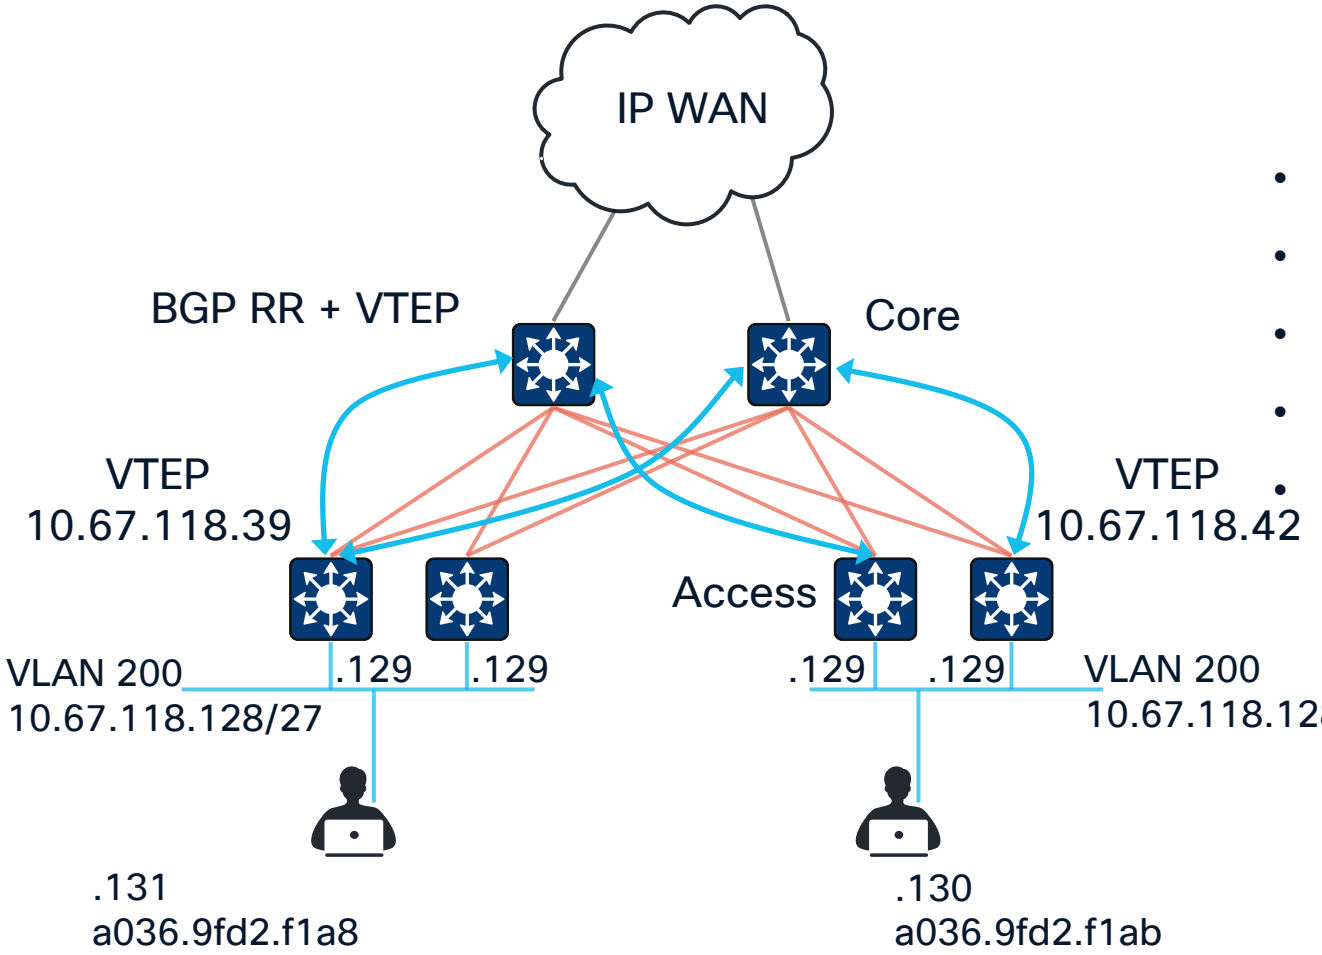

- iBGP between Access and Core
- BGP peering is between Loopback0 interfaces
- Core switches are BGP RRs
- OSPF in the underlay
- Anycast Gateways at each access switch for VLAN 200

Overlay Control Plane – BGP EVPN

Layer 2 Stretch with Distributed Anycast Gateway (DAG) – Local Advertisement

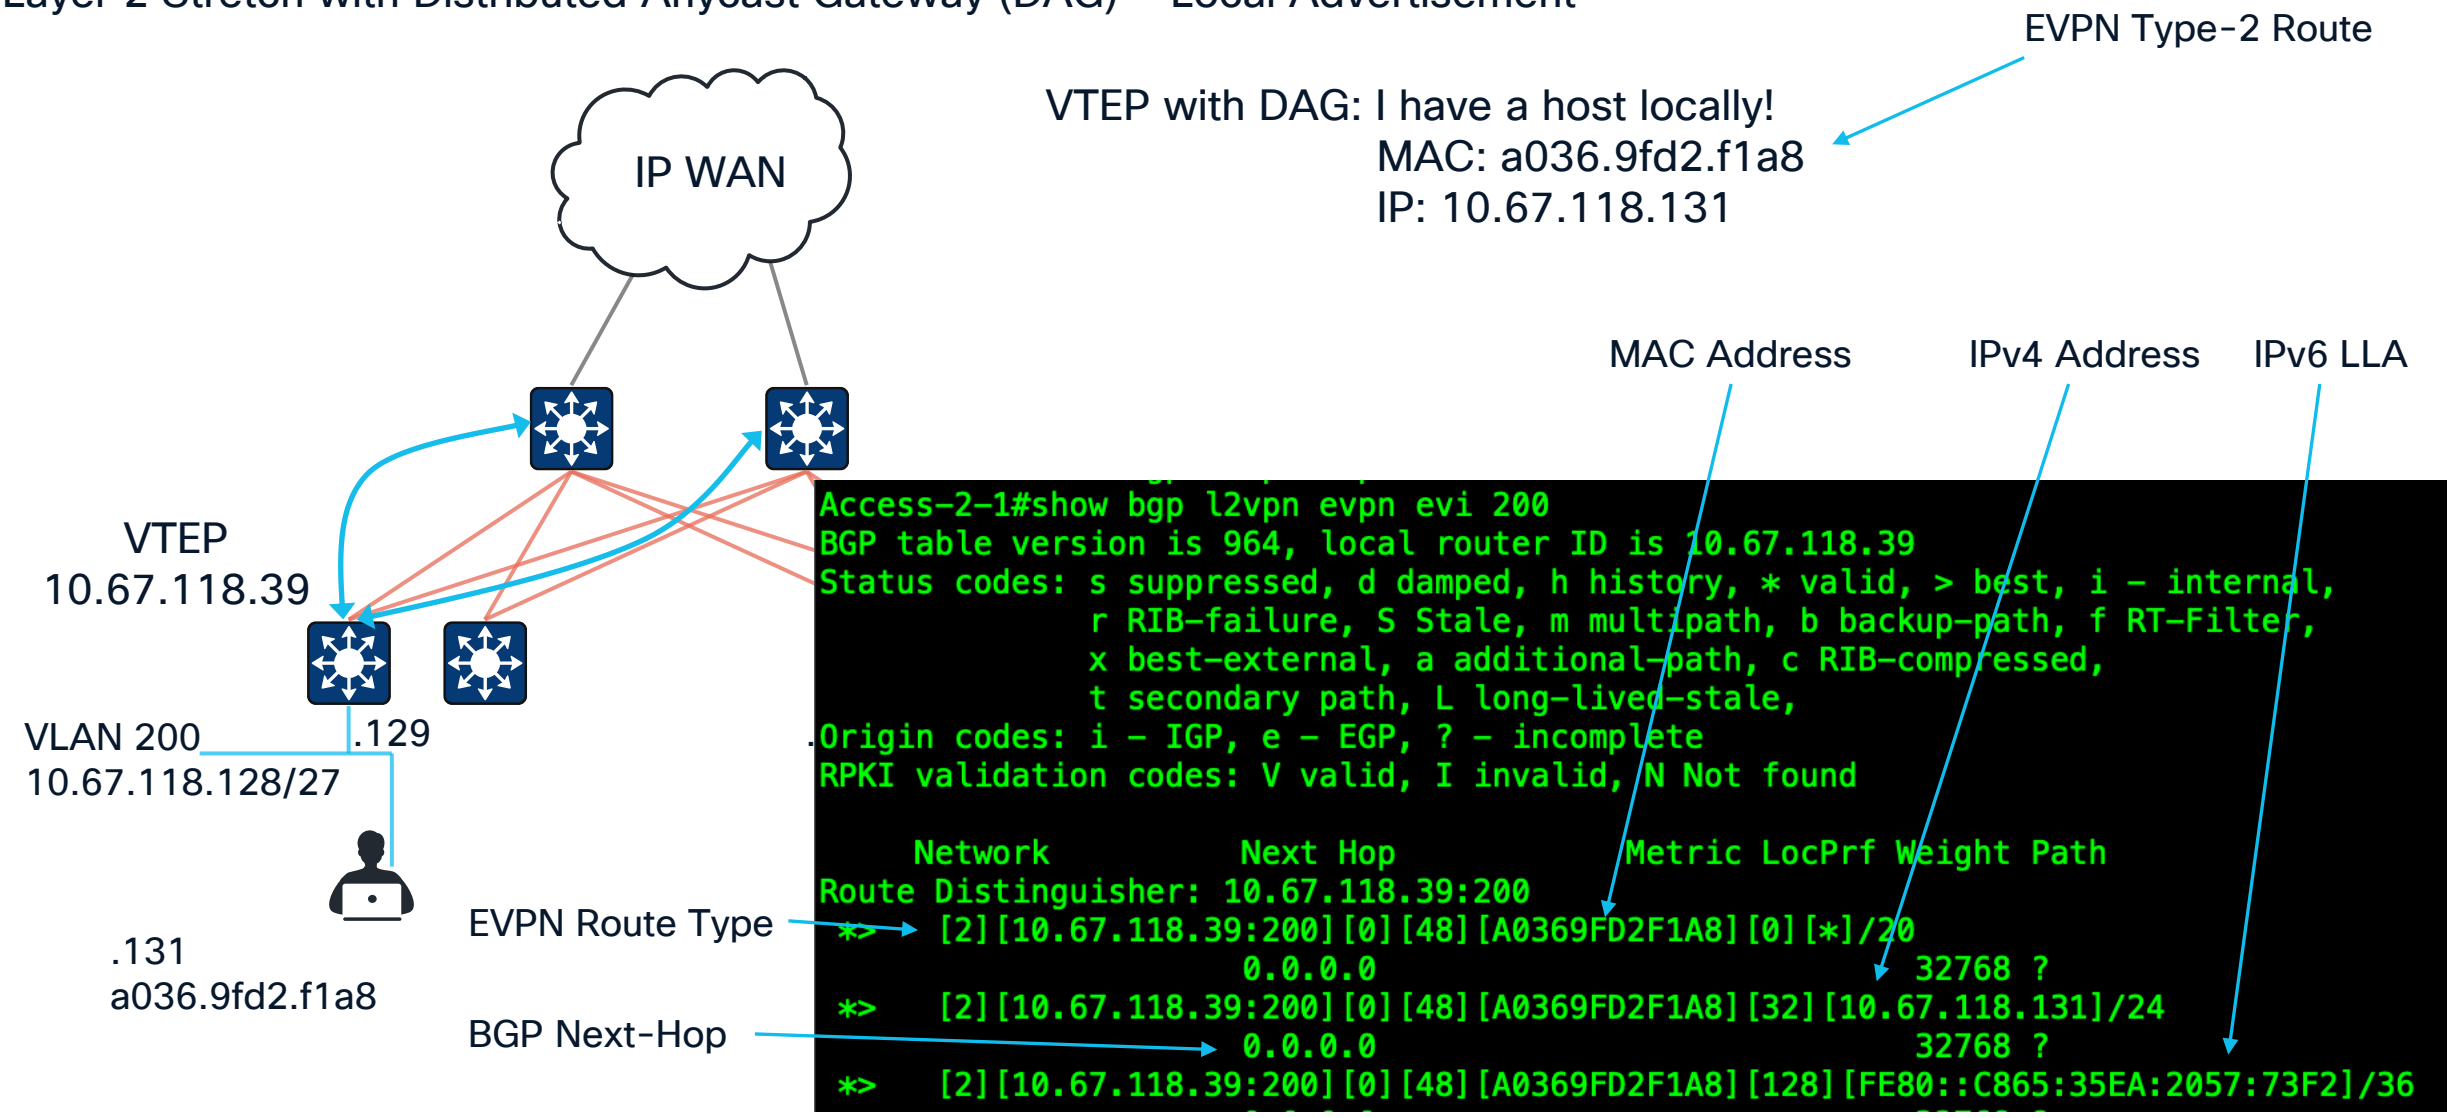

cisco

Overlay Control Plane – BGP EVPN

Layer 2 Stretch with Distributed Anycast Gateway (DAG) – Rest of BGP speakers/VTEPs

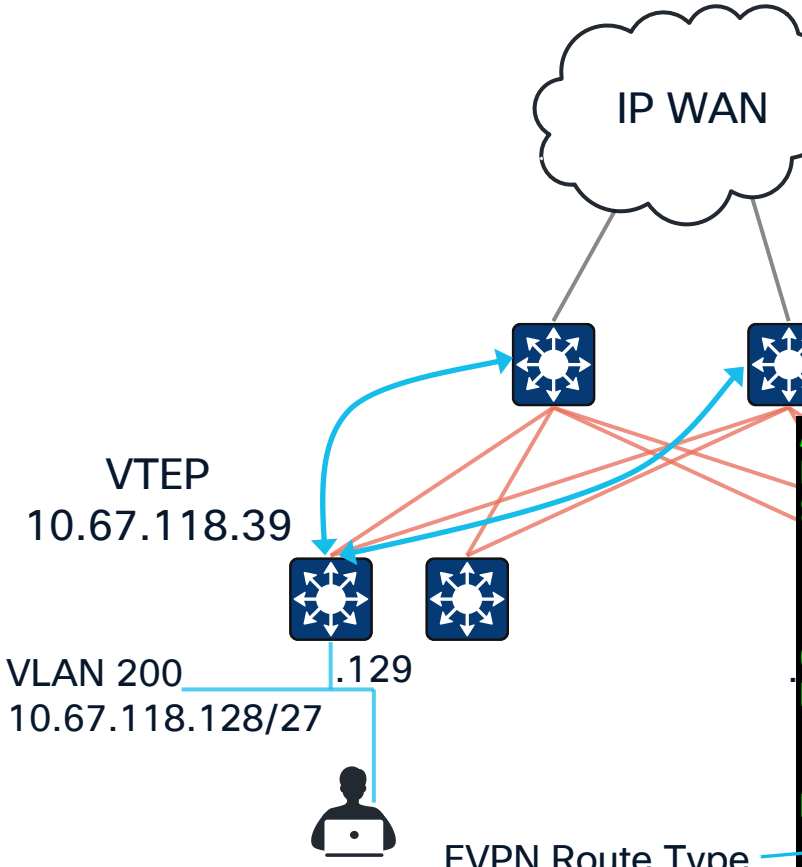

MAC Address

IPv4 Address

IPv6 LLA

EVPN Route Type

BGP Next-Hop

```
Access-2-2#show bgp l2vpn evpn evi 200
BGP table version is 1131, local router ID is 10.67.118.40
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 10.67.118.40:200
 *>i  [2][10.67.118.40:200][0][48][A0369FD2F1A8][0][*]/20
                       10.67.118.39             0    100      0 ?
 *>i  [2][10.67.118.40:200][0][48][A0369FD2F1A8][32][10.67.118.131]/24
                       10.67.118.39             0    100      0 ?
 *>i  [2][10.67.118.40:200][0][48][A0369FD2F1A8][128][FE80::C865:35EA:2057:73F2]/36
                       10.67.118.39             0    100      0 ?
```


cisco

Overlay Control Plane – BGP EVPN

Problem with BUM Traffic – what happens if control plane does not have IP address mapping for a given MAC?

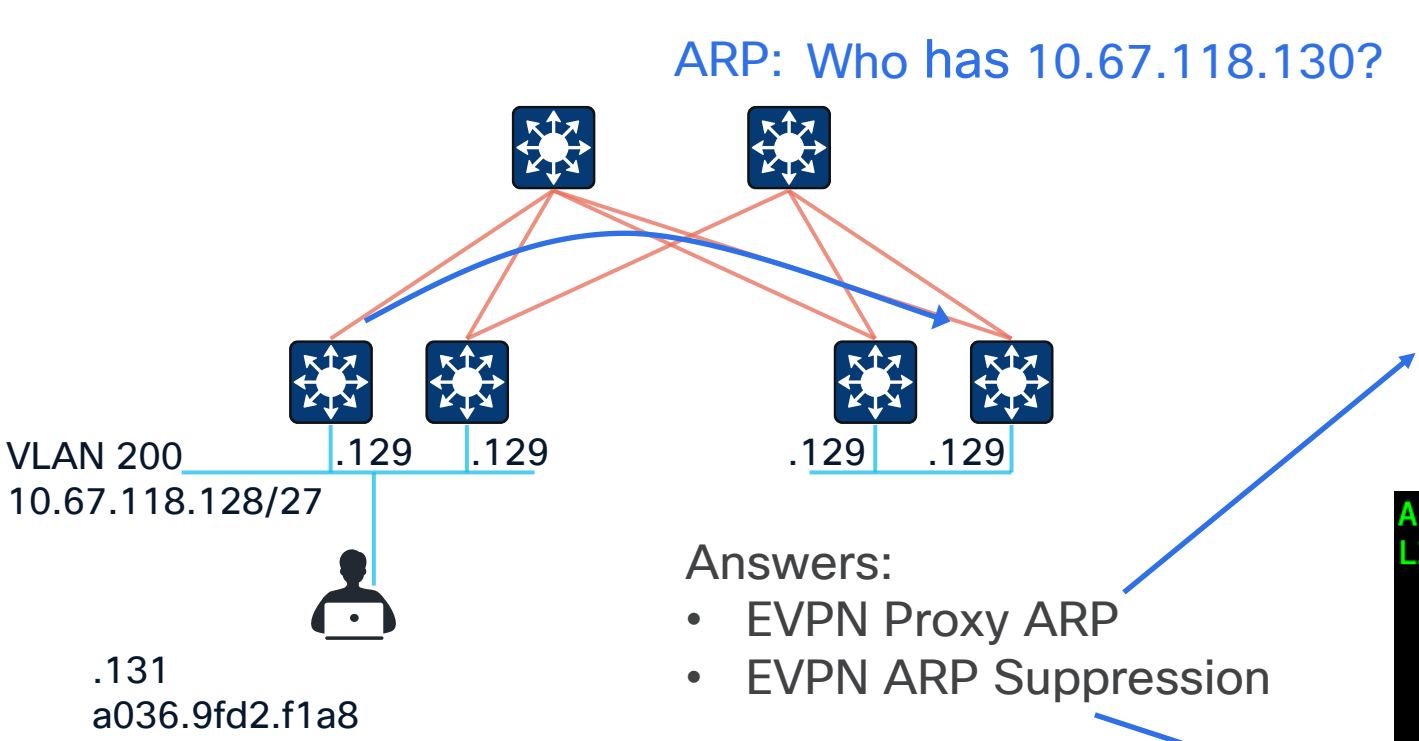

Answers:

- EVPN Proxy ARP
- EVPN ARP Suppression

```
interface Vlan100
 mac-address 0000.0000.0100
 vrf forwarding CORP
 ip dhcp relay information option vpn-id
 ip address 10.67.118.65 255.255.255.224
 ip helper-address global 10.66.119.56
 ip helper-address 10.66.119.57
 no ip redirects
 ip local-proxy-arp
 ip pim sparse-mode
 ip route-cache same-interface
 ip igmp version 3
```

```
Access-2-2#show l2vpn evpn summary
L2VPN EVPN
  EVPN Instances (excluding point-to-point): 3
    VLAN Based:   3
  Vlans: 3
  BGP: ASN 64513, address-family l2vpn evpn configured
  Router ID: 10.67.118.40
  Global Replication Type: Static
  ARP/ND Flooding Suppression: Enabled
  Connectivity to Core: UP
  MAC Duplication: seconds 180 limit 5
```

Overlay Control Plane – BGP EVPN

Problem with BUM Traffic – what happens if control plane does not have IP address mapping for a given MAC?

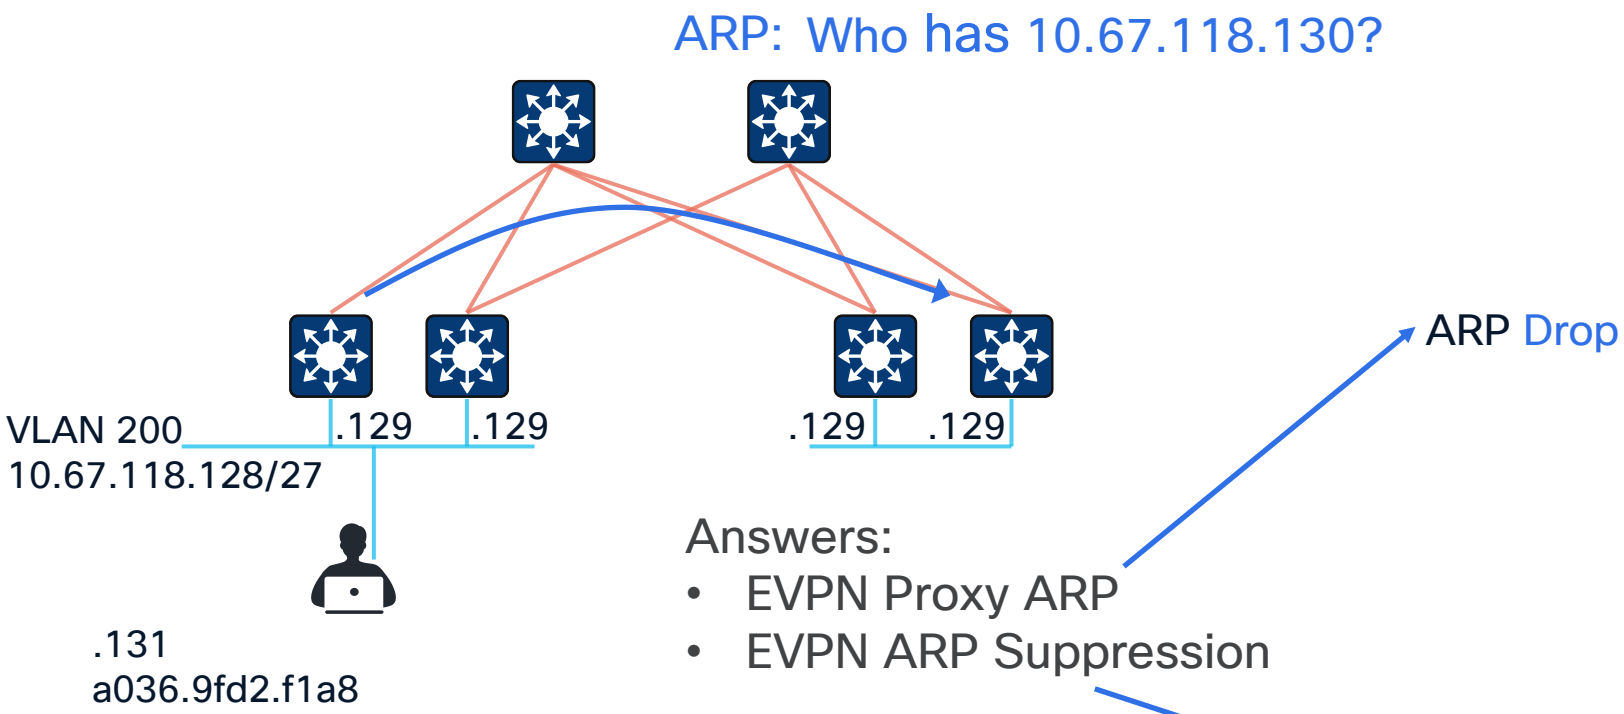

ARP Drop

Answers:

- EVPN Proxy ARP
- EVPN ARP Suppression

ARP Flood → How can we flood ARP in the overlay?

cisco

Overlay Control Plane – BGP EVPN

Flooding in Overlay

Ingress Replication (IR) / Head-End Replication (HER)

interface nve1

member vni 10300 ingress-replication

Pros: one-line configuration, no multicast state in the core

Cons: low scalability, need to learn/troubleshoot EVPN Type-3 route

Multicast-enabled underlay

interface nve1

member vni 10100 mcast-group 239.0.17.1

Pros: very high scalability

Cons: multicast-enabled underlay

Putting It All Together – Layer 2 Overlay

1. Create VLAN

```
vlan 200
 name Campus-MRI
```

```
interface Vlan200
 mac-address 0000.0000.0200
 vrf forwarding MEDICAL
 ip dhcp relay information option vpn-id
 ip address 10.67.118.129 255.255.255.224
```

2. Create SVI (Optional)

```
l2vpn evpn instance 200 vlan-based
 encapsulation vxlan
```

3. Create MAC VRF (EVI)

```
vlan configuration 200
 member evpn-instance 200 vni 10200
```

4. Map VLAN to MAC VRF to VNI

```
interface nve1
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
 group-based-policy
 member vni 10200 mcast-group 239.0.17.1
```

5. Configure VTEP interface

```
router bgp 64513
```

```
address-family l2vpn evpn
 neighbor 10.67.118.32 activate
 neighbor 10.67.118.32 send-community both
```

6. Advertise reachability information in BGP

```
address-family ipv4 vrf MEDICAL
 advertise l2vpn evpn
 redistribute connected
 maximum-paths 2
```

Behind The Scenes – Bridging

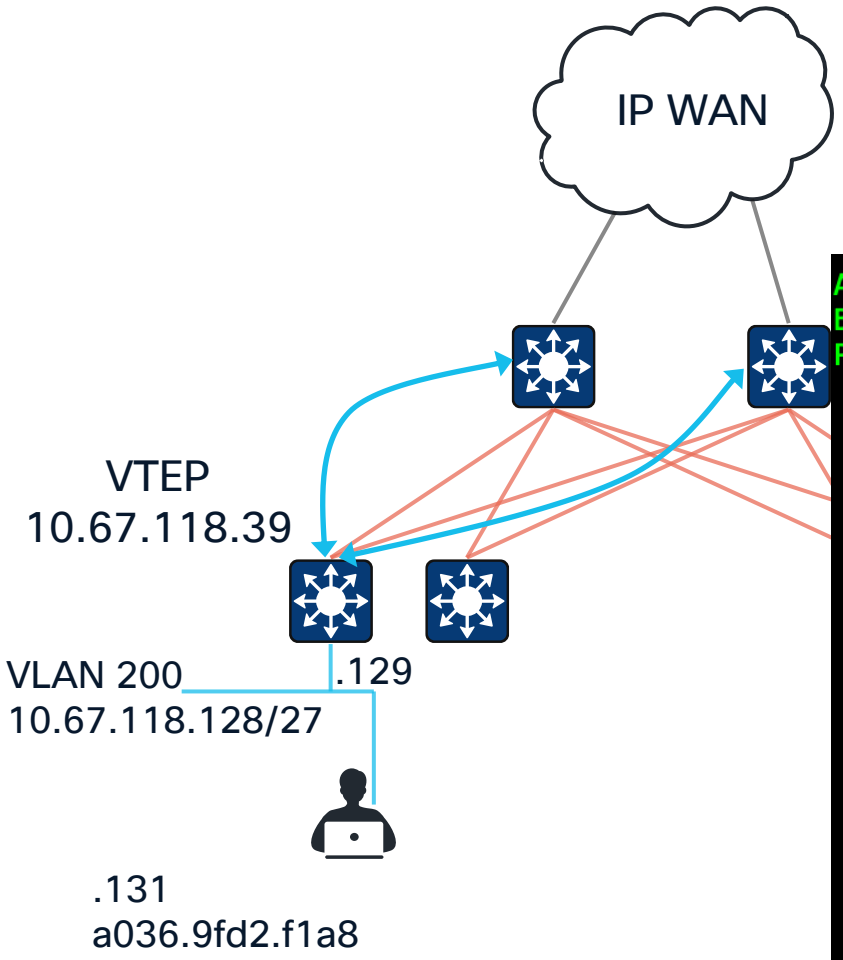

MAC Address

No IP Address (separate route)

Destination VTEP

```
Access-2-2#show bgp l2vpn evpn A0369FD2F1A8
BGP routing table entry for [2][10.67.118.39:200][0][48][A0369FD2F1A8][0][*]/20, version 1177
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    10.67.118.39 (metric 81) (via default) from 10.67.118.33 (10.67.118.33)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Label1 10200
      Extended Community: RT:64513:200 ENCAP:8
      Originator: 10.67.118.39, Cluster list: 10.67.118.33
      rx pathid: 0, tx pathid: 0
      Updated on May 15 2025 00:40:01 AEST
  Refresh Epoch 2
  Local
    10.67.118.39 (metric 81) (via default) from 10.67.118.32 (10.67.118.32)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10200
      Extended Community: RT:64513:200 ENCAP:8
      Originator: 10.67.118.39, Cluster list: 10.67.118.32
      rx pathid: 0, tx pathid: 0x0
      Updated on May 15 2025 00:40:01 AEST
```

VNI (why label?)

Automatic RT (ASN:EVI)

Encapsulation 8: VXLAN

cisco

How Does Routing in the Overlay Work?

- L3 VNI is unique per VRF
- L3 VNI is configured at every switch with given VRF
- During routing lookup, source VTEP will populate L3 VNI in VXLAN header and send the packet to destination VTEP
- Destination VTEP will do bridging lookup across local L2 VNIs and bridge traffic in the correct VLAN.

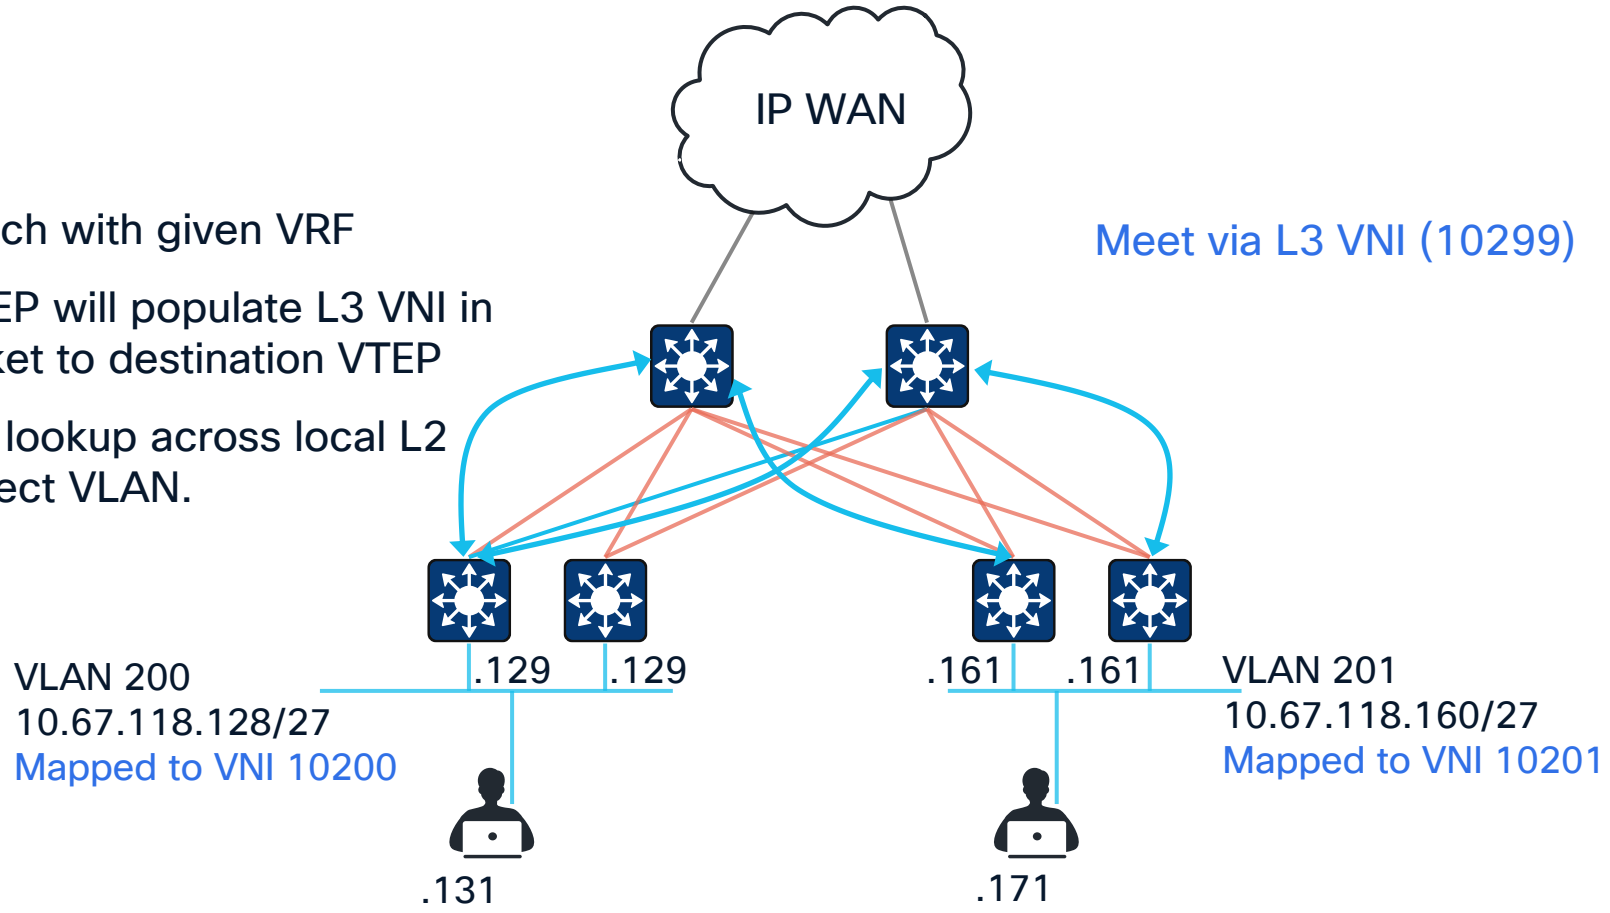

Think of L3 VNI as an inner MPLS label in VPNv4 MPLS VPNs; It is used to identify correct IP VRF at the destination switch.

Behind The Scenes – Routing

VTEP BGP advertisement:

- I have a subnet 10.67.118.128/27
- My VTEP address is 10.67.118.39

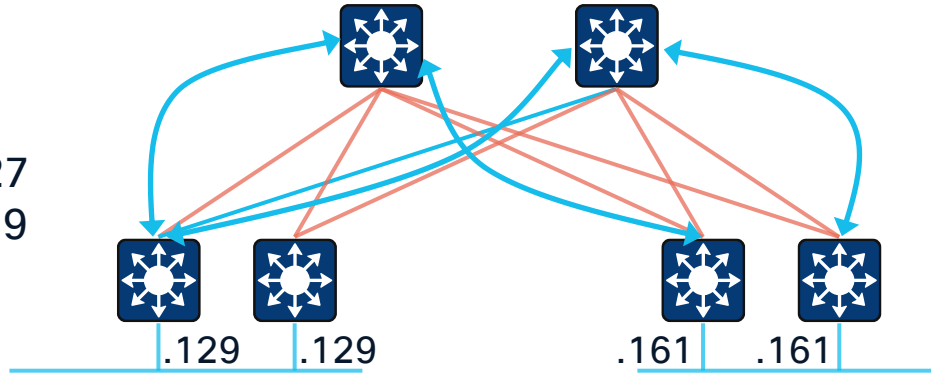

VLAN 200
10.67.118.128/27

Mask/Subnet

L3 VNI

EVPN Route Type

BGP NH / VTEP IP Address

```
Access-2-2#show bgp l2vpn evpn 10.67.118.128/27
BGP routing table entry for [5][10.67.118.39:10299][0][27][10.67.118.128]/17, version 1264
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    10.67.118.39 (metric 81) (via default) from 10.67.118.33 (10.67.118.33)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 10299, MPLS VPN Label 0
      Extended Community: RT:64513:10299 ENCAP:8 Router MAC:706B.B928.4D52
      Originator: 10.67.118.39, Cluster list: 10.67.118.33
      rx pathid: 0, tx pathid: 0
      Updated on May 15 2025 00:40:03 AEST
  Refresh Epoch 2
  Local
    10.67.118.39 (metric 81) (via default) from 10.67.118.32 (10.67.118.32)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 10299, MPLS VPN Label 0
      Extended Community: RT:64513:10299 ENCAP:8 Router MAC:706B.B928.4D52
      Originator: 10.67.118.39, Cluster list: 10.67.118.32
      rx pathid: 0, tx pathid: 0x0
      Updated on May 15 2025 00:40:03 AEST
```

Putting It All Together – Routing in Overlay

1. Define IP VRF and L3 RTs

```
vrf definition MEDICAL
 rd 10.67.118.39:10299
 !
 address-family ipv4
  route-target export 64513:10299
  route-target import 64513:10299
  route-target export 64513:10299 stitching
  route-target import 64513:10299 stitching
 exit-address-family
```

2.Define Core VLAN/SVI

```
vlan 299
 name VRF_MEDICAL_CORE_VLAN
```

```
interface Vlan299
 vrf forwarding MEDICAL
 ip unnumbered Loopback0
 ip pim sparse-mode
 no autostate
```

3. Map Core VLAN to L3 VNI

```
vlan configuration 299
 member vni 10299
```

4. Configure L3 VNI under VTEP interface

```
interface nve1
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
 group-based-policy
 member vni 10199 vrf CORP
 member vni 10299 vrf MEDICAL
```

5. Advertise reachability information in BGP

```
router bgp 64513
```

```
address-family l2vpn evpn
 neighbor 10.67.118.32 activate
 neighbor 10.67.118.32 send-community both
```

```
address-family ipv4 vrf MEDICAL
 advertise l2vpn evpn
 redistribute connected
 maximum-paths 2
```

Connecting to the Outside World

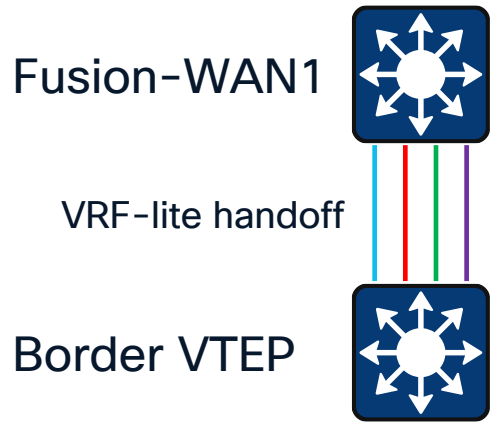

VRF:

```
vrf definition MEDICAL
 rd 10.67.118.32:10299
 !
 address-family ipv4
  route-target export 64513:10299
  route-target import 64513:10299
  route-target export 64513:10299 stitching
  route-target import 64513:10299 stitching
 exit-address-family
```

Core VLAN / VNI:

```
vlan 299
 name VRF_MEDICAL_CORE_VLAN
```

```
vlan configuration 299
 member vni 10299
```

```
interface Vlan299
 vrf forwarding MEDICAL
 ip unnumbered Loopback0
 ip pim sparse-mode
 no autostate
```

VTEP interface:

```
interface nve1
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
 group-based-policy
 member vni 10199 vrf CORP
 member vni 10299 vrf MEDICAL
```

BGP:

```
router bgp 64513
```

```
address-family l2vpn evpn
 neighbor 10.67.118.32 activate
 neighbor 10.67.118.32 send-community both
```

```
address-family ipv4 vrf MEDICAL
 advertise l2vpn evpn
 network 10.67.118.14 mask 255.255.255.254
 aggregate-address 10.67.118.128 255.255.255.192 summary-only
 neighbor 10.67.118.14 remote-as 64512
 neighbor 10.67.118.14 fall-over bfd
 neighbor 10.67.118.14 activate
 neighbor 10.67.118.14 route-map WAN1-INBOUND-FILTER in
 maximum-paths 2
exit-address-family
```

Campus EVPN Architecture

Leaf and Spine* vs Three-Tier

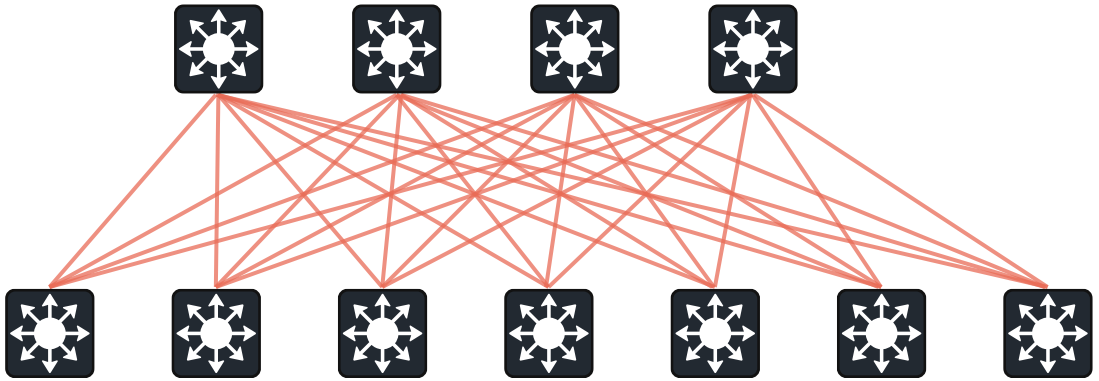

East-West optimised

*Also known as CLOS

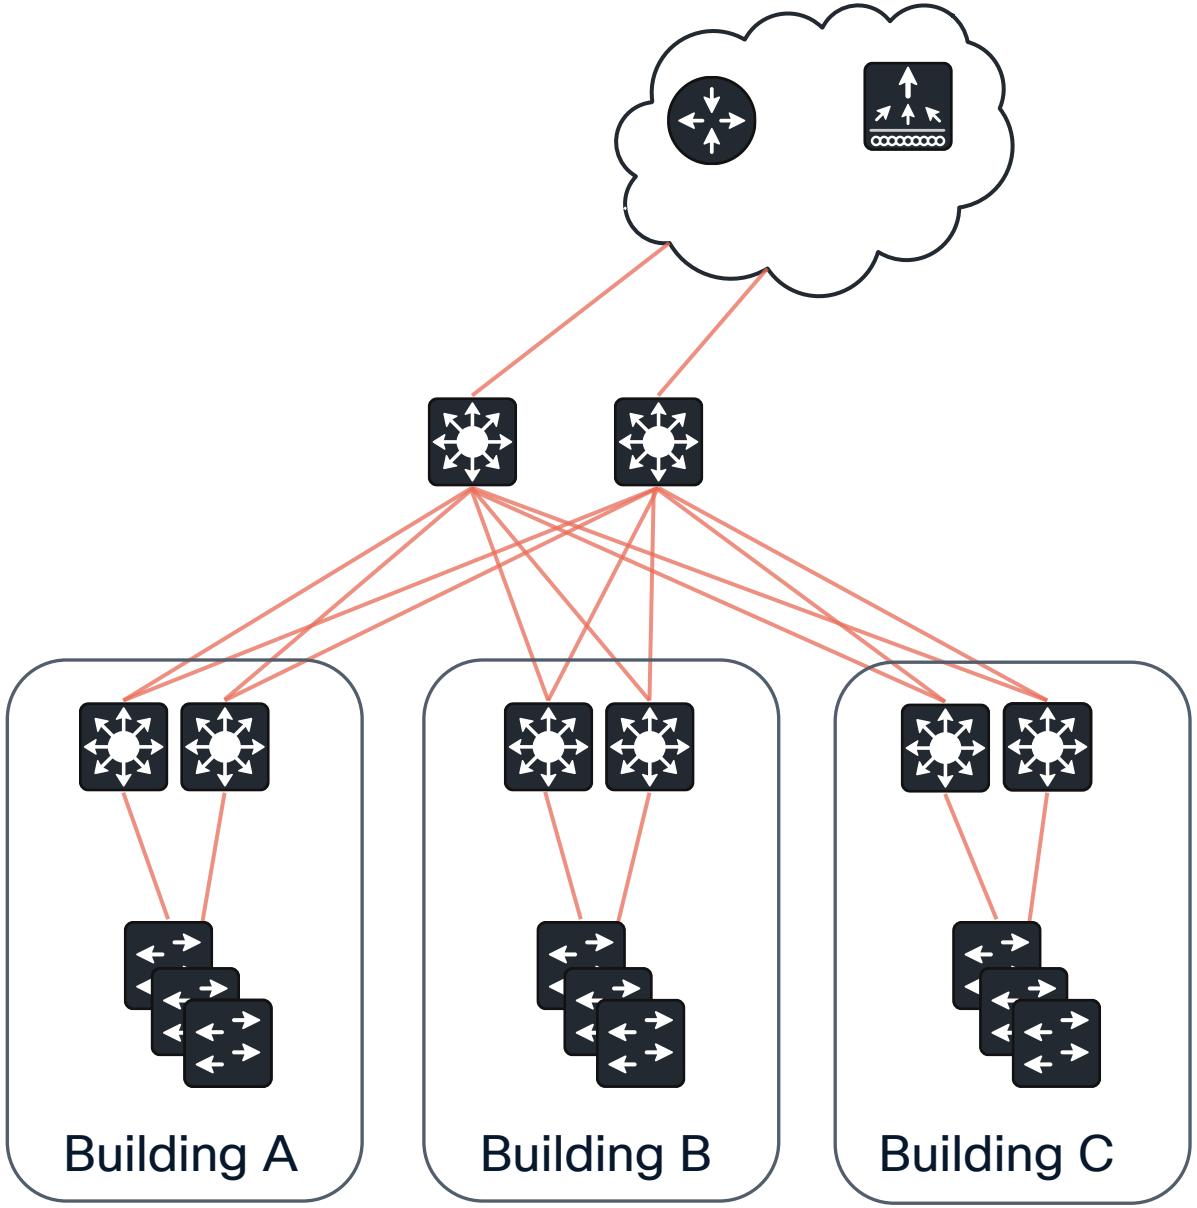

North-South optimised

Campus EVPN Deployment Options

IOS-XE can only support EVPN with VXLAN dataplane

Two main deployment options: fabric to access vs fabric to distribution and L2 access

Routed Access

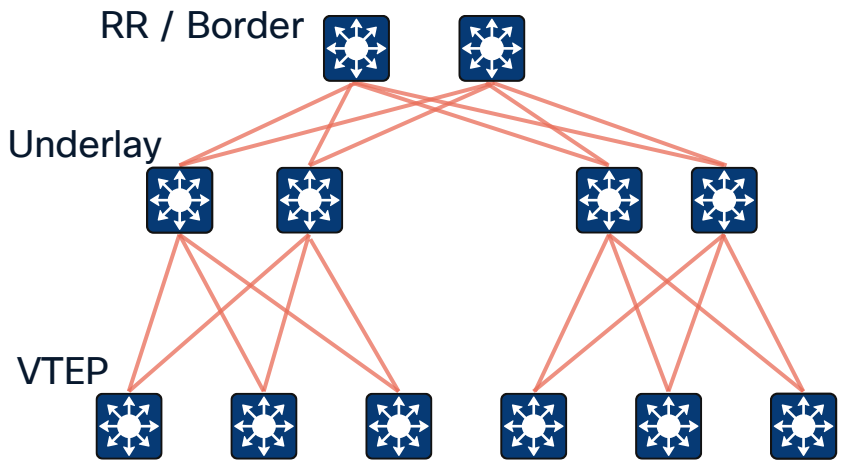

L2 Access with MLAG

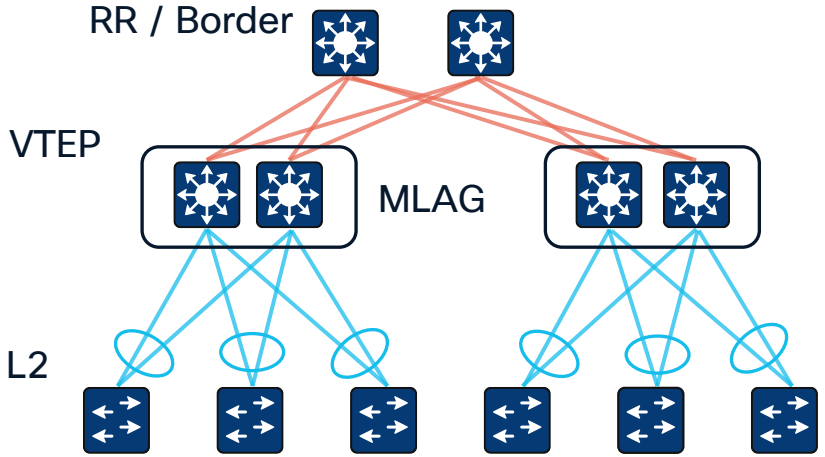

MLAG Options:

- StackWise Virtual
- ESI Single-Active
- ESI Multihoming (coming in Nov'25 /IOS-XE 17.18.2)

Fabric To Distribution and Layer 2 Access

Pros and Cons

L2 Access with MLAG

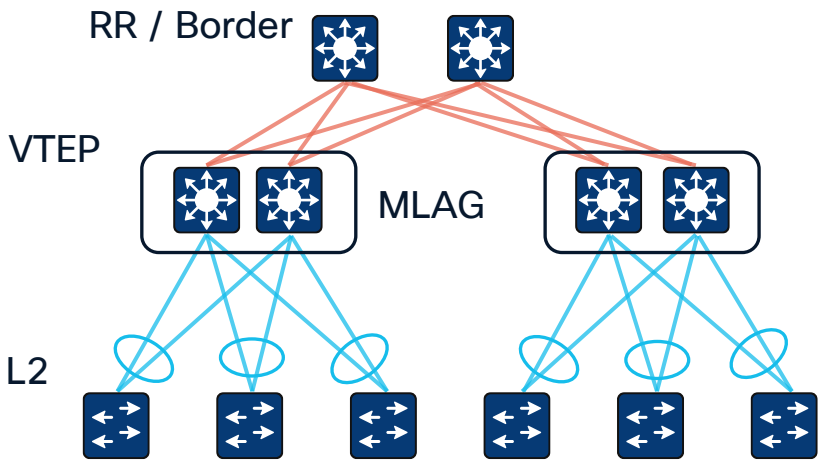

Pros:

- Familiar Layer 2 at access.
- Cheap – can use L2-only switches at access.
- Less EVPN-related TCAM scaling issues.

Cons:

- Familiar Layer 2 at access (STP + UDLD + IGMP + Trunking + VLAN Pruning + Port-Channels).
- StackWise Virtual = single control plane at distribution.
- ESI Active/Active Multihoming (17.18.2) is cutting edge capability & requires additional configuration.

Fabric to Routed Access

Pros and Cons

Routed Access

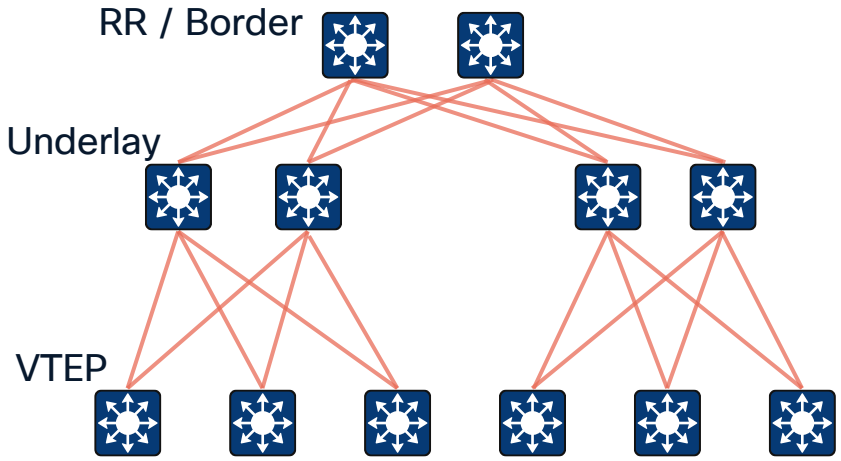

Pros:

- No Layer 2 protocols in the campus (same as in LISP-SDA).
- No SPOF at any point in the network and rapid network convergence.

Cons:

- Automation is highly desired.
- Extra design work is required to make sure your TCAM is used efficiently.

Fabric to Routed Access

Design Considerations

Routed Access

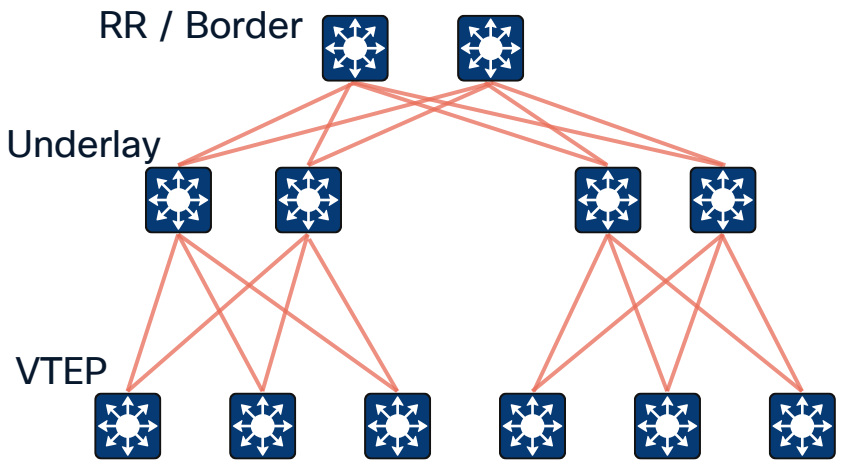

- **Access Layer TCAM resources:**
  - Every access switch will have a complete and consistent routing table across the fabric.
  - Fabric size is limited by TCAM resources in the access switch (lowest common denominator).
  - Maximum number of endpoints in the fabric will be limited by host route TCAM table.

- **Note: Each fabric endpoint consumes:**
  - 1 x TCAM slot per endpoint MAC address (EVPN Type2 MAC route)*
  - 1 x TCAM slot per endpoint IPv4 address (EVPN Type2 MAC/IP route)
  - 2 x TCAM slots per endpoint IPv6 address (EVPN Type2 MAC/IP route)

*without EVPN Proxy ARP

cisco

Fabric to Routed Access

Design Considerations

Routed Access

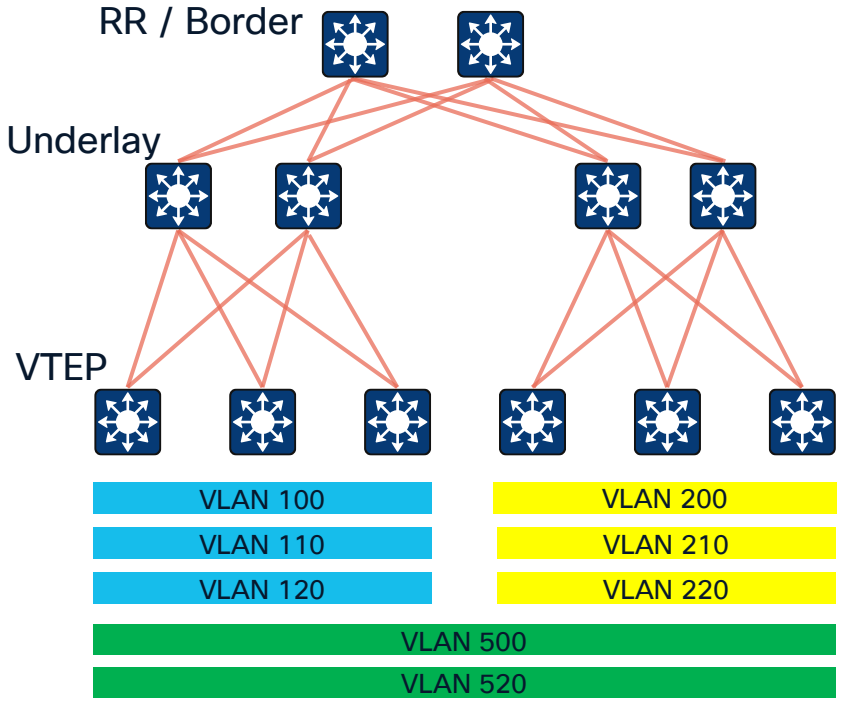

- Route everywhere
- Contain subnets to distribution blocks
  - Reduces TCAM scale concern
  - Wired clients typically do not need campus mobility
- Restrict the number of stretched Layer 2 segments with flooding
  - Use it case-by-case only
- Leverage EVPN Proxy ARP (create routed-only segments)

Implementation Details

Underlay

Fully routed, fabric-ready

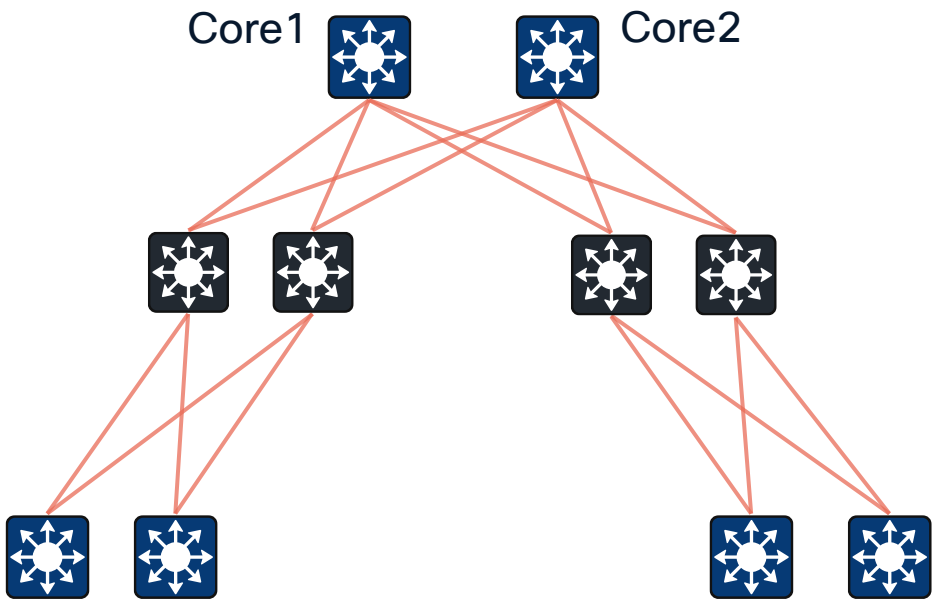

50 bytes overhead for VXLAN

```
system mtu 9100
```

Routing, management and VXLAN loopback

```
interface Loopback0
 ip address <IPv4 address> 255.255.255.255
 ip pim sparse-mode
 ip ospf network point-to-point
 ip ospf 1 area 0
```

Underlay routed interface

```
interface TenGigabitEthernet1/0/X
 ip address <IPv4 address> 255.255.255.252
 ip pim sparse-mode
 ip ospf network point-to-point
 ip ospf 1 area 0
```

OSPF optimisations

```
router ospf 1
 nsf ietf
 auto-cost reference-bandwidth 400000
 fast-reroute per-prefix enable prefix-priority low
 default-information originate
 max-metric router-lsa summary-lsa external-lsa include-stub on-startup 60
```

CEF optimisations

```
ip cef load-sharing algorithm include-ports source destination
ipv6 cef load-sharing algorithm include-ports source destination
port-channel load-balance src-dst-mixed-ip-port
```

Underlay

Multicast for underlay – ASM with Anycast RPs at Core1/Core2 with MSDP

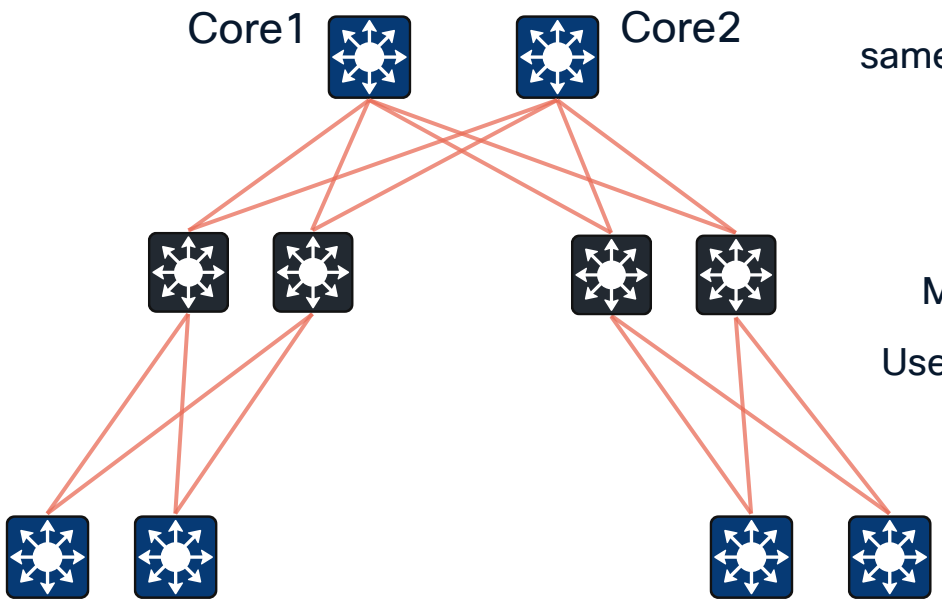

Anycast RP Loopback

same IP address on Core1/Core2

```
interface Loopback60000
 ip address <Anycast RP IPv4 address> 255.255.255.255
 ip pim sparse-mode
 ip ospf network point-to-point
 ip ospf 1 area 0
```

MSDP configuration for Core1

Use Core1 Loopback0 on Core2

```
ip msdp peer <Core2 Loopback0 IP address> connect-source Loopback0
ip msd originator-id Loopback0
```

Multicast configuration on all switches in the fabric

```
ip multicast-routing
ip pim rp-address <Loopback60000 IP address> override
ip pim register-source Loopback0
ip pim ssm default
```

CEF optimisation

```
ip multicast multipath s-g-hash next-hop-based
```

Infrastructure – BGP Configuration

Core1/Core2 are BGP router-reflectors for fabric VTEPs; no BGP configuration at distribution switches.

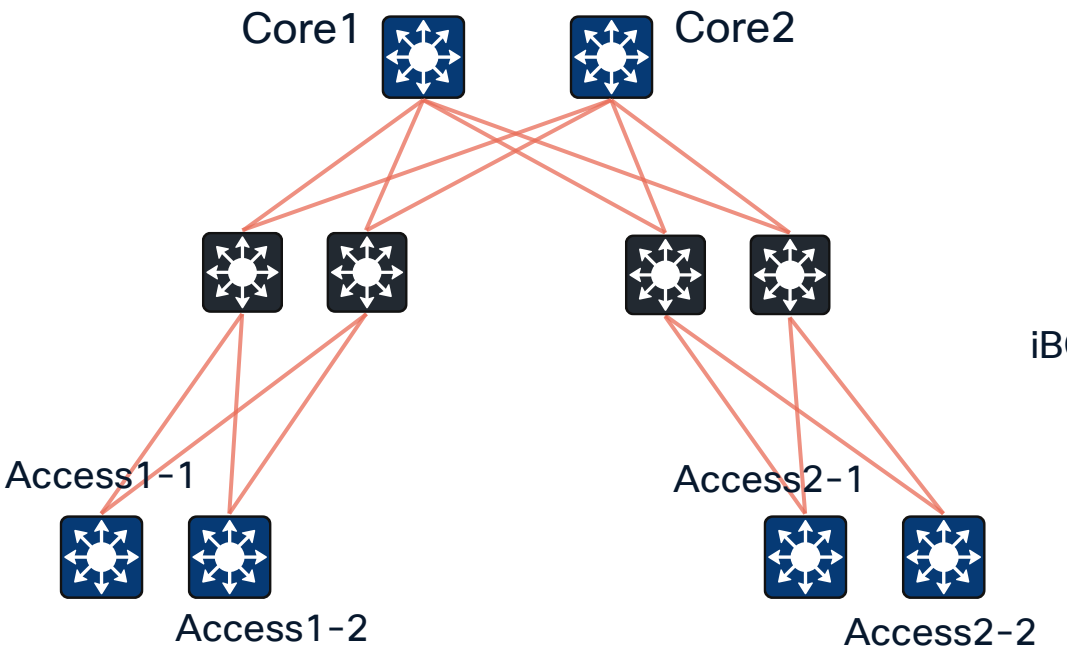

iBGP rapid failure detection

```
ip prefix-list VTEP-LOOPBACK seq 5 permit 10.67.118.32/27 ge 32
route-map VTEP-LOOPBACK permit 10
 match ip address VTEP-LOOPBACK-RANGE
```

```
router bgp 64513
 template peer-policy EVPN-PEER-POLICY
  send-community both
  route-reflector-client*
  advertisement-interval 0
 template peer-session EVPN-PEER-SESSION
  remote-as 64513
  update-source Loopback0
  fall-over route-map VTEP-LOOPBACK
 bgp router-id interface Loopback0
 bgp graceful-restart
 no bgp default ipv4-unicast
```

*On BGP RRs only

Infrastructure – BGP & NVE Interface Configuration

Core1/Core2 are BGP router-reflectors for fabric VTEPs; no BGP configuration at distribution switches.

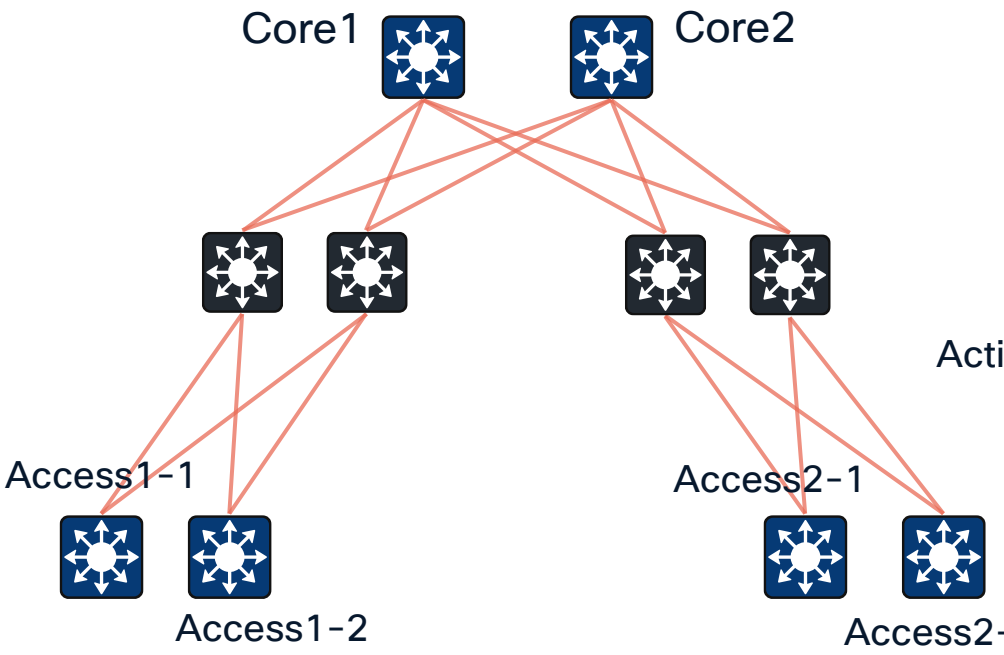

Define BGP neighbors*

```
router bgp 64513
 neighbor <Access1-1 Loopback0> inherit peer-session EVPN-PEER-SESSION
 neighbor <Access1-2 Loopback0> inherit peer-session EVPN-PEER-SESSION
 neighbor <Access1-1 Loopback0> inherit peer-session EVPN-PEER-SESSION
 neighbor <Access2-2 Loopback0> inherit peer-session EVPN-PEER-SESSION
```

Activate EVPN AF

```
address-family l2vpn evpn
  bgp nexthop trigger delay 0
  neighbor <Access1-1 Loopback0> activate
  neighbor <Access1-1 Loopback0> inherit peer-policy EVPN-PEER-POLICY
  neighbor <Access1-2 Loopback0> activate
  neighbor <Access1-2 Loopback0> inherit peer-policy EVPN-PEER-POLICY
```

Configure VXLAN interface

```
interface nve1
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
```

*Use of BGP Dynamic Neighbors is possible but makes it harder to troubleshoot BGP sessions.

Overlay – L2-only

At every access VTEP where this overlay is needed

- IP + non-IP traffic
- Gateway outside the fabric
- Limited built-in loop protection

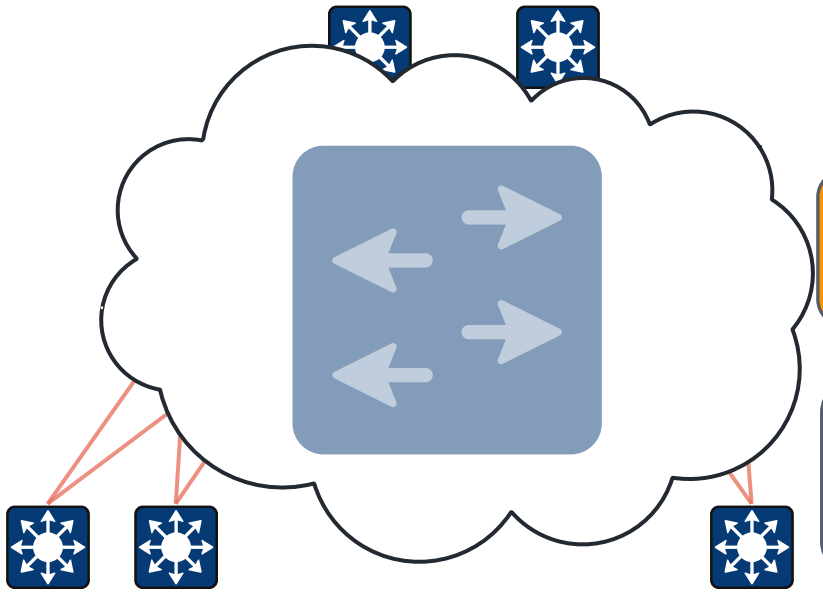

```
l2vpn evpn
 replication-type static
```

EVPN Global Settings:
Set BUM replication to multicast

```
l2vpn evpn instance <EVI ID> vlan-based
 encapsulation vxlan
```

Create MAC VRF / EVI

```
vlan <ACCESS VLAN ID>
 name <L2-OVERLAY-VLAN-NAME>
```

Create VLAN

```
vlan configuration <ACCESS VLAN ID>
 member evpn-instance <EVI ID> vni <L2-OVERLAY-VNI>
```

Map MAC VRF / VLAN / VNI to each other

```
interface nve1
 member vni <L2-OVERLAY-VNI> mcast-group <BUM-MCAST-GRP>
```

Add L2 VNI to VXLAN tunnel and map to BUM group if required

Overlay – L2 + L3, Part 1

Anycast Gateway - Bridged

Create L2 overlay – same as previous slide.

- IP + non-IP traffic
- Anycast Gateway at every access VTEP
- ARP Suppression is enabled (default)
- EVPN Proxy ARP is not used (default)

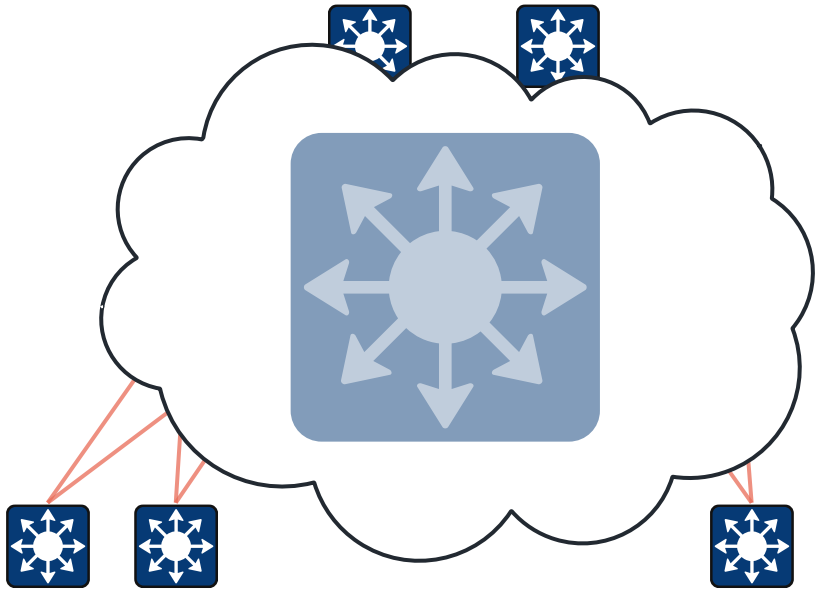

Add L3 part

```
vrf definition DAG-VRF-NAME
 rd <VTEP-Loopback0:L3VNI>
 route-target import/export <VTEP-BGP-ASN-ID:L3VNI>
 route-target import/export <VTEP-BGP-ASN-ID:L3VNI> stitching
```

Create IP VRF

```
vlan <DAG VRF CORE VLAN ID>
 name <DAG-VRF-CORE-VLAN-NAME>
```

Create Core VLAN

```
interface vlan <DAG VRF CORE VLAN ID>
 vrf forwarding DAG-VRF-NAME
 ipv6 enable
 ip pim sparse-mode
 ip unnumbered Loopback0
 no autostate
```

Create Core SVI in target VRF

```
vlan configuration <DAG VRF CORE VLAN ID>
 member vni <DAG VRF L3 VNI>
```

Map Core VLAN to L3 VNI

Overlay – L2 + L3, Part 2

Anycast Gateway - Bridged

- IP + non-IP traffic
- Anycast Gateway at every access VTEP
- ARP Suppression is enabled (default)
- EVPN Proxy ARP is not used (default)

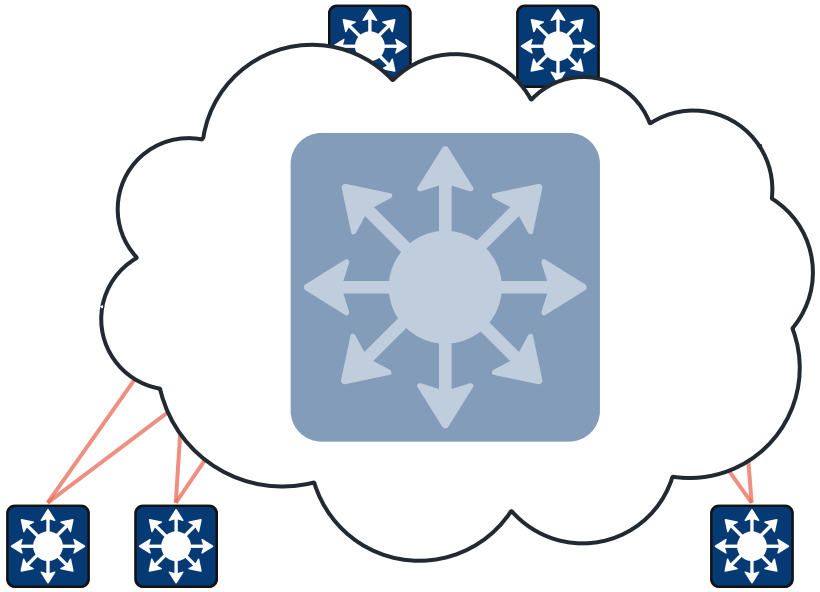

```
l2vpn evpn instance <EVI ID> vlan-based
 default-gateway advertise enabled
```

Enable Anycast Gateway MAC/IP route advertisement

```
interface vlan <ACCESS VLAN ID>
 mac-address 0000.0000.<L2VNI>
 vrf forwarding <DAG-VRF-NAME>
 ip address <ACCESS SUBNET GW IP Address>
```

Create Anycast Gateway

```
interface nve1
 member vni <DAG VRF L3 VNI> vrf <DAG-VRF-NAME>
```

Add L3 VNI to VXLAN interface

```
router bgp 64513
 address-family ipv4 vrf <DAG-VRF-NAME>
  advertise l2vpn evpn
  redistribute connected
  maximum-paths 2
```

Advertise DAG VRF in EVPN

Overlay – L3 Only

Anycast Gateway - Routed

Create L2+L3 overlay as in the previous slide

Make it routed only

- IP traffic only
- Anycast Gateway at every access VTEP
- ARP Suppression is disabled
- EVPN Proxy ARP is enabled

```
l2vpn evpn instance <EVI ID> vlan-based
 flooding-suppression address-resolution disable
```

Disable ARP suppression

```
interface vlan <ACCESS VLAN ID>
 mac-address 0000.0000.<L2VNI>
 vrf forwarding <DAG-VRF-NAME>
 ip address <ACCESS SUBNET GW IP Address>
 ip proxy-arp
 ip local-proxy-arp
```

Enable EVPN Proxy ARP

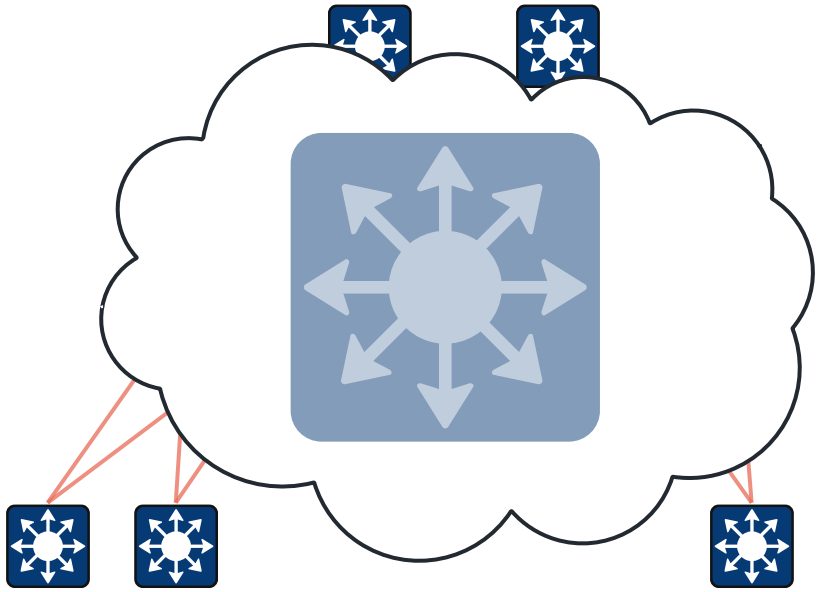

Benefit: EVPN Type 2 MAC routes are not generated (50% less TCAM utilisation).

Caveat: No flooding in the overlay (unknown unicast, link-local multicast, etc).

Overlay – Centralised Gateway

- IP traffic only
- L2 VTEP at access
- L3 VTEP at service attachment
- Bring traffic to L3 handoff (firewall)

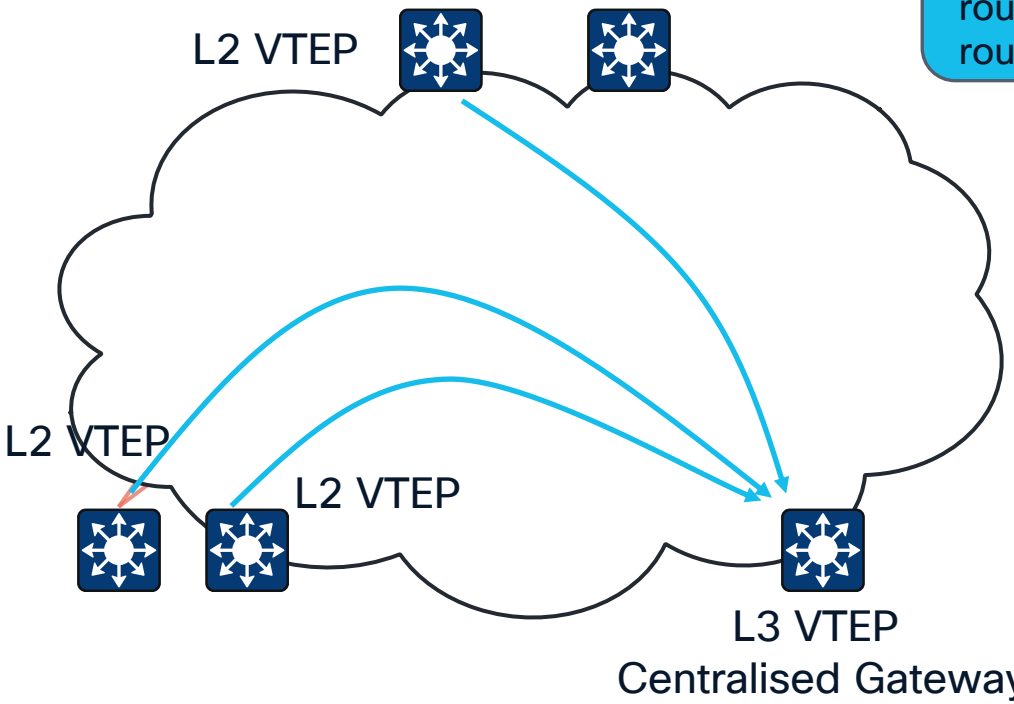

Create L2 overlay at access VTEPs

Create CGW at service attachment VTEP

```
vrf definition CGW-VRF-NAME
 rd <VTEP-Loopback0:L3VNI>
 route-target import/export <VTEP-BGP-ASN-ID:L3VNI>
 route-target import/export <VTEP-BGP-ASN-ID:L3VNI> stitching
```

Create IP VRF

```
vlan <CGW VRF CORE VLAN ID>
 name <CCW-VRF-CORE-VLAN-NAME>
```

Create Core VLAN

```
interface vlan <CGW VRF CORE VLAN ID>
 vrf forwarding CGW-VRF-NAME
 ip pim sparse-mode
 ip unnumbered Loopback0
 no autostate
```

Create Core SVI in target VRF

```
vlan configuration <CGWVRF CORE VLAN ID>
member vni <CGW VRF L3 VNI>
```

Map Core VLAN to L3 VNI

Overlay – Centralised Gateway, Part 2

- IP traffic only
- L2 VTEP at access
- L3 VTEP at service attachment
- Bring traffic to L3 handoff (firewall)

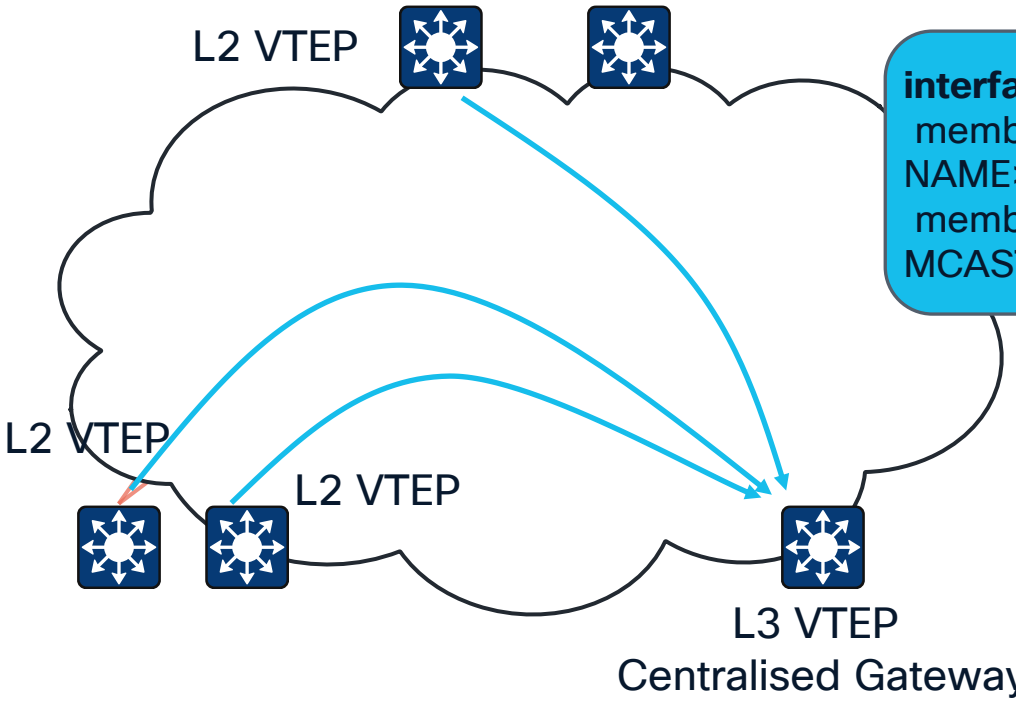

Create CGW at service attachment VTEP

```
l2vpn evpn instance <EVI ID> vlan-based
 default-gateway advertise enabled
```

Enable CGW MAC/IP route advertisement

```
interface vlan <CCW VLAN ID>
 mac-address 0000.0000.<L2VNI>
 vrf forwarding <CCW-VRF-NAME>
 ip address <ACCESS SUBNET GW IP Address>
```

Create Centralised Gateway

```
interface nve1
 member vni <CCW VRF L3 VNI> vrf <CCW-VRF-NAME>
 member vni <CCW VRF L3 VNI> mcast-group <BUM-MCAST-GRP> local-routing
```

Add L3 VNI to VXLAN interface

```
router bgp 64513
 address-family ipv4 vrf <DAG-VRF-NAME>
  advertise l2vpn evpn
  redistribute connected
  maximum-paths 2
```

Advertise CCW VRF in EVPN

ESI Multihoming

Infrastructure

- iBGP peering over L3 port-channel between VTEPs
- Supports routed Anycast Gateways
- Supports bridged Anycast Gateways

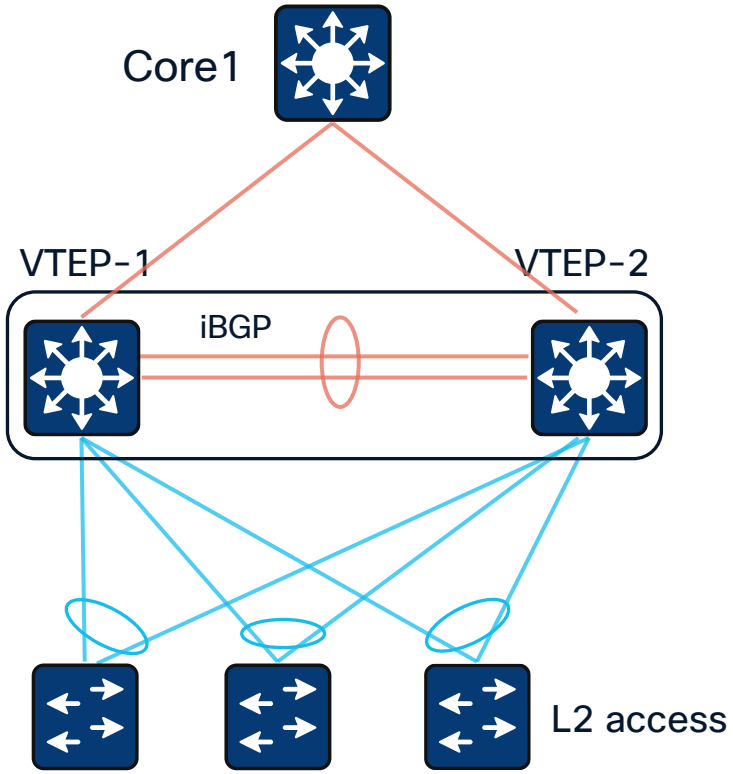

VTEP-1 configuration, replicate for VTEP-2

```
router bgp 64513
 neighbor <VTEP-2 Loopback0> inherit peer-session EVPN-PEER-SESSION
 address-family l2vpn evpn
  neighbor <VTEP-2 Loopback0> activate
  neighbor <VTEP-2 Loopback0> inherit peer-policy EVPN-PEER-POLICY
```

```
l2vpn evpn
 multihoming aliasing disable
 multicast advertise sync-only
```

```
interface Port-channel <ID>
 description CONNECTED TO L2 ACCESS
 evpn multihoming core-tracking
 evpn ethernet-segment auto lacp df-election wait-time 1
 switchport mode trunk
```

```
interface te1/1/1
 description CONNECTED TO CORE
 evpn multihoming core-tracking
```


cisco

Overlay – DHCP

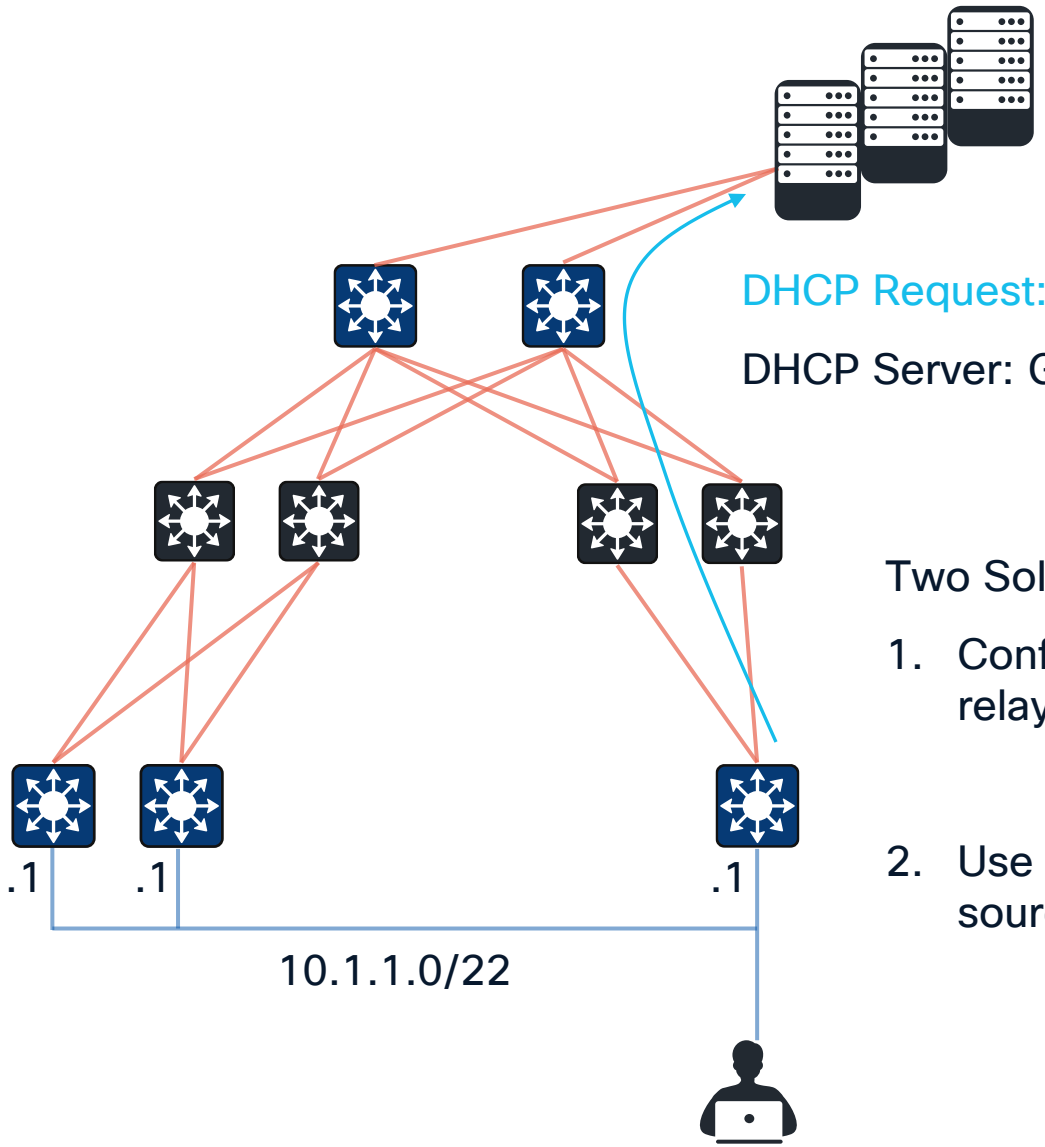

DHCP Request: GiAddr = 10.1.1.1

DHCP Server: GiAddr = 10.1.1.1, where to send the response?

Two Solutions:

1. Configure additional, per-VTEP unique Loopback interface as DHCP-relay source. Repeat for each VRF where you need DHCP service.

2. Use per-VTEP unique Loopback interface (Loopback0) as DHCP-relay source. No extra per-VRF configuration – recommended.

Overlay – DHCP – Option 2

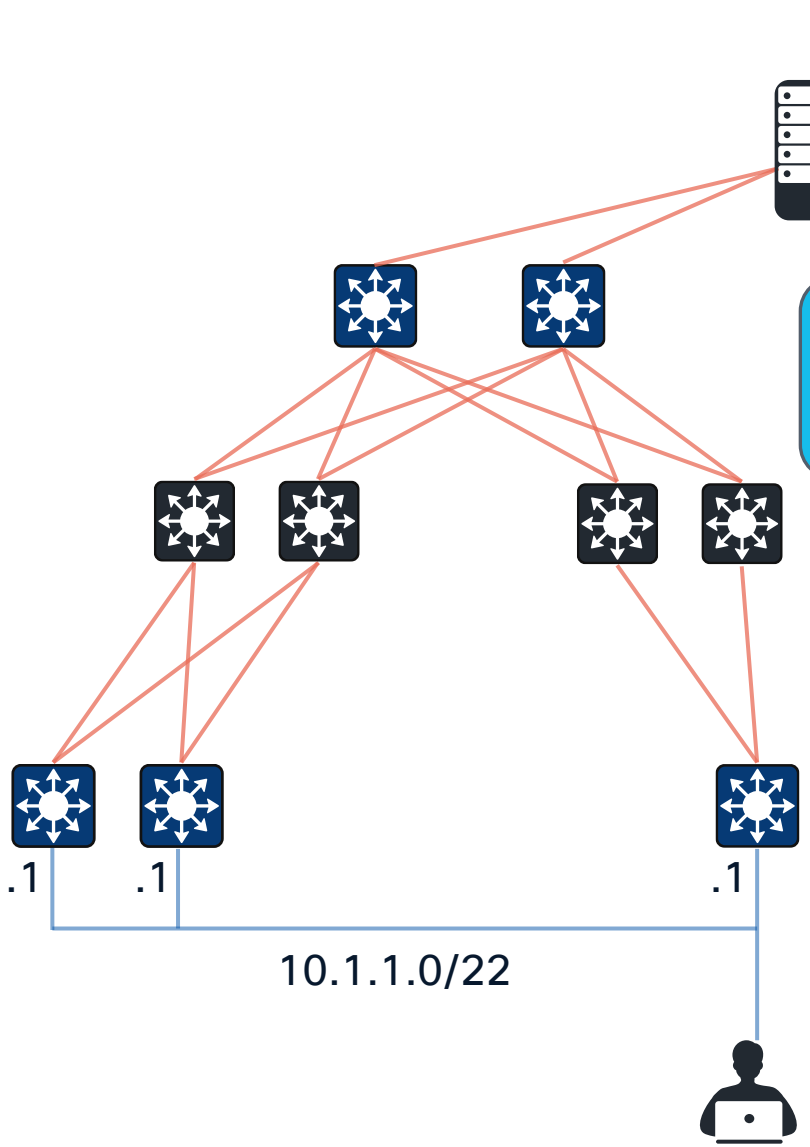


```
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
```

Insert Option 82, VRF ID and set instruction for DHCP server

```
ip dhcp-relay source-interface Loopback0
```

Set DHCP relay as Loopback0

```
ip dhcp snooping
ip dhcp snooping vlan <ANYCAST GW VLAN ID>
```

Enable DHCP snooping at access VLAN (required)

```
interface vlan <ANYCAST GW VLAN ID>
 ip helper-address global <DHCP SERVER1 IP ADDR>
 ip helper-address global <DHCP SERVER2 IP ADDR>
```

Set DHCP relay at Anycast GW

Overlay – Group Policy / TrustSec Implementation

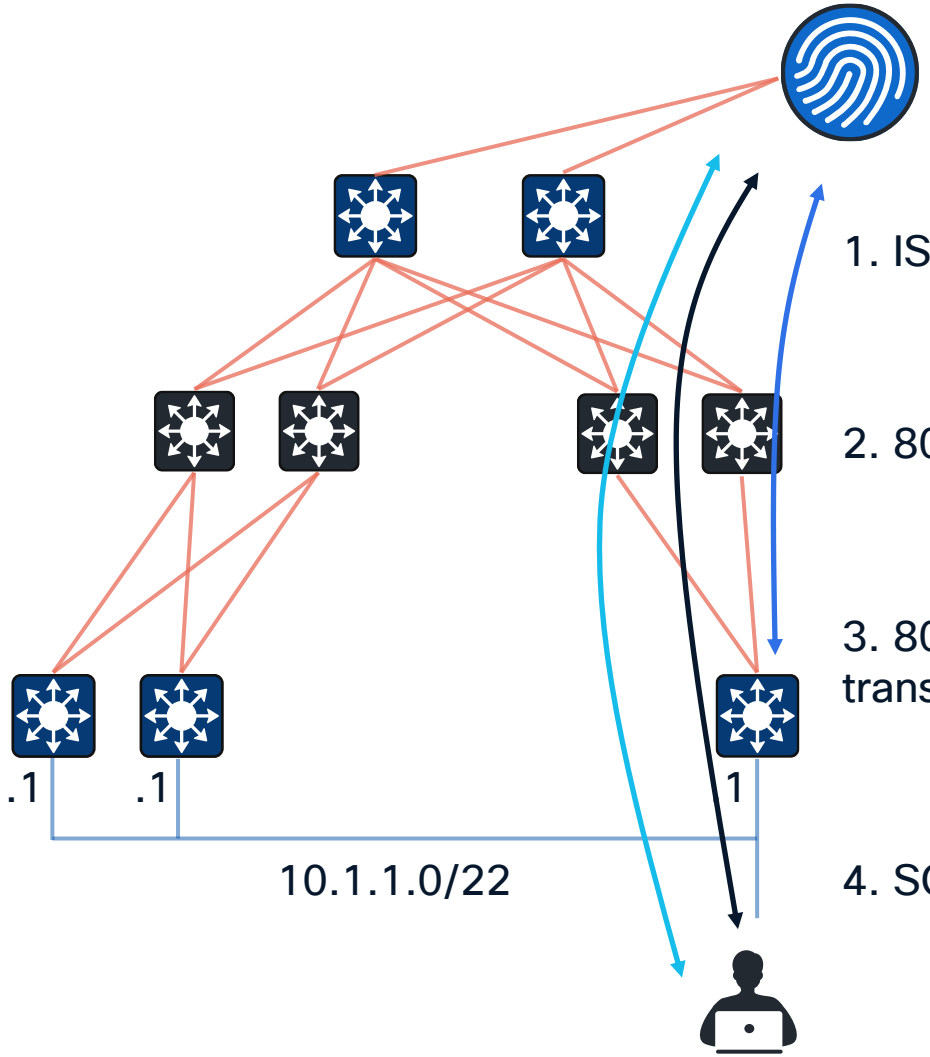

1. ISE <-> NAD RADIUS & TrustSec connection

2. 802.1X or MAB – Authentication

3. 802.1X or MAB – Authorization (VLAN and SGT assignment). Enable SGT transport in VXLAN header.

4. SGACL download and enforcement

Overlay – Group Policy Implementation

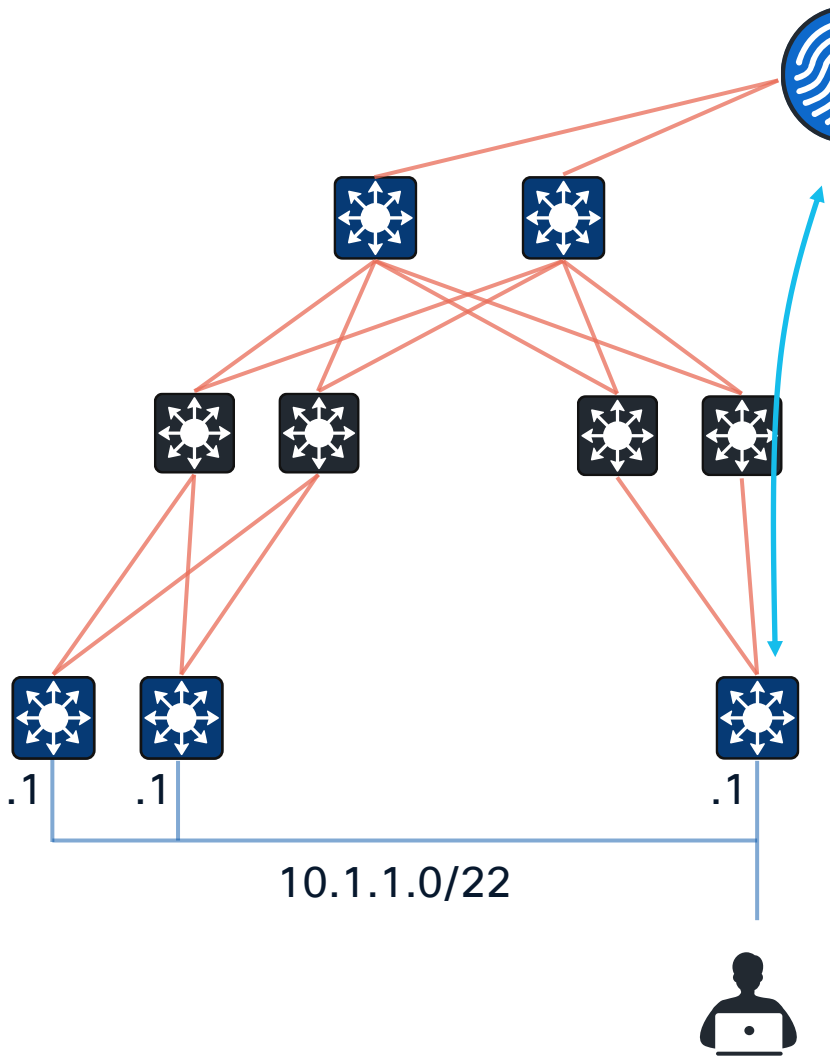

1. ISE <-> NAD RADIUS & TrustSec connection

```
aaa group server radius TRUSTSEC
 server name ISE
aaa server radius dynamic-author
 client <ISE PSN IP Address> server-key <RADIUS PSK>
ip radius source-interface Loopback0
radius server ISE
 address ipv4 <ISE PSN IP Address> auth-port 1812 acct-port 1813
aaa authentication dot1x default group TRUSTSEC
aaa authorization network default group TRUSTSEC
aaa authorization network cts-mlist group TRUSTSEC
cts authorization list cts-mlist
```

RADIUS configuration

PAC-less authentication*

```
#cts credentials id <CTS-ID> password <CTS-PASSWORD>
```

Enable mode

*Needs ISE 3.4+ & IOS-XE 17.15.1+

Overlay – Group Policy Implementation

2. 802.1X or MAB - Authentication

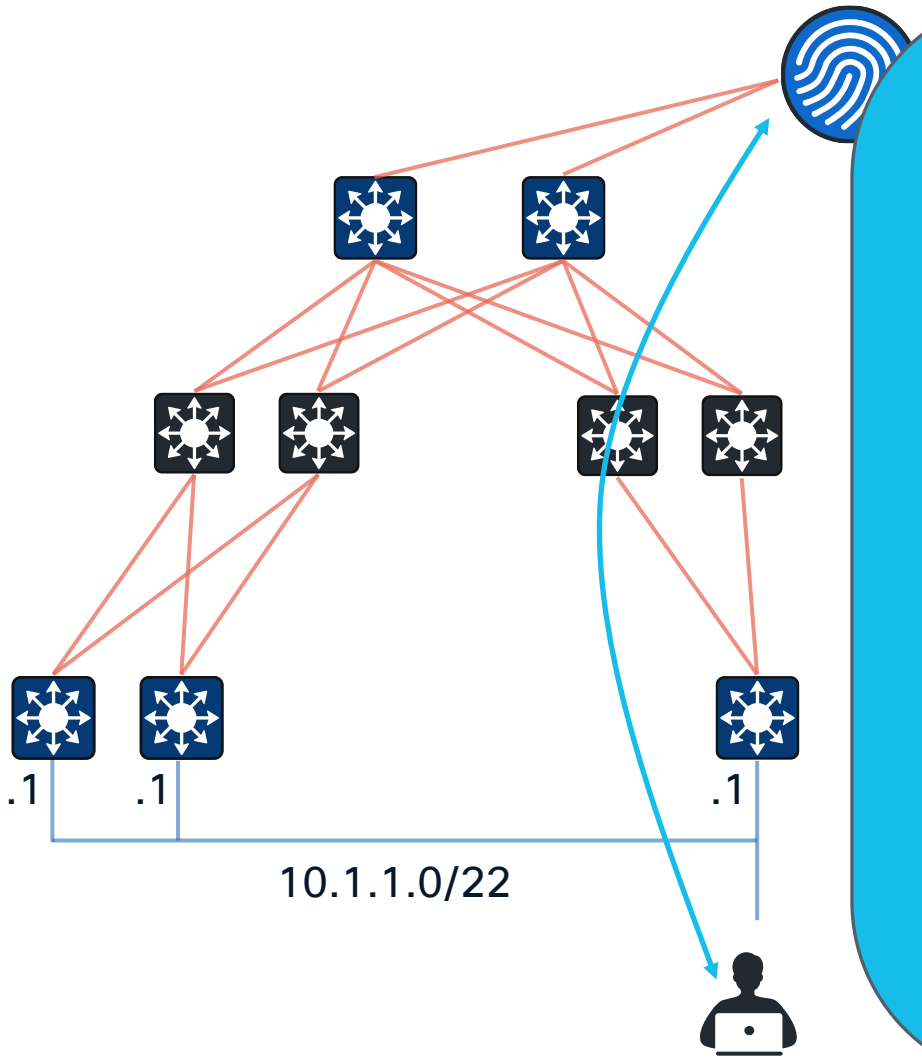


```
interface GigabitEthernet1/0/1
 source template Dot1xMAB
template Dot1xMAB
 dot1x pae authenticator
 dot1x timeout tx-period 3
 dot1x timeout supp-timeout 2
 dot1x max-req 3
 dot1x max-reauth-req 3
 switchport mode access
 mab
 access-session closed
 access-session port-control auto
 authentication periodic
 authentication timer reauthenticate server
 service-policy type control subscriber IBNS2-EVPN*
```


*service-policy example is in hidden slide

Overlay – Group Policy Implementation

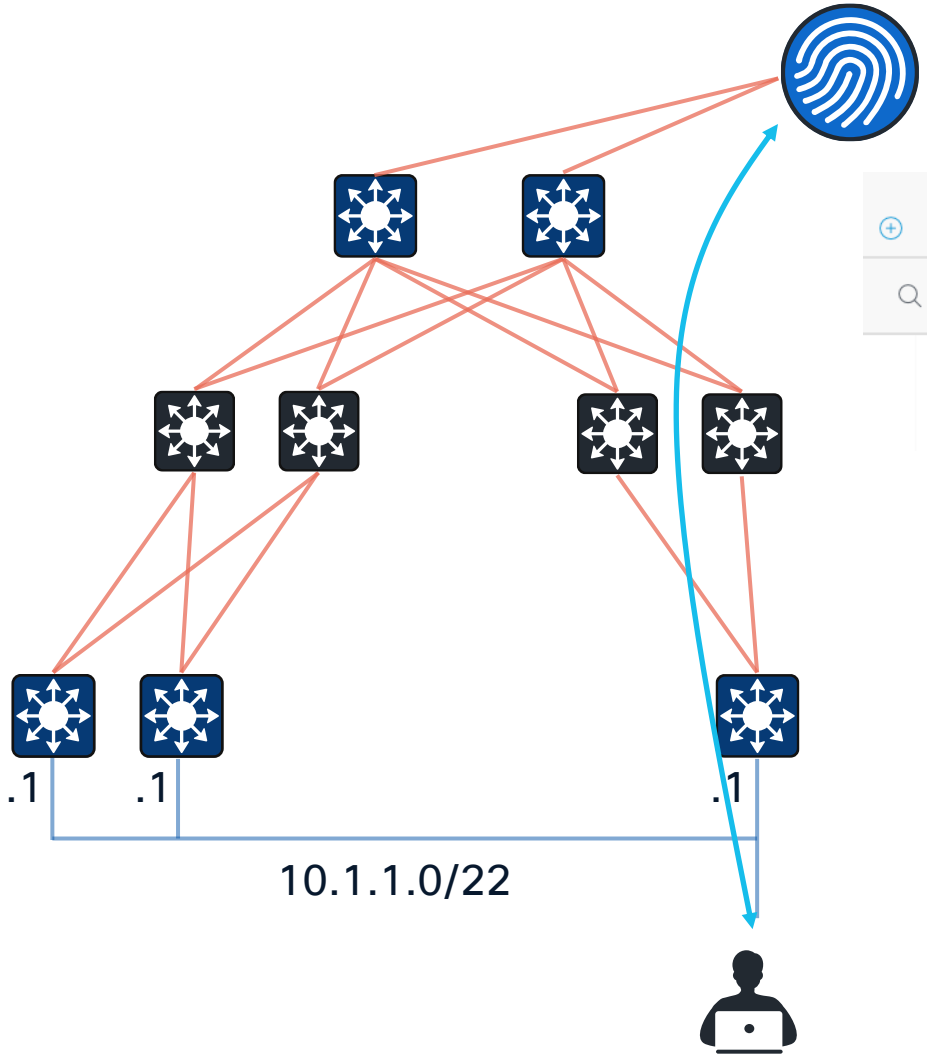

3. 802.1X or MAB – Authorization (VLAN and/or SGT assignment)

Status	Rule Name	Conditions	Profiles	Security Groups
Search				
Staff	AND	Wired_802.1X evpn-ExternalGroups EQUALS evpn.local/SecurityGroups/Staff	Campus-STAFF	Staff

Enable SGT transport in VXLAN header (GPO).

```
interface nve1
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
 group-based-policy
```

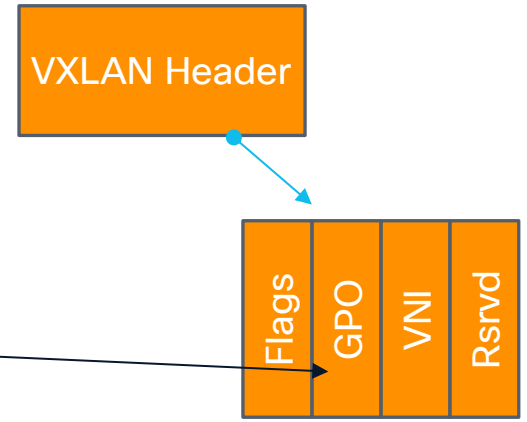

Overlay – Group Policy Implementation

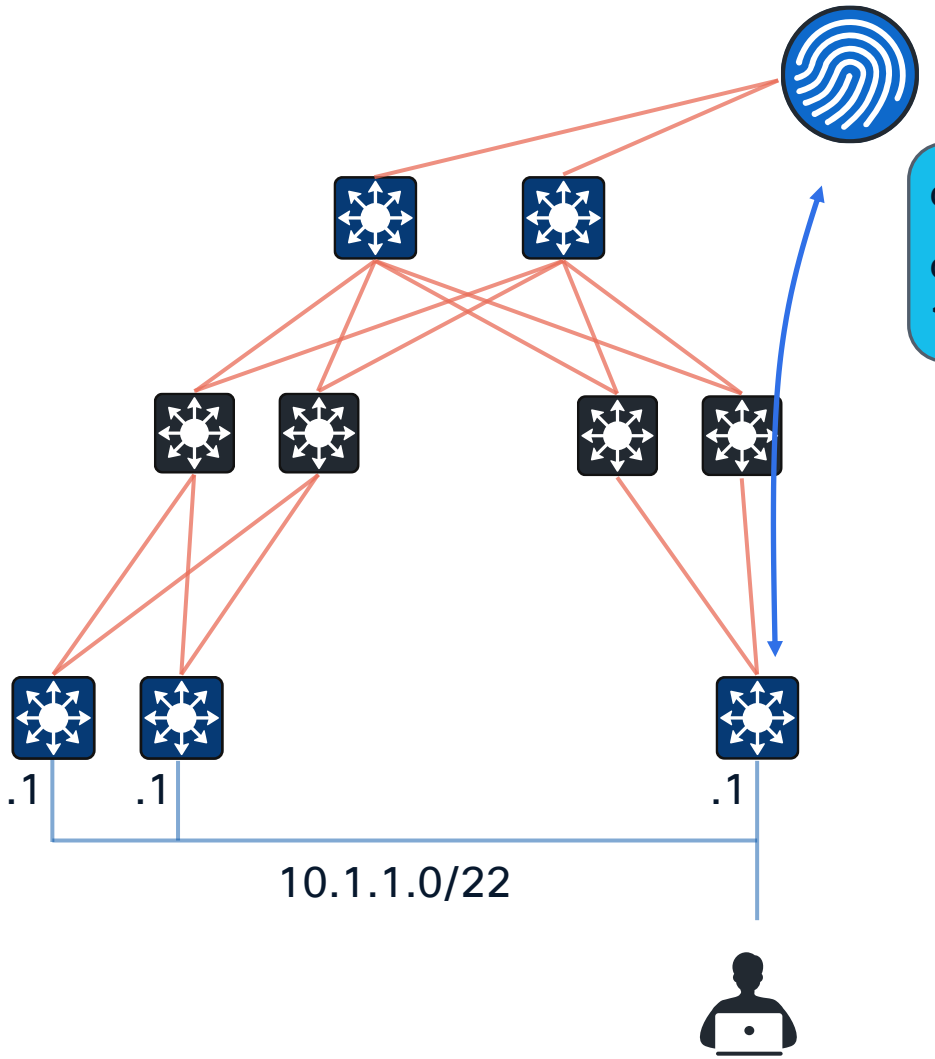

4. SGACL download and enforcement

```
cts role-based enforcement

cts role-based enforcement vlan-list <ACCESS VLAN ID-1>,
<ACCESS VLAN ID-N>
```

Enable enforcement on access VLANs

```
interface TenGigabitEthernet1/1/1
 no cts role-based enforcement
```

Disable enforcement on routed uplinks

Overlay – Loop Prevention

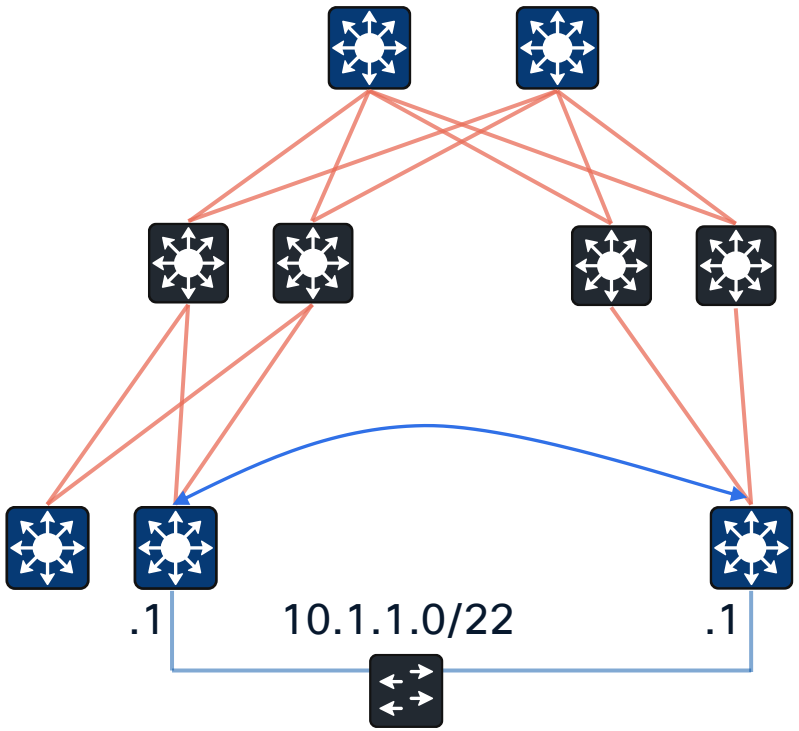

5 roams within 3 minutes

```
Access-1-1#show l2vpn evpn summary
L2VPN EVPN
  EVPN Instances (excluding point-to-point): 3
    VLAN Based:   3
  Vlans: 3
  BGP: ASN 64513, address-family l2vpn evpn configured
  Router ID: 10.67.118.38
  Global Replication Type: Static
  ARP/ND Flooding Suppression: Enabled
  Connectivity to Core: UP
  MAC Duplication: seconds 180 limit 5
  MAC Addresses: 1
    Local:     1
    Remote:    0
    Duplicate: 0
  IP Duplication: seconds 180 limit 5
  IP Addresses: 2
    Local:     2
    Remote:    0
    Duplicate: 0
```

If the limit is reached, MAC/IP is blacklisted in the fabric and will require manual re-enablement.

Multicast in EVPN

Multicast For Overlays

Default MDT and Data MDT

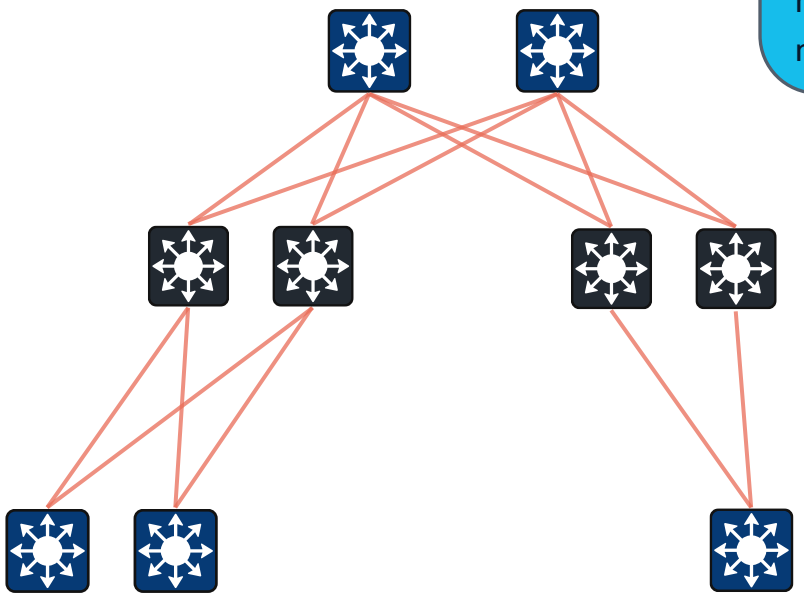

```
router bgp 64513
 address-family ipv4 mvpn
 neighbor <Access1-1 Loopback0> activate
 neighbor <Access1-1 Loopback0> inherit peer-session EVPN-PEER-POLICY
 neighbor <Access1-2 Loopback0> activate
 neighbor <Access1-2 Loopback0> inherit peer-session EVPN-PEER-POLICY
```

Add MVPN AF

```
vrf definition <DAG-VRF-NAME>
 address-family ipv4
  mdt auto-discovery vxlan
  mdt default vxlan <DEFAULT MDT GROUP-1>
  mdt data vxlan <DATA MDT GROUP RANGE> <WILDCARD MASK>
  mdt overlay use-bgp spt-only
```

MDT VRF Configuration

```
ip multicast-routing vrf <DAG-VRF-NAME>
ip pim vrf <DAG-VRF-NAME> ssm default
interface vlan <DAG VRF CORE VLAN ID>
 ip pim sparse-mode

interface vlan <ANYCAST GW VLAN ID>
 ip pim sparse-mode
 ip igmp version 3
```

Multicast and PIM Configuration

Multicast Inside the Fabric VRF

Fabric Anycast RP

- Each VTEP hosts Anycast RP for the VRF
- Underlay VTEP interface (L0) is used as PIM source
- No inter-VRF multicast
- Can interoperate with external RPs using MSDP

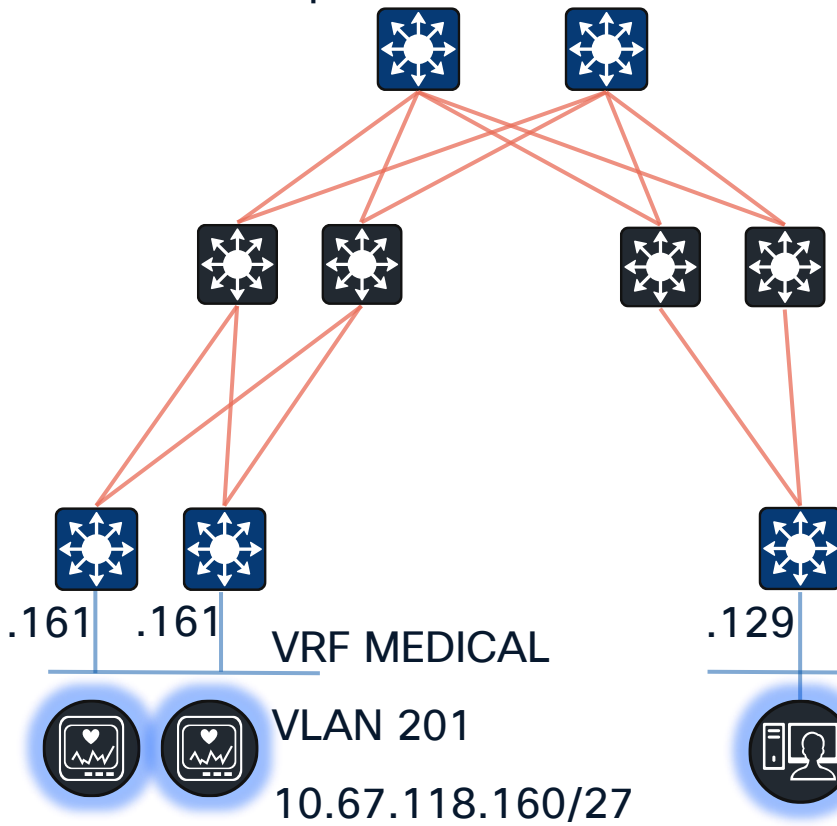

Anycast RP interface

```
interface Loopback<L3VNI>
 vrf forwarding <DAG-VRF-NAME>
 ip address <Anycast-IPv4-Address> 255.255.255.255
 ip pim sparse-mode
```

RP configuration

```
ip pim vrf <DAG-VRF-NAME> rp-address <Anycast-IPv4-Address>
```

Multicast Between the Fabric VRFs

External RP

- Fabric VRF is using external RP
- Each VTEP hosts in-VRF unique Loopback interface as PIM Source
- Allows inter-VRF multicast (same external RP for different VRFs)

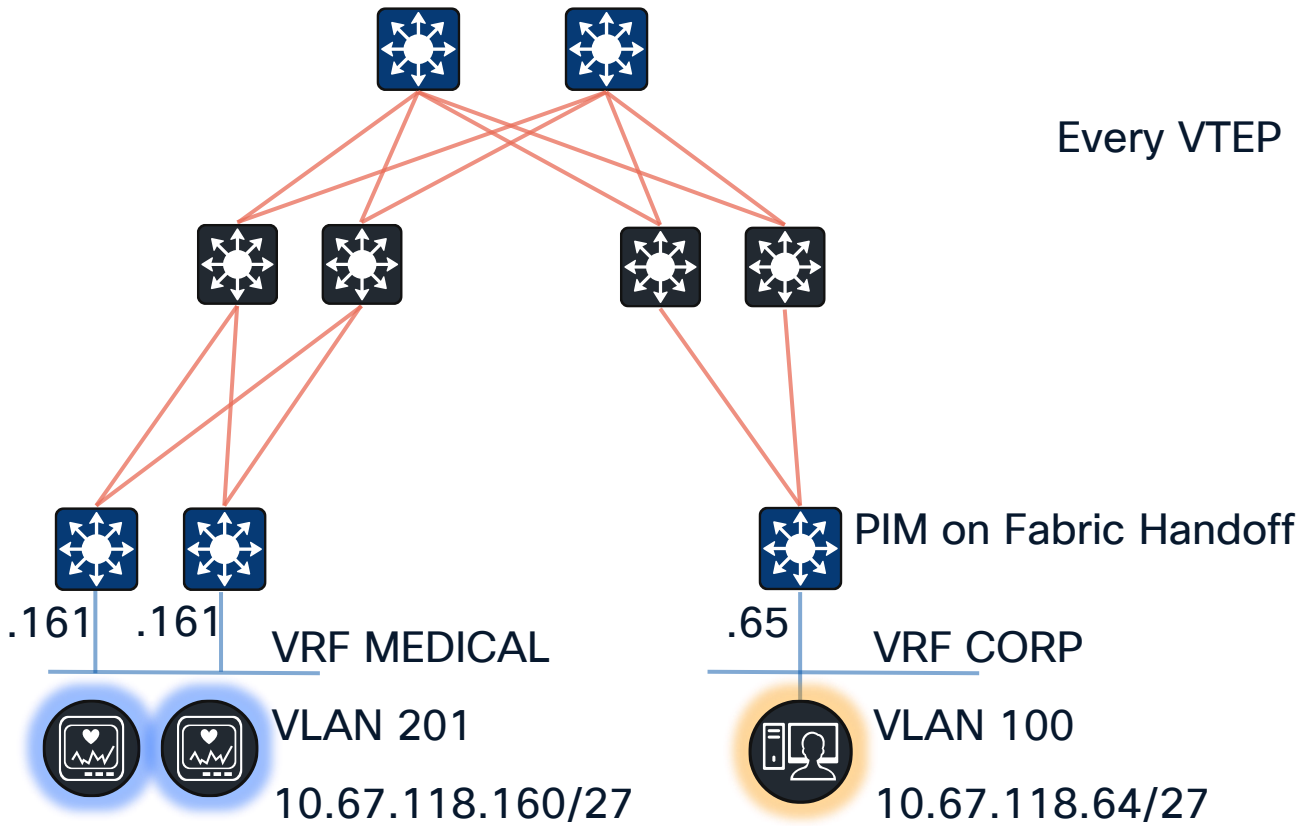

Every VTEP

```
ip pim vrf <DAG-VRF-NAME> register-source Loopback<L3VNI>
ip pim vrf <DAG-VRF-NAME> rp-address <EXTERNAL-IP-ADDRESS>
interface Loopback<L3VNI>
 vrf forwarding <DAG-VRF-NAME>
 ip address <Unique-IPv4-Address> 255.255.255.255
 ip pim sparse-mode
```

```
interface GigabitEthernet1/0/1
 vrf forwarding <DAG-VRF-NAME>
 ip pim sparse-mode
```

Closing

Key Takeaways

- BGP EVPN address-family provides flexible overlays (L2, L2+L3, L3 only, Distributed, Centralised) that are suited to any business need.
- EVPN overlays in the campus can be either:
  - To distribution, keeping L2 access in place
  - To access layer, removing L2 protocols from the network
- Both options come with a set of trade-offs; There is no right or wrong answer.
- Functionally, the C9K EVPN fabric meets all requirements of a modern campus:
  - Overlays
  - DHCP
  - Multicast
  - Loop prevention
  - Multihoming / MLAG
  - Micro-segmentation (SGTs)

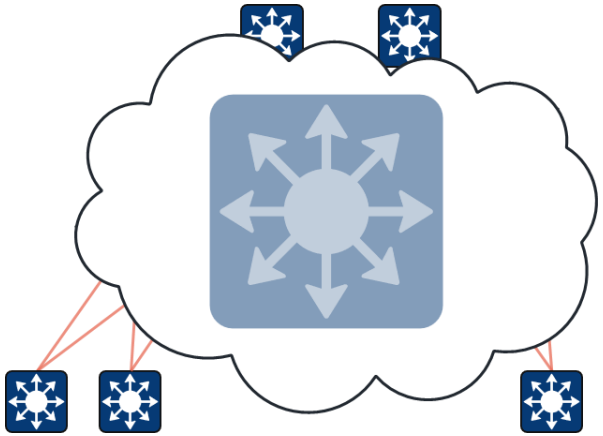

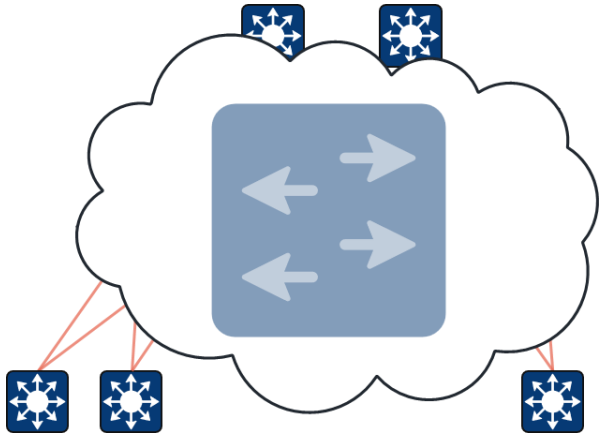

cisco

Complete your session evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to claim a Cisco Live T-Shirt.

Earn up to 800 points by completing all surveys and climb the Cisco Live Challenge leaderboard.

Level up and earn exclusive prizes!

Complete your surveys in the Cisco Live Events app.

Continue your education

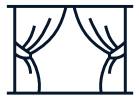

Visit the Cisco Stand for related demos

Book your one-on-one Meet the Expert meeting

Attend the interactive education with Capture the Flag, and Walk-in Labs

Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

Thank you