



Building Scalable VXLAN BGP EVPN Fabrics for Enterprise networks

Catalyst 9000 Switching Platforms

Raj Kumar Goli - Technical Marketing Engineer

BRKENS-2092



Agenda

- What is BGP EVPN ?
- BGP EVPN in Enterprise Campus
- Underlay and Overlay Networks
- Scaling Multicast in Fabric
- BGP EVPN Interworking
- SD-Access with BGP EVPN

Webex App

Questions?

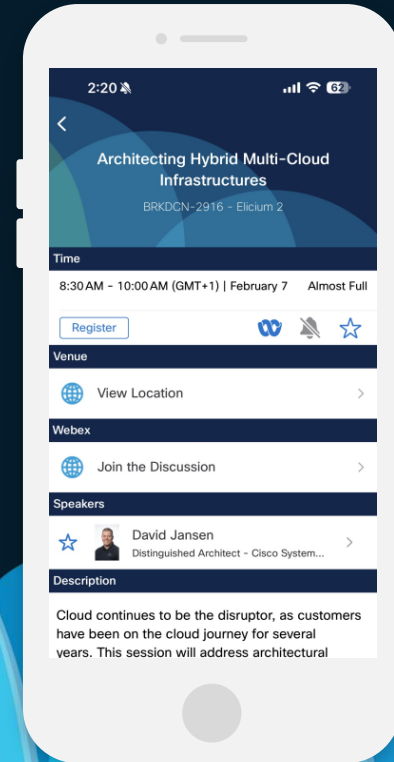
Use the Webex app to chat with the speaker after the session

How

- 1 Find this session in the Cisco Events 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 February 28, 2025.

CISCO *Live!*



Catalyst 9000 Leaf | Border Support
 Overlay Type : Layer 3 Overlay
 Overlay Type : Distributed AnyCast Gateway
 IPv4 host in overlay
 Ingress Replication for BUM
 DHCPv4 Relay in EVPN VRF
 Border : Multi-VRF IPv4 Handoff
 Border : L2 VLAN Handoff

Catalyst 9000 Spine | Leaf | Border Support
 Overlay Type : Layer 2, Centralized Gateway IPv4 ARP / IPv6 ND Suppression for Distributed AnyCast Gateway, Layer 2 Leaf
 EVPN L2 Multi-homing with Cisco StackWise Virtual
 Central Switching Wireless support
 Firewall integration for Service-Insertion
 IPv6 host in overlay
 DHCPv4/v6 Relay in Default VRF
 VLAN Aware Flexible Netflow - IPv4 / IPv6 Unicast/Multicast overlay
 Border : Multi-VRF IPv6 Handoff
 Border : EVPN to MPLS VPNv4 Integration
 Border : EVPN to MPLS VPNv6 Integration
 Border : EVPN to VPLS Bridge Interworking

Hybrid Role - Spine + Leaf + Border Support
 BGP EVPN Layer 2 Multi-home ESI support
 Distributed AnyCast RP for TRM
 IPv4/IPv6 host overlay Multicast with Default MDT, Fabric External Domain RP
 L3/L3 Overlay Topologies : Full-mesh | Partial-Mesh | Hub-n-Spoke | P2P
 Per-VNI Multicast BUM Rate-Limiter
 BGP EVPN PVLAN based Segmentation
 Multicast DNS Service-Routing over BGP EVPN VXLAN
 Border Multi-VRF v4/v6 Multicast Handoff
 Border : EVPN to VPLS IRB Interworking domain
 Border : EVPN to MPLS Multicast VPNv4 Integration

BGP EVPN RT2 to RT5 re-originate support
 Increased VNI Scale in 2X (512)
 Catalyst 9500-H Custom SDM Template for large scale MAC/IP routes
 Increased up to 500 Leaf scale per Fabric Domain
 Optimized L2 Multicast with IGMP/MLD Snooping for Centralized Gateway
 IPv4 host overlay Multicast with Data MDT
 IPv6 host overlay Multicast with Data MDT
 Border : Multi-VRF IPv4 and IPv6 Data MDT to external Multicast Handoff
 Border : EVPN IPv4 and IPv6 Data MDT to external Multicast VPN Handoff
 Border : EVPN to Global VN Extranet support
 Border : EVPN to Non-EVPN VN Extranet support

16.9.1

16.12.1

17.3.1

17.6.1

17.9.1

17.12.1

17.15.1


EVPN VXLAN with NAT44 support
 Per-VLAN Peer-to-Peer Protected Mode
 BGP EVPN VXLANv6 Control-Plane
 VXLANv6 Underlay - IPv6 BGP EVPN AF Peering support
 VXLANv6 Overlay - Layer 3 Overlay, Distributed AnyCast GW & Layer 2 Overlay over VXLANv6
 VXLANv6 Overlay - Border Layer 2, Layer 3 & MPLS VPNv6 802.1Q IPv6 Handoff
 VXLANv6 Overlay - VXLANv4 to VXLANv6 Migration w/ IR
 VXLANv6 Overlay - IPv6 Multicast Replication for VXLANv6
 VXLANv6 Overlay - VXLANv4 to VXLANv6 Migration w/ Multicast Rep

RT-2 and RT-5 BGP EVPN Route-Map support
 Per-VLAN ESI Layer 2 Multi-home support
 EVPN Micro-Segmentation
 CLI Simplicity - Dynamic BGP EVPN Address-Family Peering Support
 CLI Simplicity - IP VRF Auto RD and Auto RT
 Scalable Fabric - 1024 VNI Scale
 Scalable Fabric - Multi-Tenant IEEE 802.3ad Layer 2 Overlay Networks
 Secure First Hop Layer 2 overlay with DHCP Snooping and ARP Inspection
 VXLANv6 Overlay - TRM Multicast - v4 Support for DAG, AnyCast RP over VXLANv6 support
 VXLANv6 Underlay - IPv6 BGP mVPN AF Peering support
 802.1Q VLAN over Layer 2 VNI Overlay

VXLANv6 Overlay - TRM Multicast - v6 Support for DAG
 C9500X | C9600X - Per-VLAN BGP EVPN ESI Layer 2 Multi-home support
 C9500X | C9600X - Centralized Gateway / Asymmetric IRB
 Tenant Routed Multicast SSO High Availability
 CLI Simplicity - L2VPN Profile
 IPv6 Neighbor Discovery Proxy for BGP EVPN Fabric
 Programmable EVPN Fabric with OpenConfig models



Cisco Catalyst 9000 BGP EVPN VXLAN Fabric



Enterprise



Healthcare



Education



Financial



Public Sector



Manufacturing



Hospitality



Media

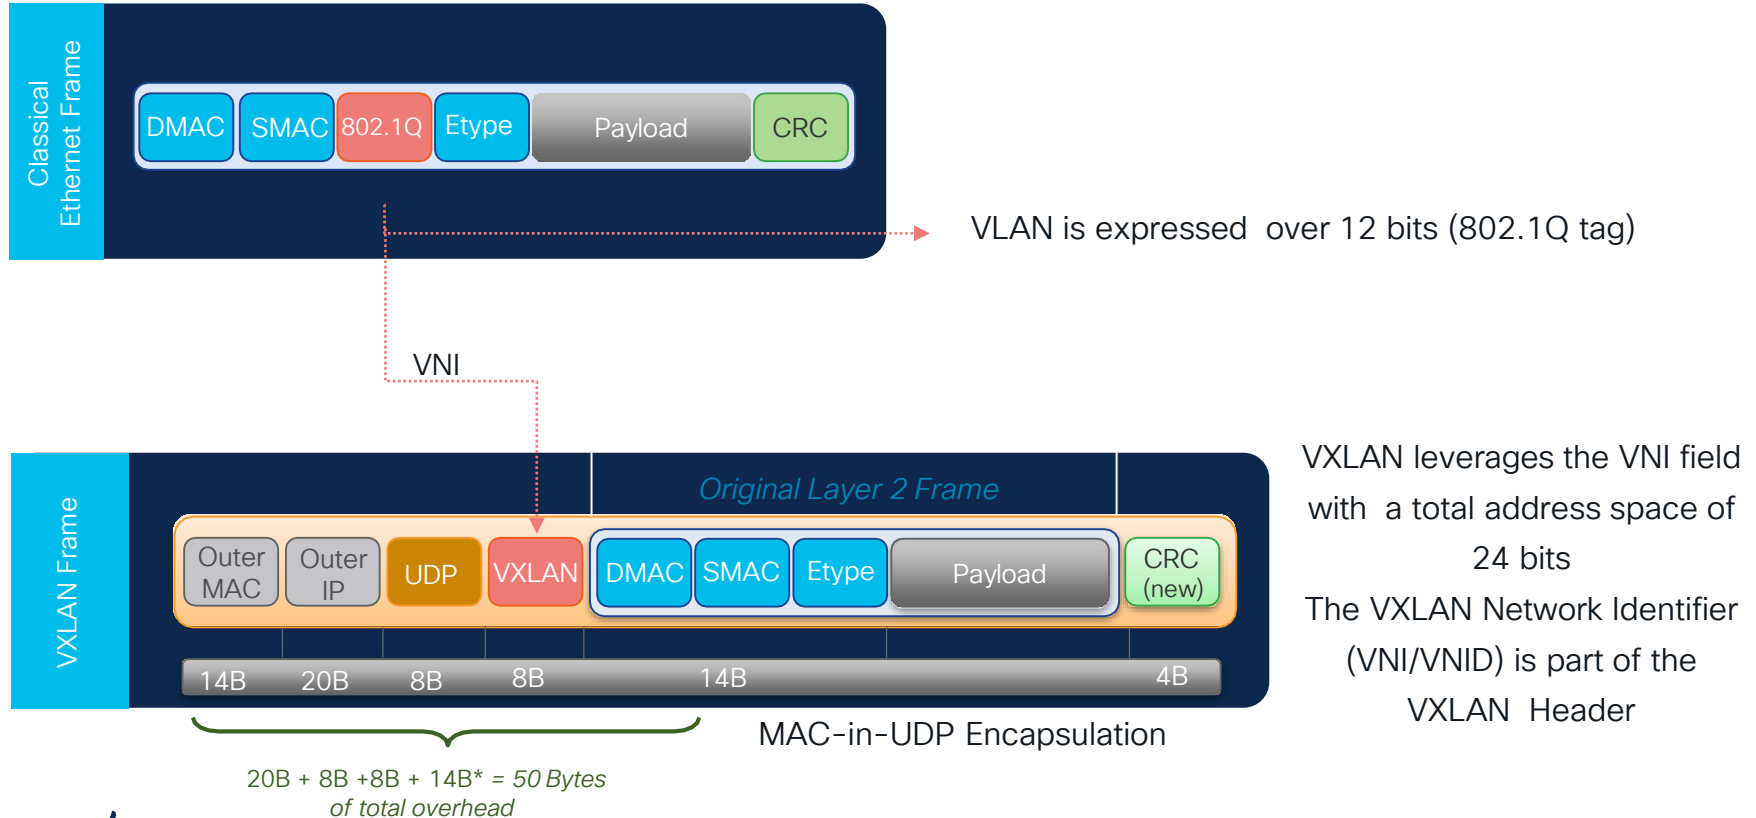


Transportation

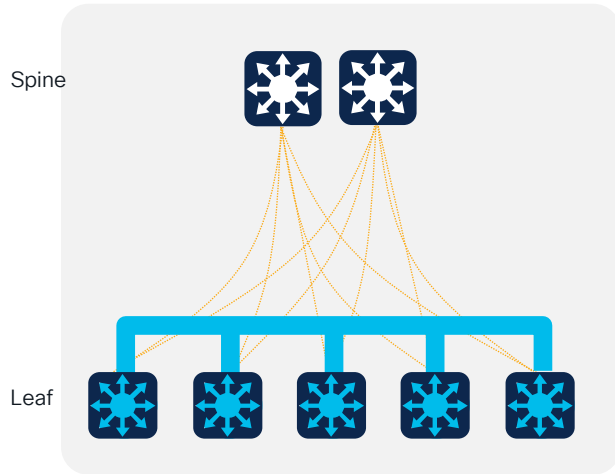


Retail

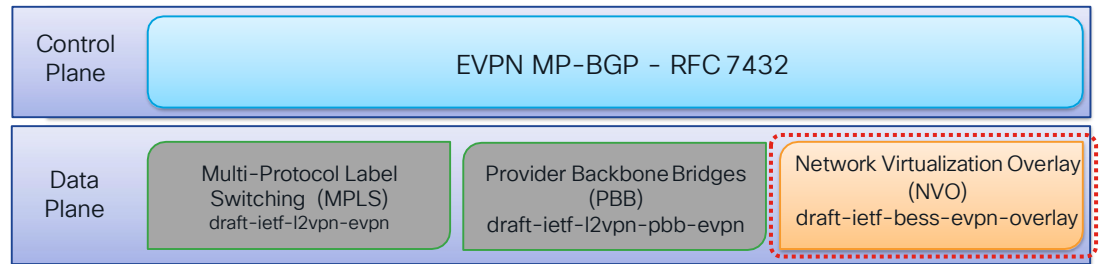
VXLAN Overview



VXLAN with BGP EVPN



- Standards based Overlay (VXLAN) with Standards based Control-Plane (BGP)
- Layer-2 MAC and Layer-3 IP information distribution by Control-Plane (BGP)
- Forwarding decision based on Control-Plane (minimizes flooding)
- Integrated Routing/Bridging (IRB) for Optimized Forwarding in the Overlay
- Multi-Tenancy At Scale




EVPN over NVO Tunnels (VXLAN)

Provides Layer-2 and Layer-3 Overlays over simple IP Networks

BGP EVPN System Role

Catalyst EVPN Scale and Performance Matrix



Cisco Catalyst BGP EVPN Configuration Guide
Scale and Performance Chapter

BORDER-GATEWAY:

A gateway point of between two or more BGP EVPN administrative domain boundary.

BORDER :

A gateway point of between EVPN fabric and external network domain.

INTERMEDIATE :

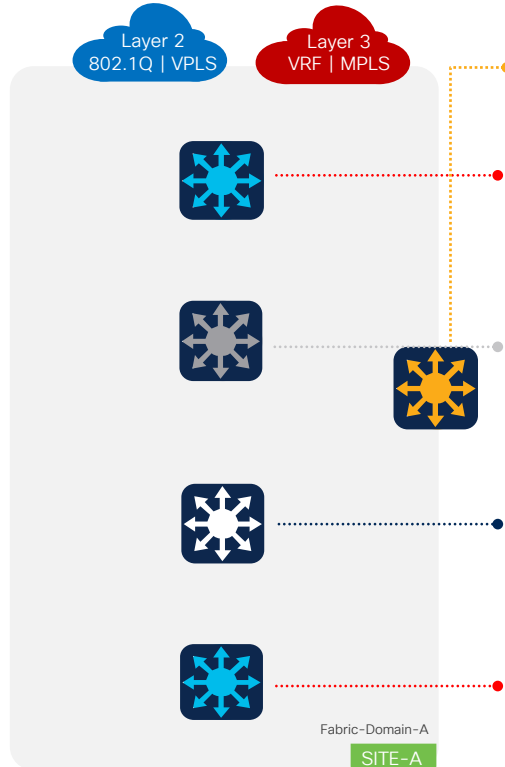
A Layer 2 or Layer 3 (IP/MPLS) Underlay network system providing basic transport and forwarding plane.

SPINE :

An BGP EVPN reflects the L2/L3 VPN prefixes providing hierarchical neighbor peering, learning and distribution point.

VTEP (LEAF) :

An origination and termination point of VXLAN enabled overlay network.



System Support	Mode
Nexus 9000	Standalone


System Support	Mode
Catalyst 9300 – 9600 (9500-H/X/9600/X)	Standalone Stack ★
Catalyst 8000 Edge ASR 1000	Physical
Nexus 9000	Standalone
ASR 9000	Standalone

System Support	Mode
Any	Any

System Support	Mode
Catalyst 9300 – 9600 (9500-H/X & 9600/X)	Standalone Stack
Catalyst 8000 Edge ASR 1000	Physical Virtual
Nexus 9000	Standalone
ASR 9000	Standalone

System Support	Mode
Catalyst 9300L 9300 9300X Series	Standalone StackWise ★
Catalyst 9400 9400X Series	Standalone StackWise-Virtual ★
Catalyst 9500 9500X Series	Standalone StackWise-Virtual ★
Catalyst 9600 9600X Series	Standalone StackWise-Virtual ★

BGP-EVPN in Campus



Enterprise Campus BGP EVPN Drivers



Industry Standard



One Fabric Architecture



Proven and Scalable



Hierarchical Fabric Domain



Flexible Overlay



Multi-vendor IT strategy



Unified operation across – Campus | DC | WAN



BGP Protocol History. Minimum new learning curve



Multi-tier Overlay network architecture

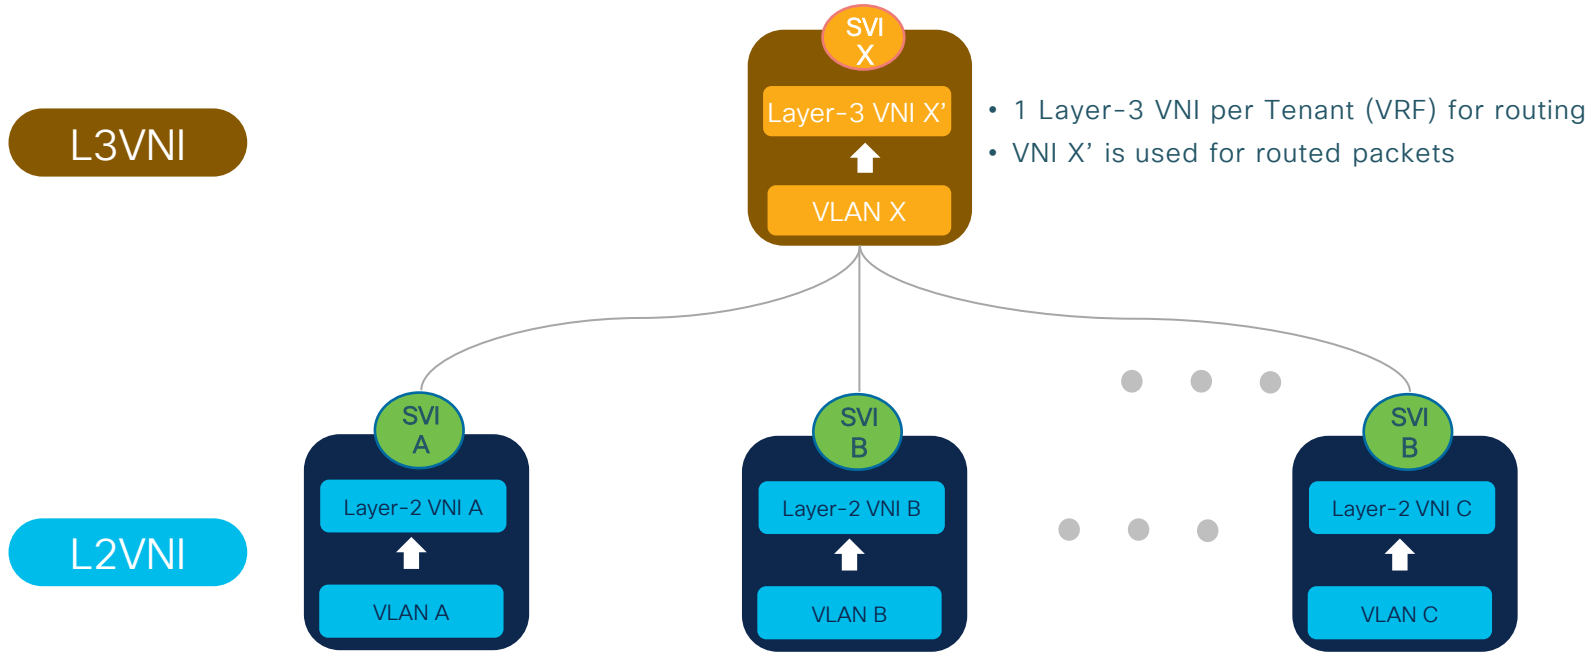


Use-case driven customize Overlay networks Types and Topologies

EVPN Basics



VXLAN Constructs



L3VNI

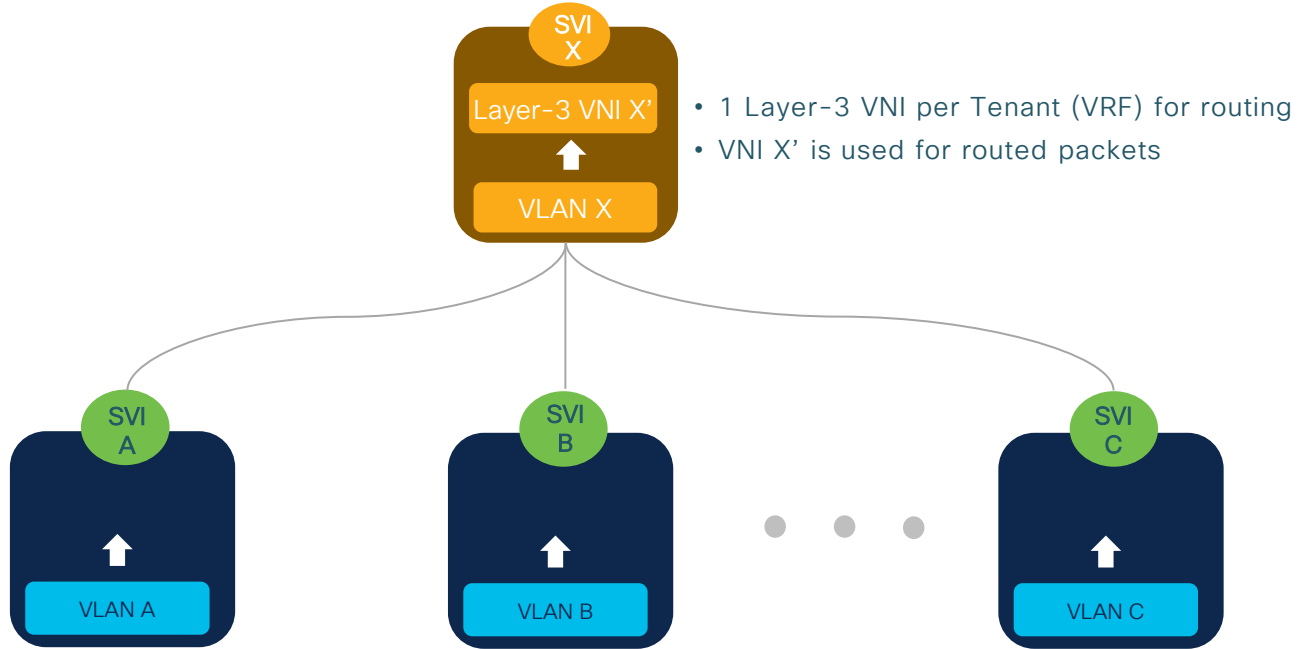
L2VNI

- 1 Layer-3 VNI per Tenant (VRF) for routing
- VNI X' is used for routed packets

- 1 Layer-2 VNI per Layer-2 segment
- L2VNI's are used for bridged packets
- Traffic between L2VNI's is routed

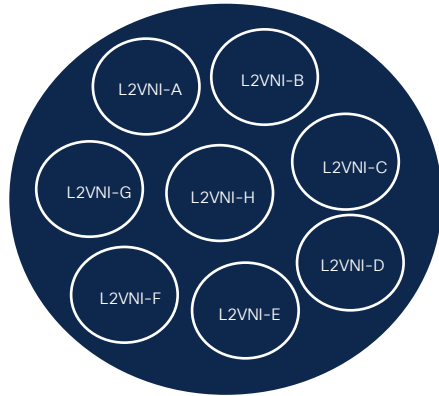
VXLAN Constructs

L3VNI

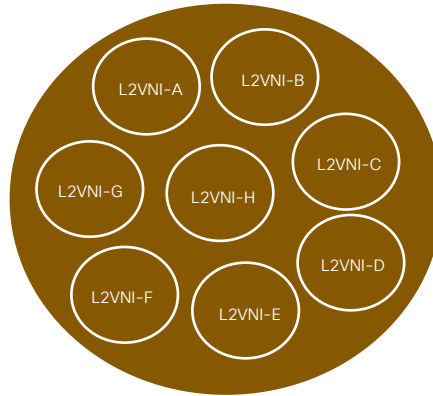


VXLAN Constructs

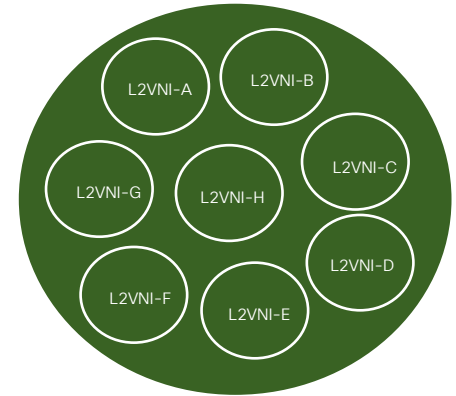
VRF-X



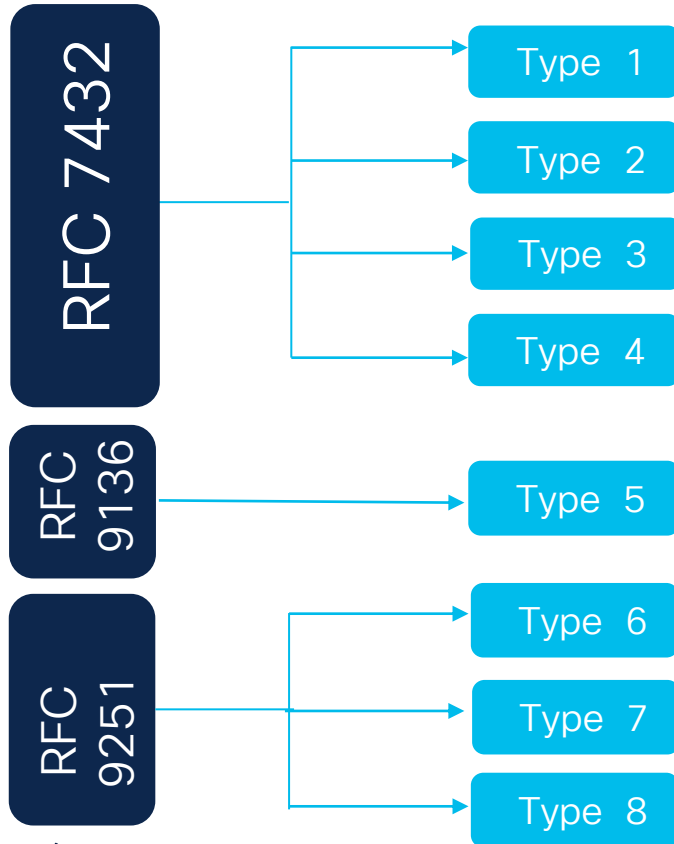
VRF-Y



VRF-Z



MP-BGP EVPN Route Type(s)



Ethernet Auto-Discovery (A-D) route

MAC/IP advertisement route

Inclusive Multicast Route

EVPN Ingress Replication (IR) (unicast mode for BUM)

Ethernet Segment Route

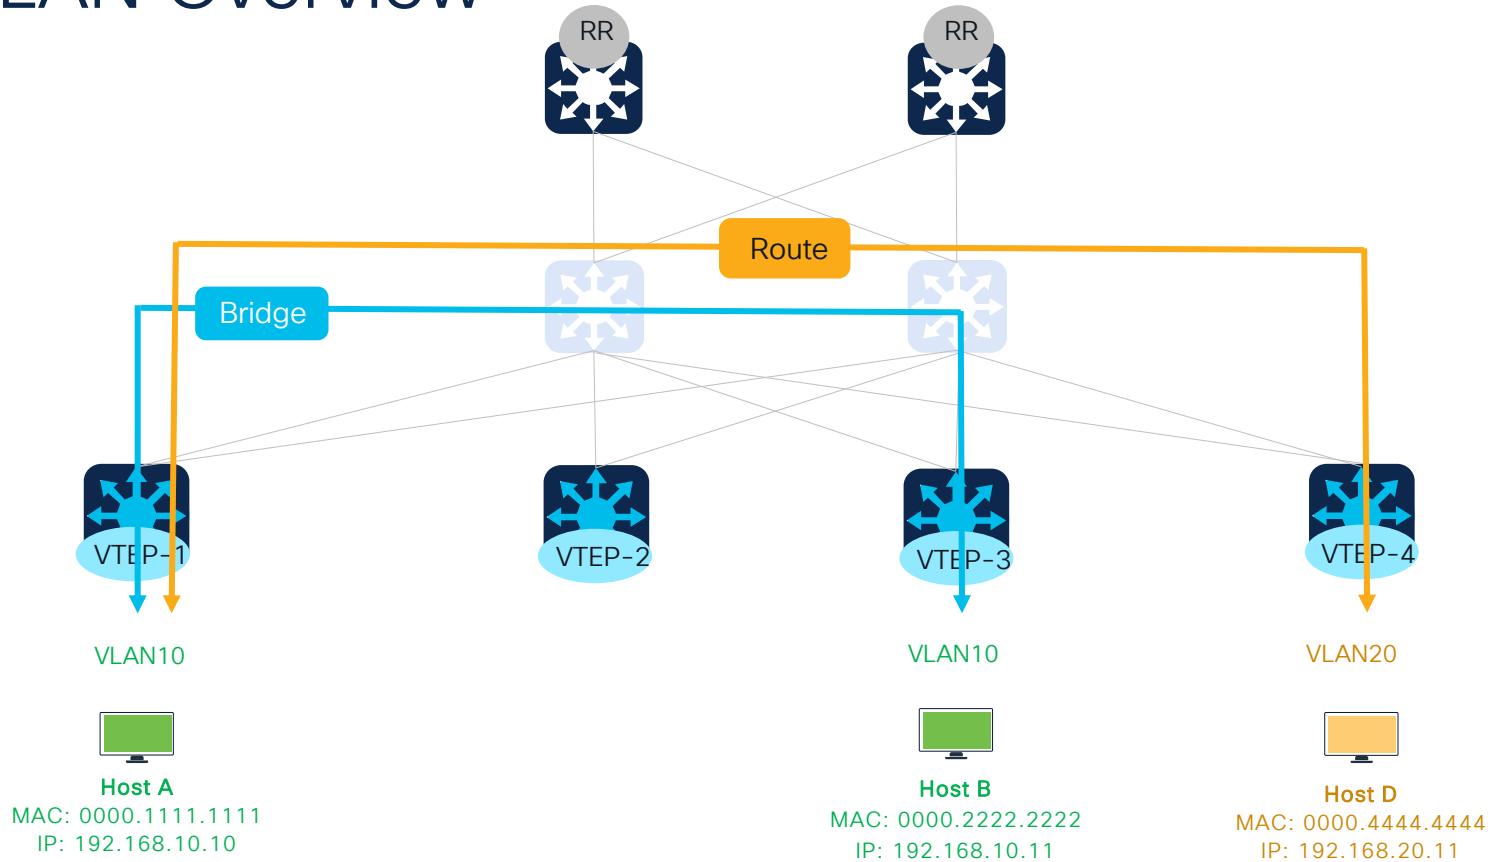
IP Prefix Route → Layer-3 VNI Route

Selective Multicast Ethernet Tag Route

IGMP Join Sync routes (Join/Leave)

IGMP Leave Sync routes (Join/Leave)

VXLAN Overview

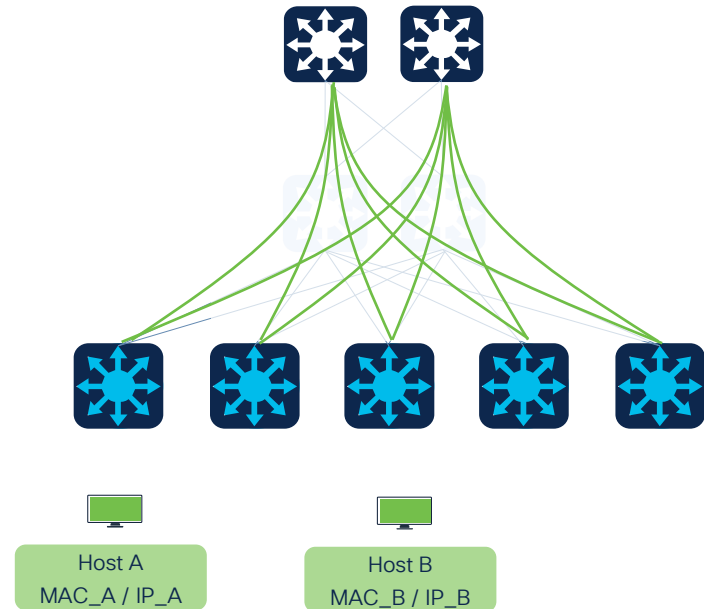


MAC/IP Advertisement route

Route-Type 2

- Host “A” attaches to Edge Device (VTEP)
- VTEP V1 advertises Host “A” reachability information
 - MAC and L2VNI [mandatory]
 - IP and L3VNI [optional]
 - depending on ARP
- Additional Attributes advertised
 - MPLS Label 1 (Layer-2 VNI)
 - MPLS Label 2 (Layer-3 VNI)
 - Extended Communities

Route Type	MAC, IP	L2VNI	Layer-3 VNI (“VRF”)	NH	Encap	Seq
2	MAC_A, IP_A	30001	50001	IP_V1	8:VXLAN	0

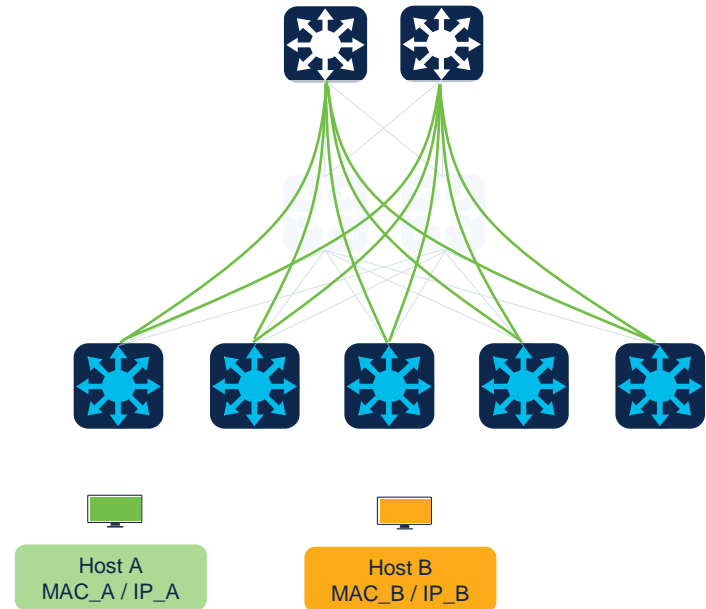


Protocol Learning & Distribution

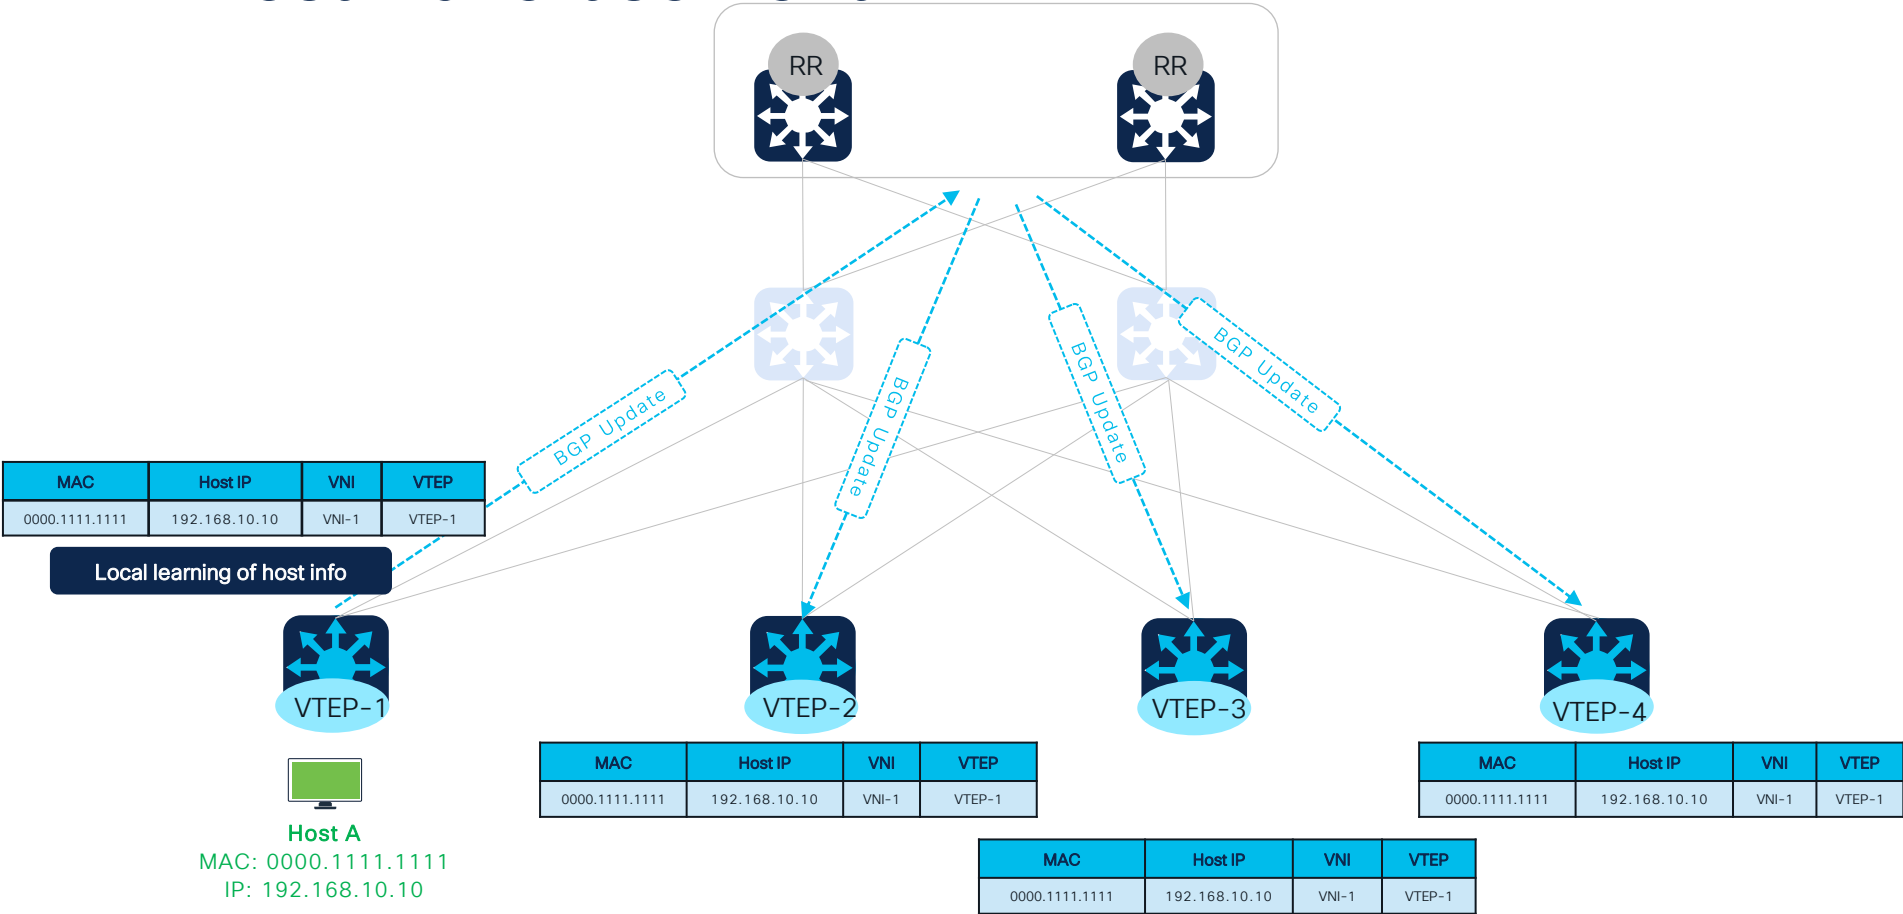
“Subnet Route Advertisement (Route-Type 5)”

- IP Prefix Redistribution
 - From “Direct” (connected), Static or dynamically learned Routes
- VTEP V1 advertises local Subnet through redistribution of “Direct” (connected) routes
 - IP Prefix, IP Prefix Length, and Layer-3 VNI
- Additional route attributes advertised
 - MPLS Label (Layer-3 VNI)
 - Extended Communities
- Multiple VTEPs can announce same IP Prefix

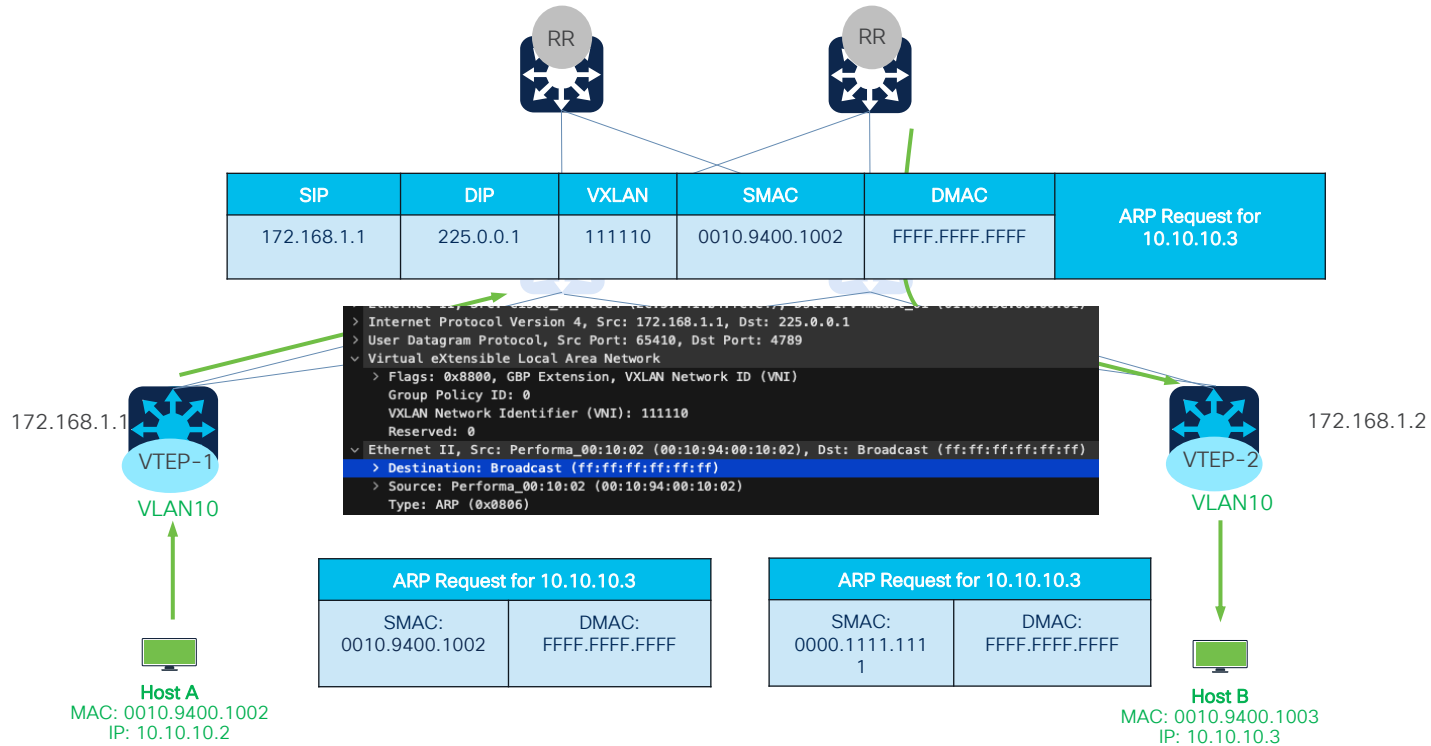
Route Type	MAC, IP	Layer-3 VNI (“VRF”)	NH	Encap
5	Subnet_A/24	50001	IP_V1	8:VXLAN



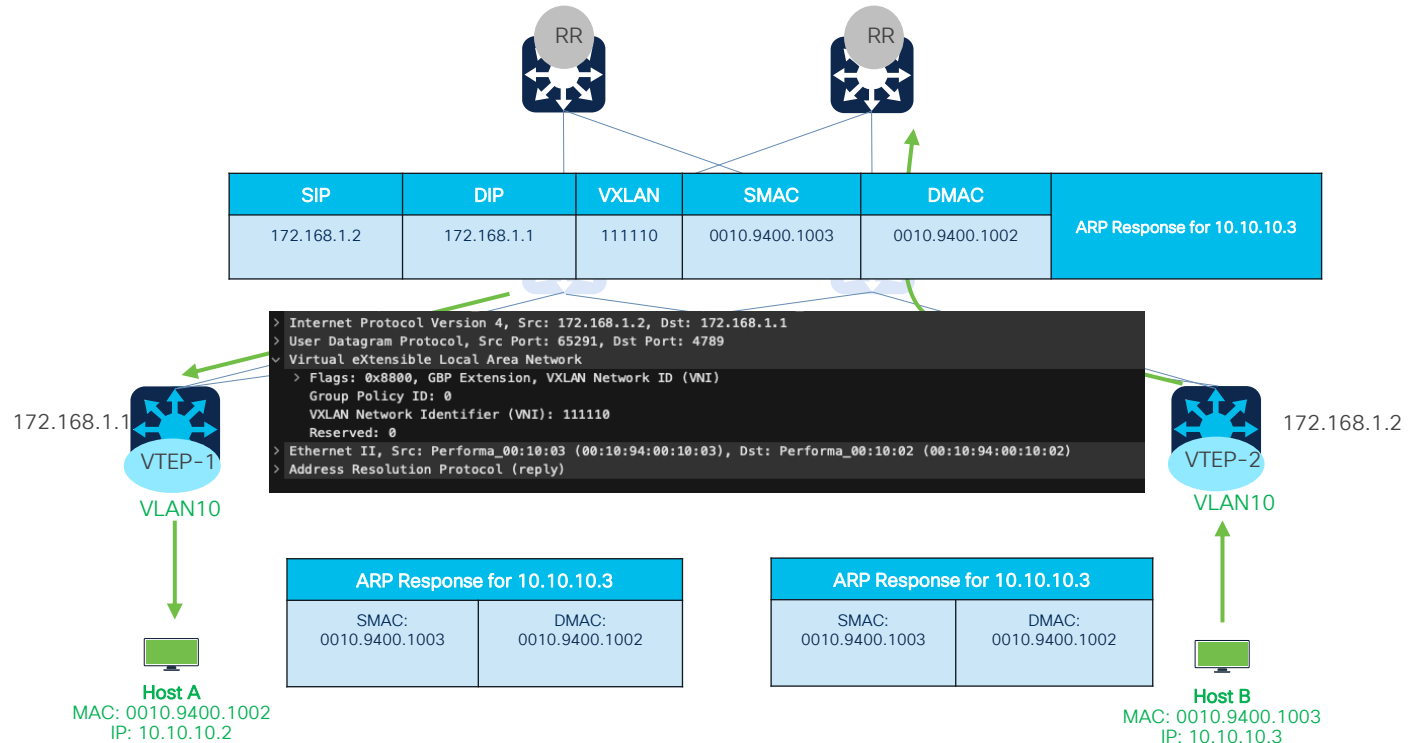
Host Advertisement



Packet Walk – ARP Request



Packet Walk - ARP Request



Overlay Leaf Configuration – BGP EVPN Control Plane

Spine-1

```
router bgp 64500
neighbor 3.3.3.3 remote-as 64500
neighbor 3.3.3.3 update-source Loopback0
neighbor 4.4.4.4 remote-as 64500
neighbor 4.4.4.4 update-source Loopback0
!
address-family l2vpn evpn
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 send-community both
neighbor 3.3.3.3 route-reflector-client
neighbor 4.4.4.4 activate
neighbor 4.4.4.4 send-community both
neighbor 4.4.4.4 route-reflector-client
maximum-paths 2
```

Leaf-1

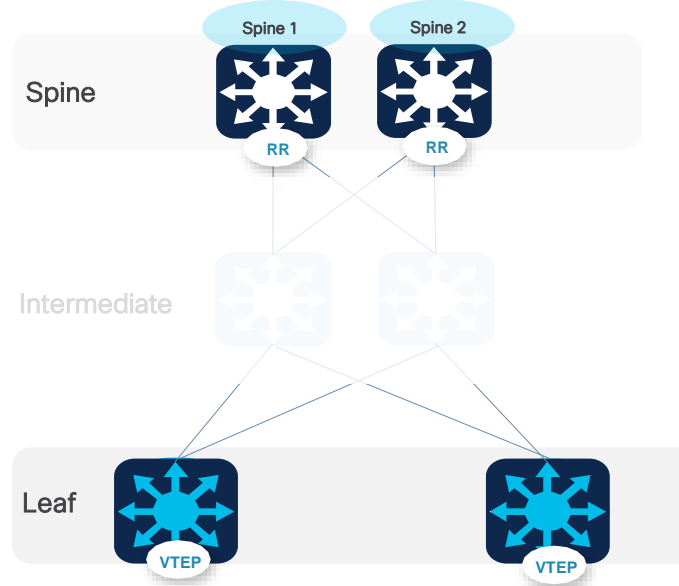
```
router bgp 65000
neighbor 1.1.1.1 remote-as 64500
neighbor 1.1.1.1 update-source Loopback0
neighbor 2.2.2.2 remote-as 64500
neighbor 2.2.2.2 update-source Loopback0
!
address-family l2vpn evpn
neighbor 1.1.1.1 activate
neighbor 1.1.1.1 send-community both
neighbor 2.2.2.2 activate
neighbor 2.2.2.2 send-community both
exit-address-family
```

Spine-2

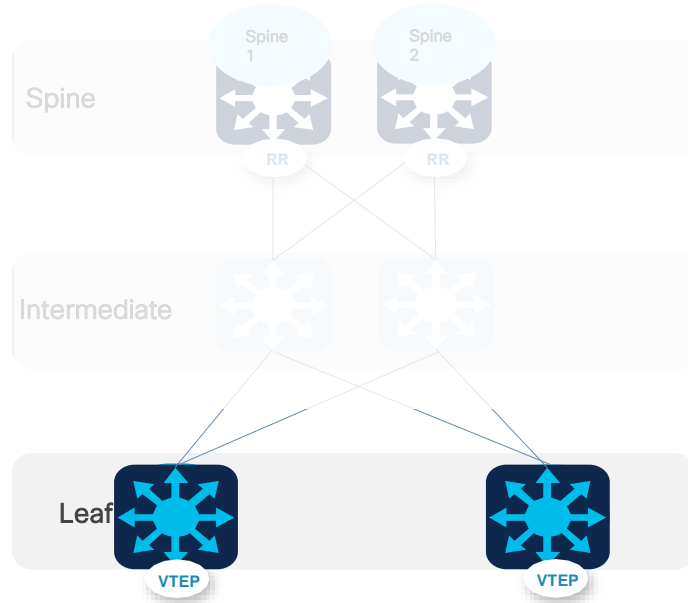
```
router bgp 64500
neighbor 3.3.3.3 remote-as 64500
neighbor 3.3.3.3 update-source Loopback0
neighbor 4.4.4.4 remote-as 64500
neighbor 4.4.4.4 update-source Loopback0
!
address-family l2vpn evpn
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 send-community both
neighbor 3.3.3.3 route-reflector-client
neighbor 4.4.4.4 activate
neighbor 4.4.4.4 send-community both
neighbor 4.4.4.4 route-reflector-client
maximum-paths 2
```

Leaf-2

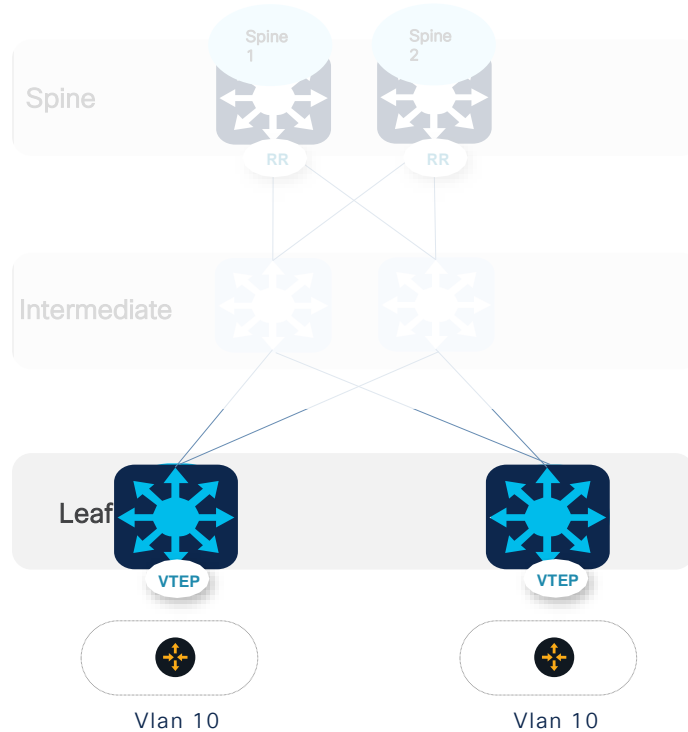
```
router bgp 65000
neighbor 1.1.1.1 remote-as 64500
neighbor 1.1.1.1 update-source Loopback0
neighbor 2.2.2.2 remote-as 64500
neighbor 2.2.2.2 update-source Loopback0
!
address-family l2vpn evpn
neighbor 1.1.1.1 activate
neighbor 1.1.1.1 send-community both
neighbor 2.2.2.2 activate
neighbor 2.2.2.2 send-community both
exit-address-family
```



VNI Configuration



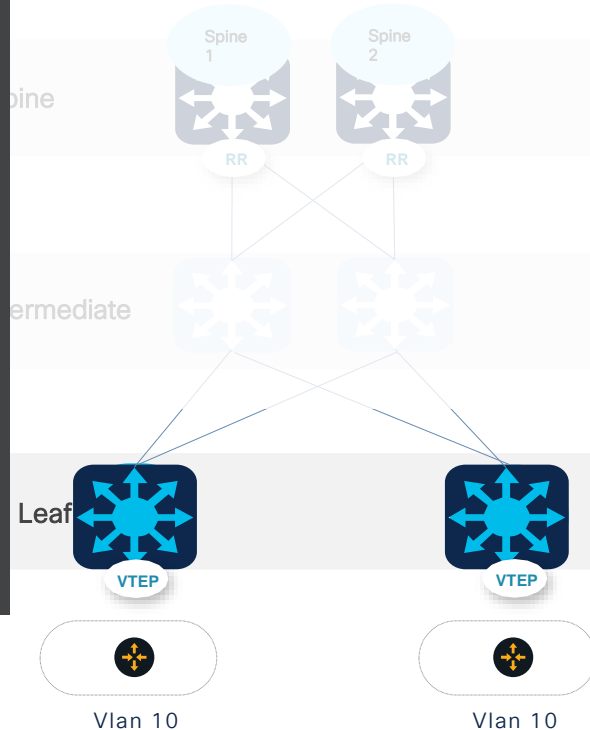
VNI Configuration



VNI Configuration

Leaf-1

```
l2vpn evpn
 replication-type static
 router-id Loopback0
 !
l2vpn evpn instance 10 vlan-based
 encapsulation vxlan
 !
vlan configuration 10
 member evpn-instance 10 vni 101000
 !
vlan configuration 102
 member vni 102102
 !
interface Vlan10
 mac-address 0010.0010.0010
 vrf forwarding S1-EVPN
 ip address 172.16.10.1 255.255.255.0
 !
interface Vlan102
 vrf forwarding S1-EVPN
 ip unnumbered Loopback0
```



Leaf-2

```
l2vpn evpn
 replication-type static
 router-id Loopback0
 !
l2vpn evpn instance 10 vlan-based
 encapsulation vxlan
 !
vlan configuration 10
 member evpn-instance 10 vni 101000
 !
vlan configuration 102
 member vni 102102
 !
interface Vlan10
 mac-address 0010.0010.0010
 vrf forwarding S1-EVPN
 ip address 172.16.10.1 255.255.255.0
 !
interface Vlan102
 vrf forwarding S1-EVPN
 ip unnumbered Loopback0
```

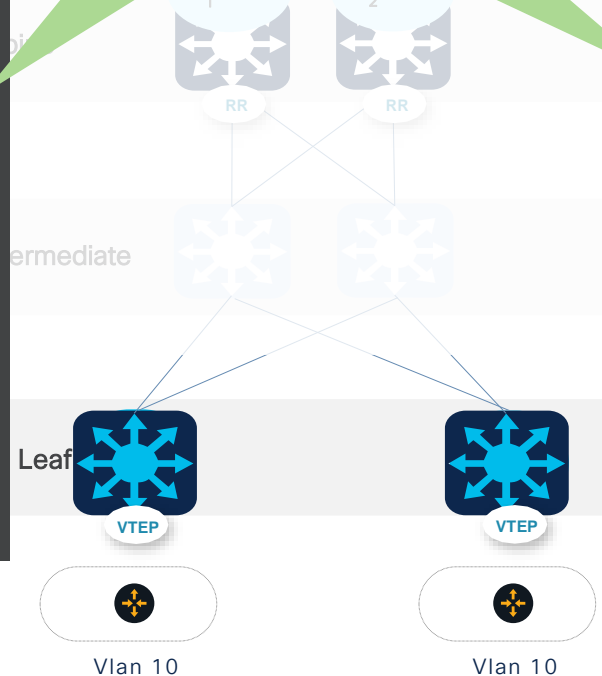
VNI Configuration

Leaf-1

```
l2vpn evpn
replication-type static
router-id Loopback0
!
l2vpn evpn instance 10 vlan-based
encapsulation vxlan
!
vlan configuration 10
member evpn-instance 10 vni 101000
!
vlan configuration 102
member vni 102102
!
interface Vlan10
mac-address 0010.0010.0010
vrf forwarding S1-EVPN
ip address 172.16.10.1 255.255.255.0
!
interface Vlan102
vrf forwarding S1-EVPN
ip unnumbered Loopback0
```

L2VNI Configuration

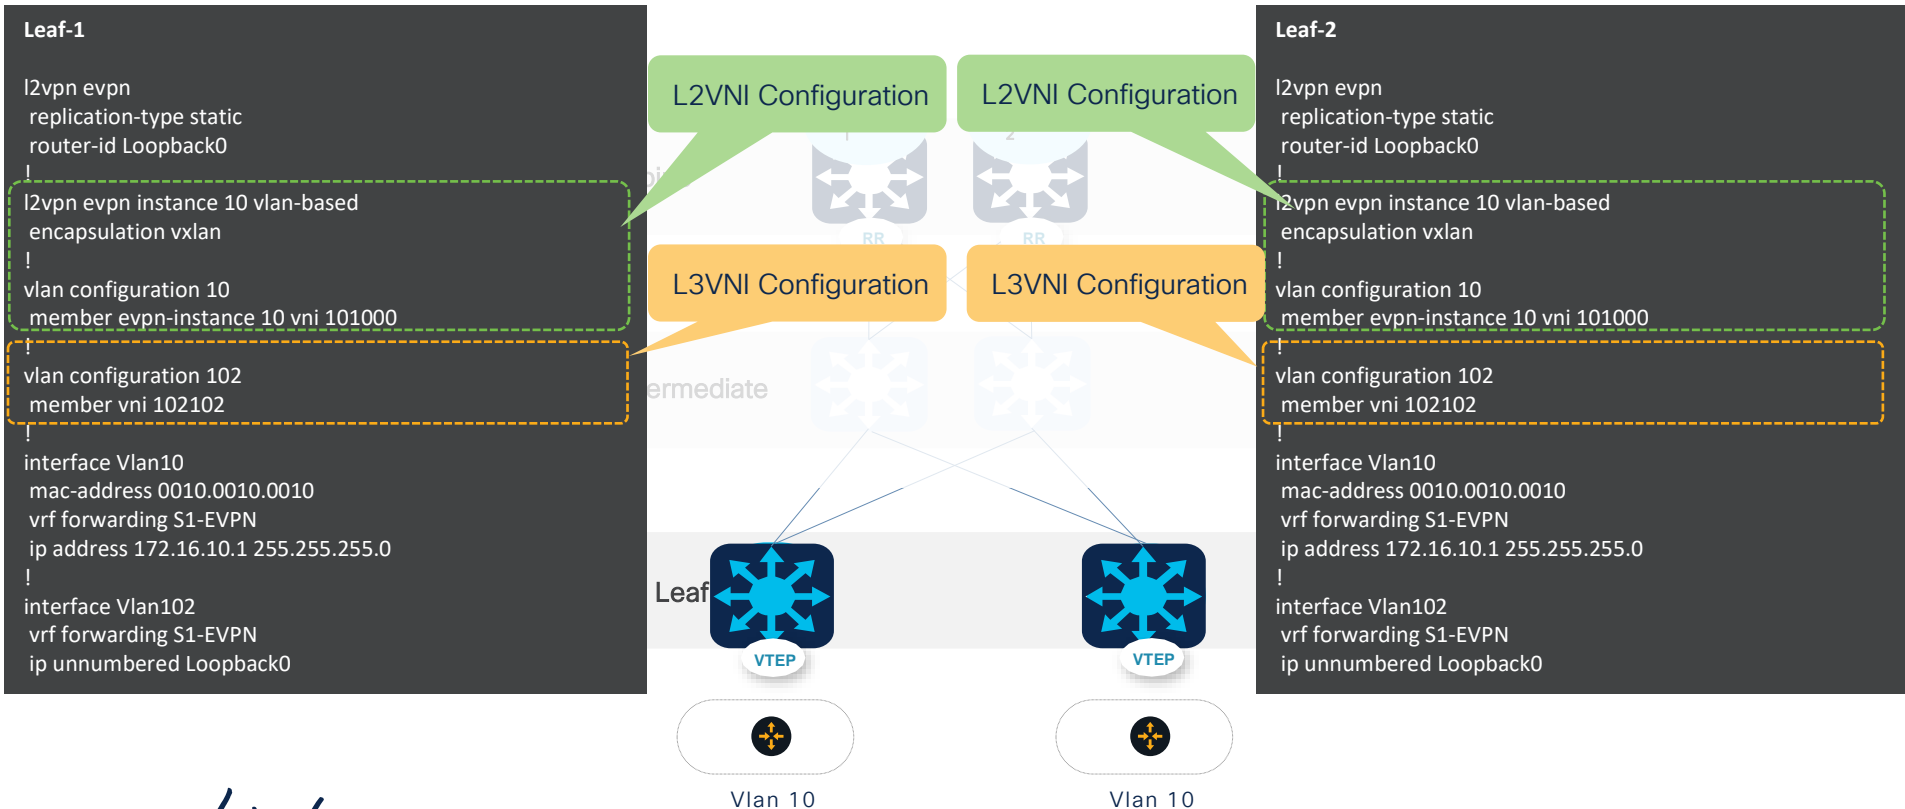
L2VNI Configuration



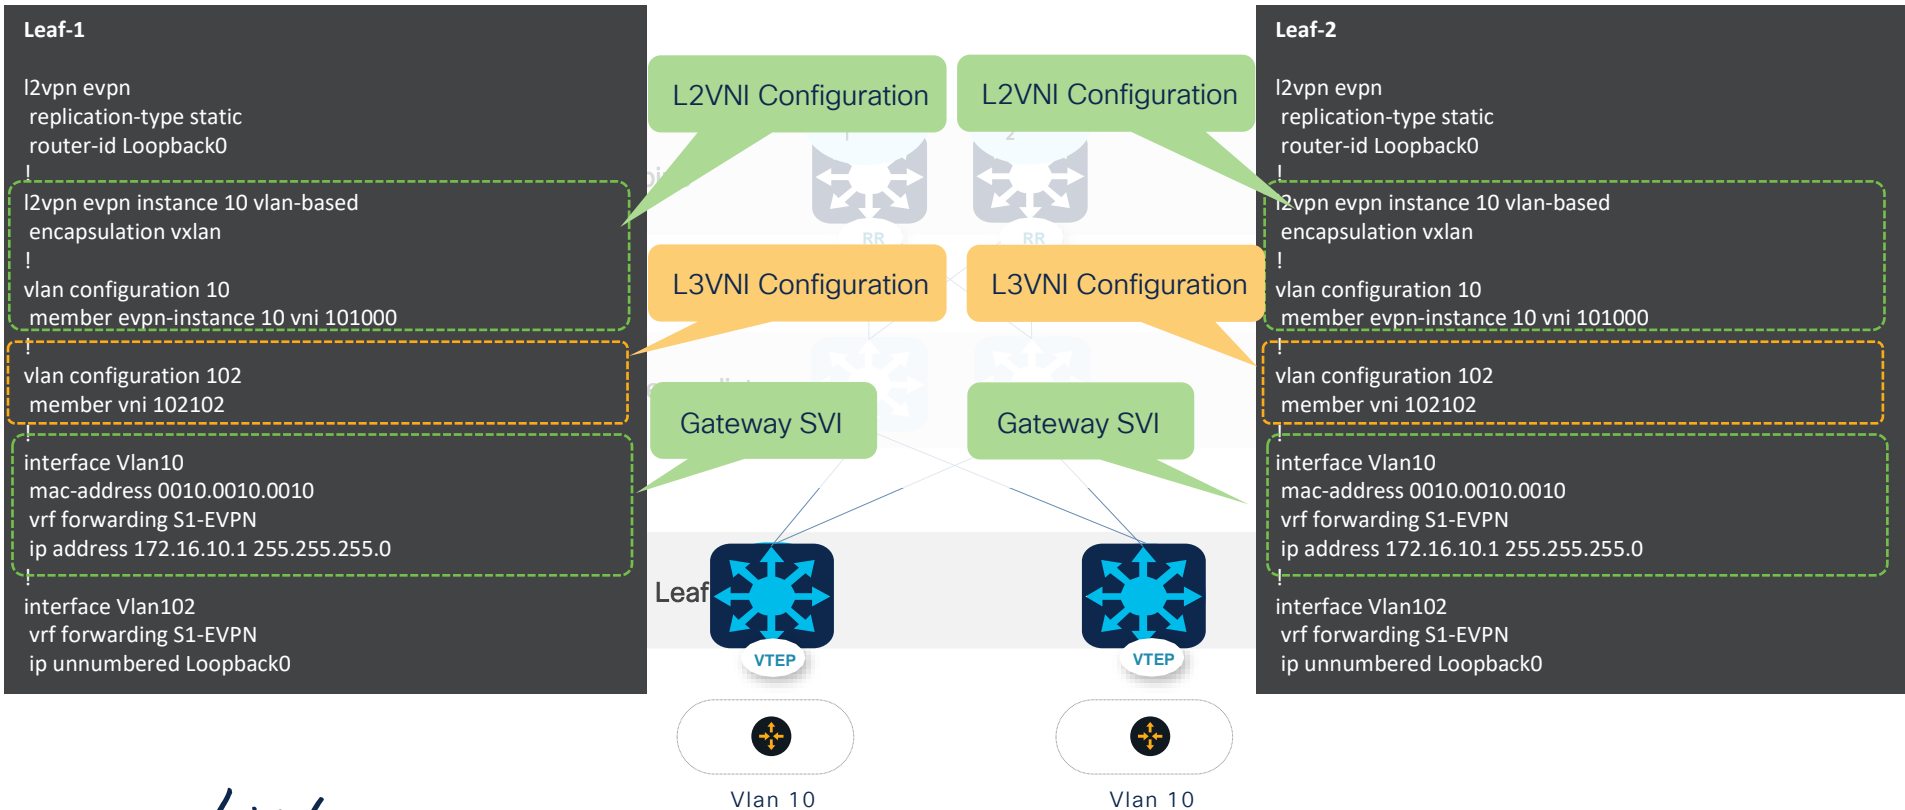
Leaf-2

```
l2vpn evpn
replication-type static
router-id Loopback0
!
l2vpn evpn instance 10 vlan-based
encapsulation vxlan
!
vlan configuration 10
member evpn-instance 10 vni 101000
!
vlan configuration 102
member vni 102102
!
interface Vlan10
mac-address 0010.0010.0010
vrf forwarding S1-EVPN
ip address 172.16.10.1 255.255.255.0
!
interface Vlan102
vrf forwarding S1-EVPN
ip unnumbered Loopback0
```

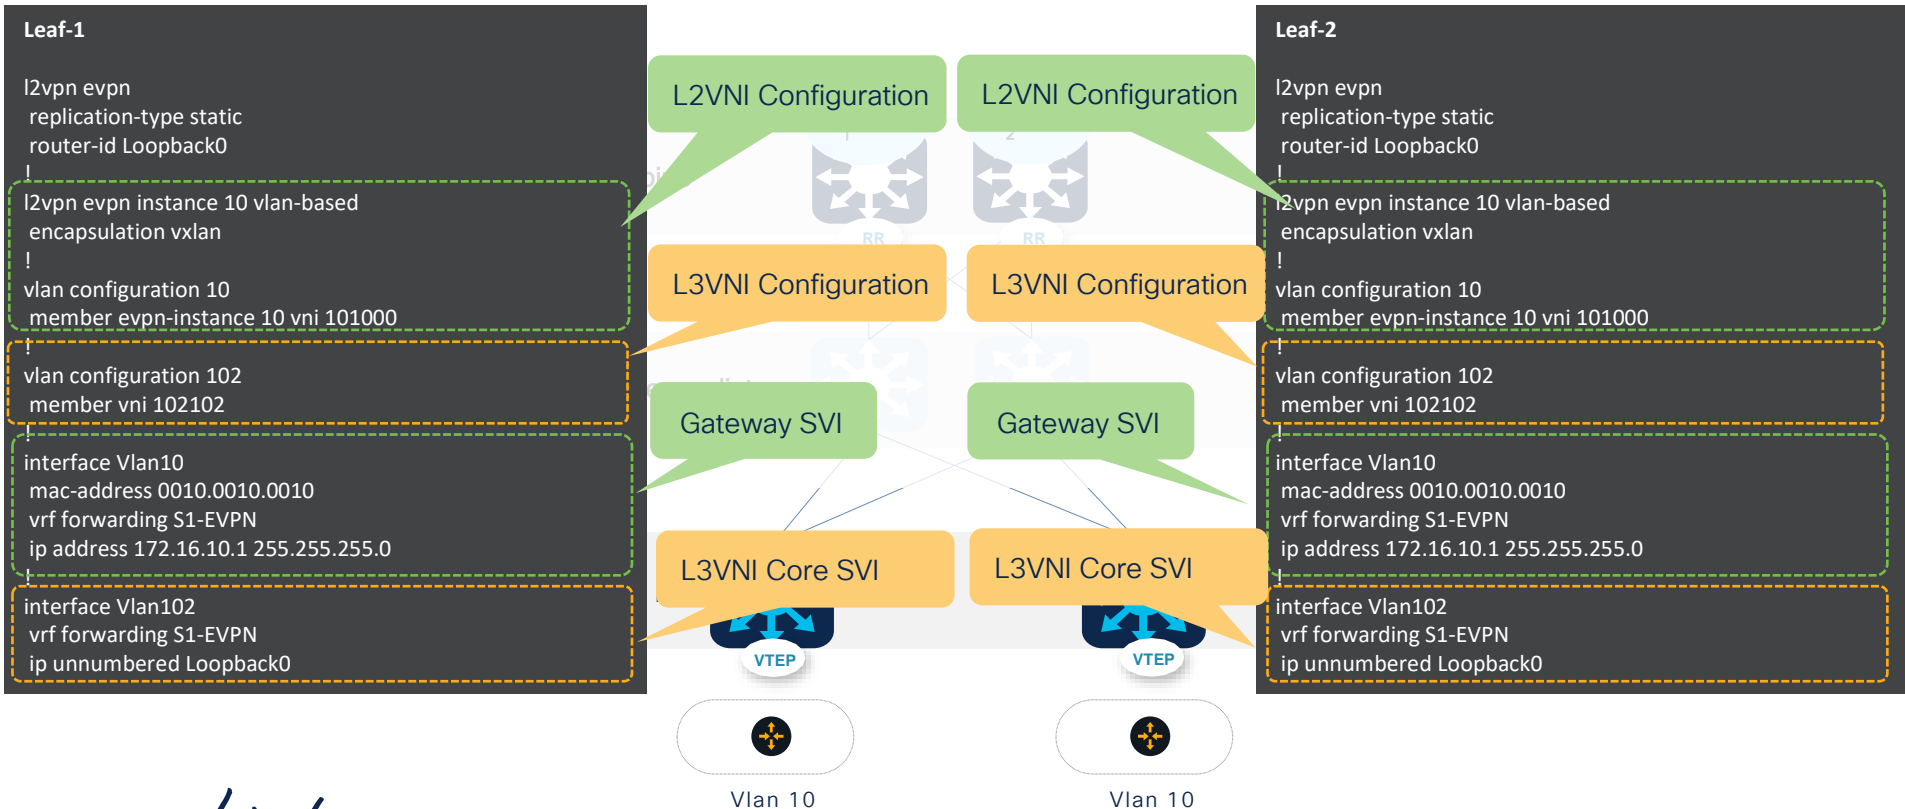
VNI Configuration



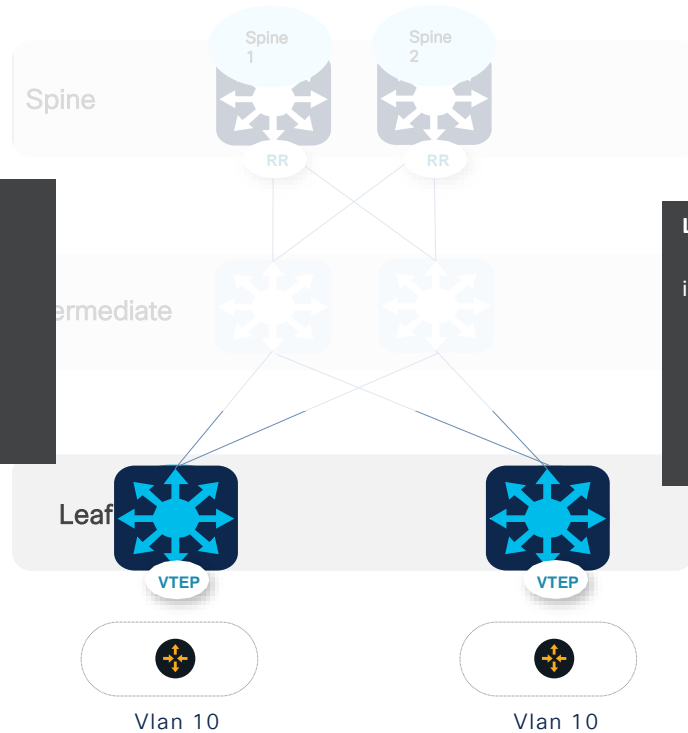
VNI Configuration



VNI Configuration



NVE Configuration



Leaf-1

```
interface nve1
no ip address
source-interface Loopback0
host-reachability protocol bgp
member vni 102102 vrf S1-EVPN
member vni 101000 ingress-replication
```

Leaf-1

```
interface nve1
no ip address
source-interface Loopback0
host-reachability protocol bgp
member vni 102102 vrf S1-EVPN
member vni 101000 ingress-replication
```

Efficient Layer 2 Broadcast domain

Ingress Replication

1 x Broadcast Packet



No Multicast in Underlay

4 x Unicast Packets



```
interface nve1
no ip address
source-interface Loopback0
host-reachability protocol bgp
member vni 10103 vrf green
member vni 10102 ingress-replication
```

Multicast Replication

1 x Broadcast Packet




Source

1 x Multicast Packet

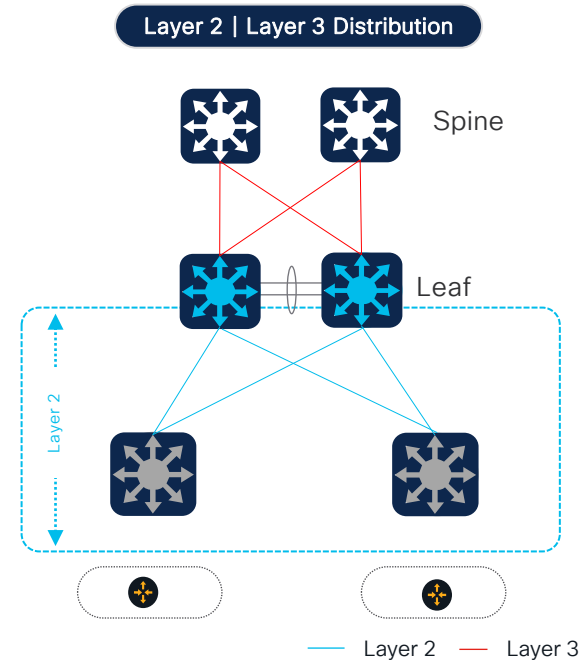
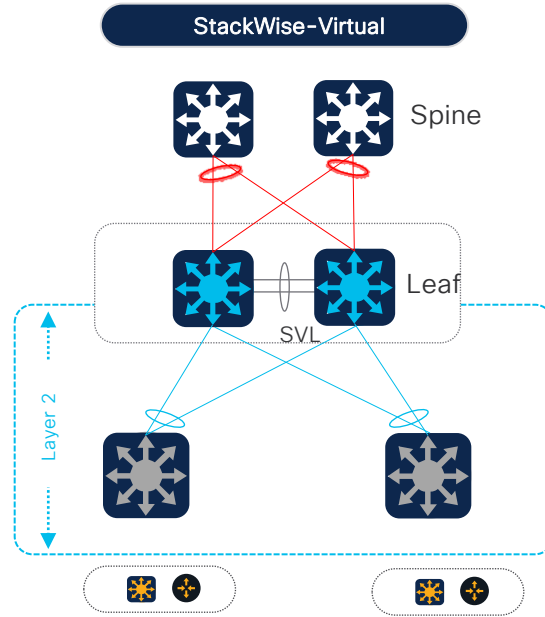
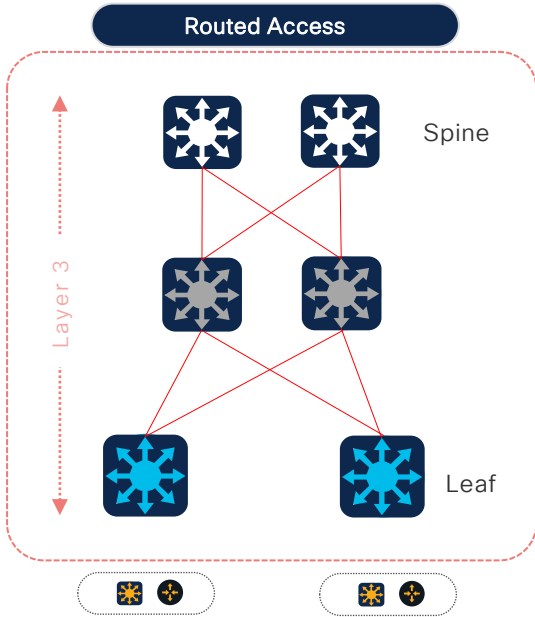


```
interface nve1
no ip address
source-interface Loopback0
host-reachability protocol bgp
member vni 10104 vrf blue
member vni 10101 mcast-group 225.0.0.1
```

Underlay Network



Underlay Design Considerations



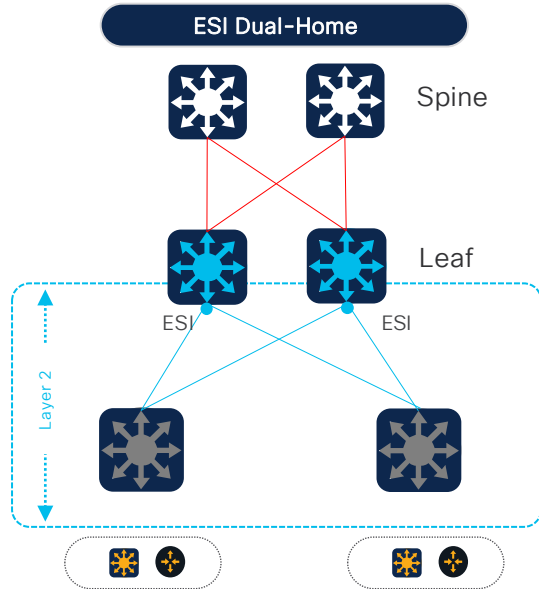
— Layer 2 — Layer 3

- Routed Access environment
- Leaf Layer – Routed Access
- Spine/RR – Direct | Multi-hop
- ECMP | Non-Blocking ports
- Underlay | Overlay IP Gateway
- L2 | L3 Overlay Support
- Multicast Support

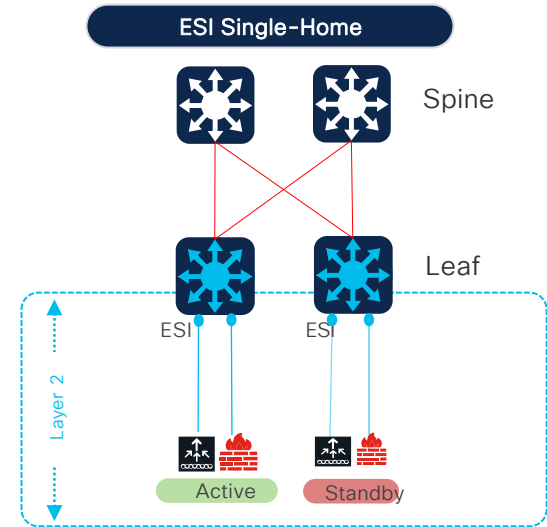
- Access – Traditional Layer 2
- Leaf Layer – Distribution
- Spine/RR – Direct | Multi-hop
- MEC | ECMP | Active/Active forwarding
- Underlay | Overlay IP Gateway
- L2 | L3 Overlay Support
- Multicast Support

- Access – Traditional Layer 2
- Leaf Layer – Distribution
- Spine/RR – Direct | Multi-hop
- FHRP | ECMP | Multicast
- Underlay | Overlay IP Gateway
- L3 Overlay Support. No L2 Extension
- Multicast Support

EVPN ESI Dual-Home (Ethernet Segment Identifier)

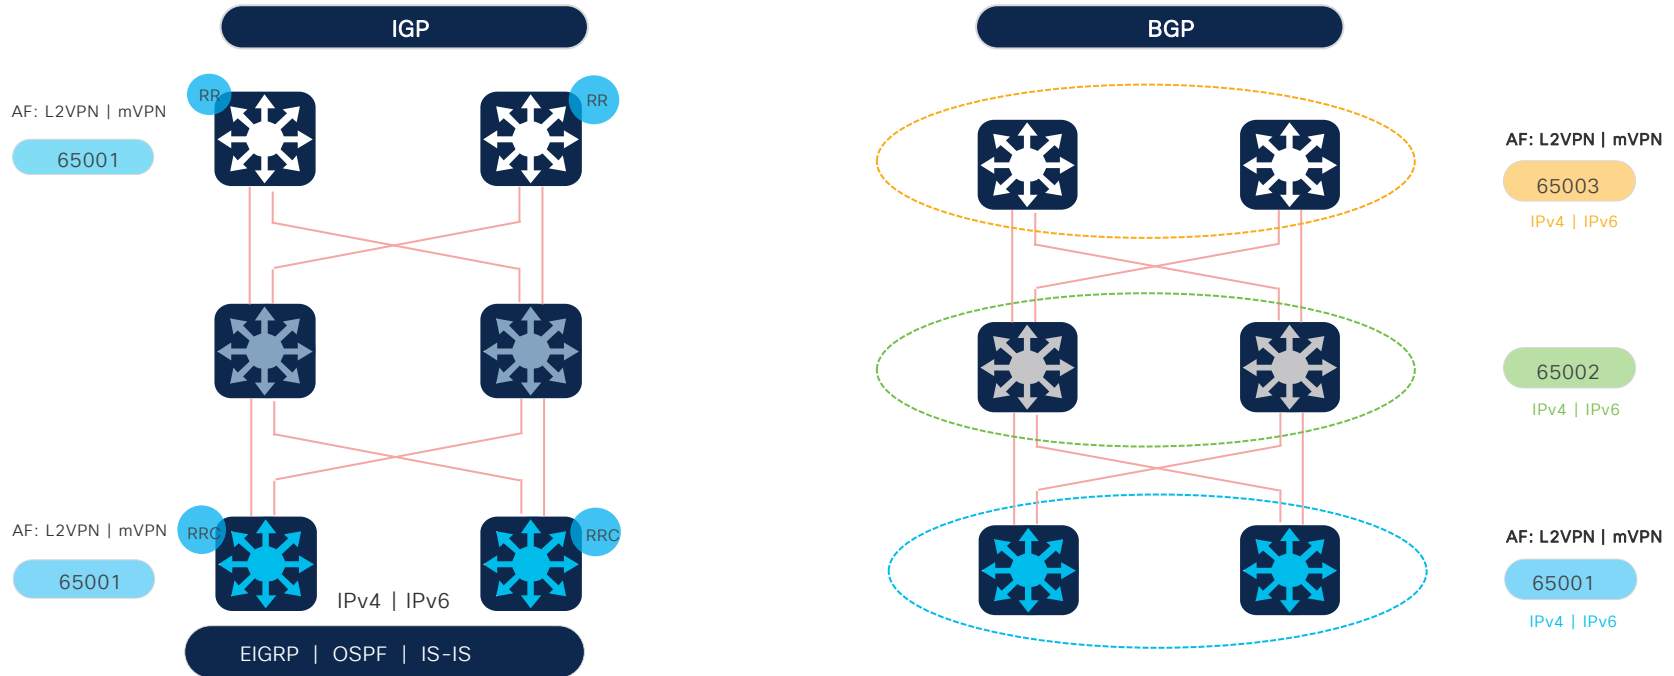


- Access – Traditional Layer 2
- Leaf Layer – Distribution
- Spine/RR – Direct | Multi-hop
- Per-ESI Anycast Gateway
- Per-VLAN | FHRP | ECMP | Multicast
- Active / Standby load-balancing
- L2 | L3 Overlay support
- Multicast Support



- Access – Traditional Layer 2
- Leaf Layer – Distribution
- Spine/RR – Direct | Multi-hop
- Per-ESI Anycast Gateway
- Per-VLAN | FHRP | ECMP | Multicast
- Active / Standby load-balancing
- L2 | L3 Overlay support
- Multicast Support

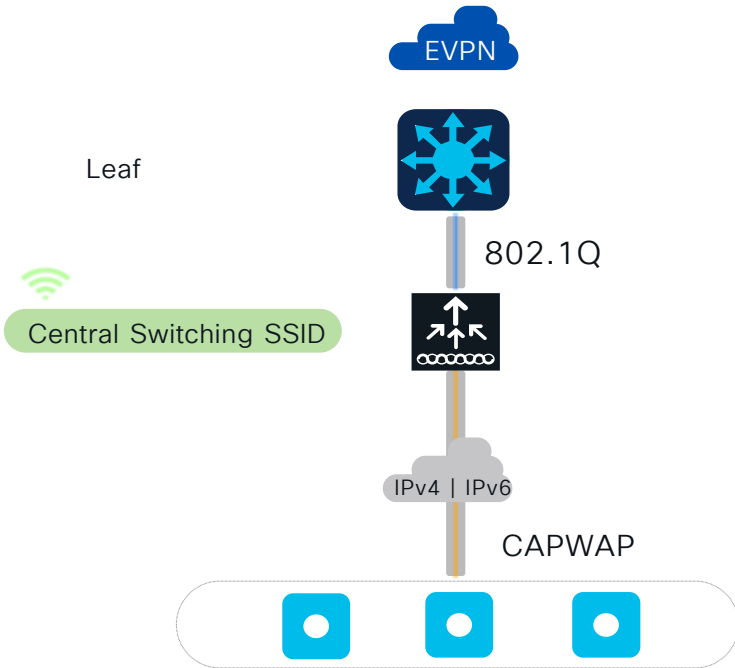
Underlay Unicast Routing Design Alternatives



- Flexible Underlay Unicast alternatives – IGP (EIGRP/OSPF/IS-IS) or BGP
- Physical/Virtual Spine RR support – IOS-XE | NXOS | XR
- Secure link-layer underlay network encryption using MACSEC
- Underlay MTU size consideration. TCP MSS adjust supported.

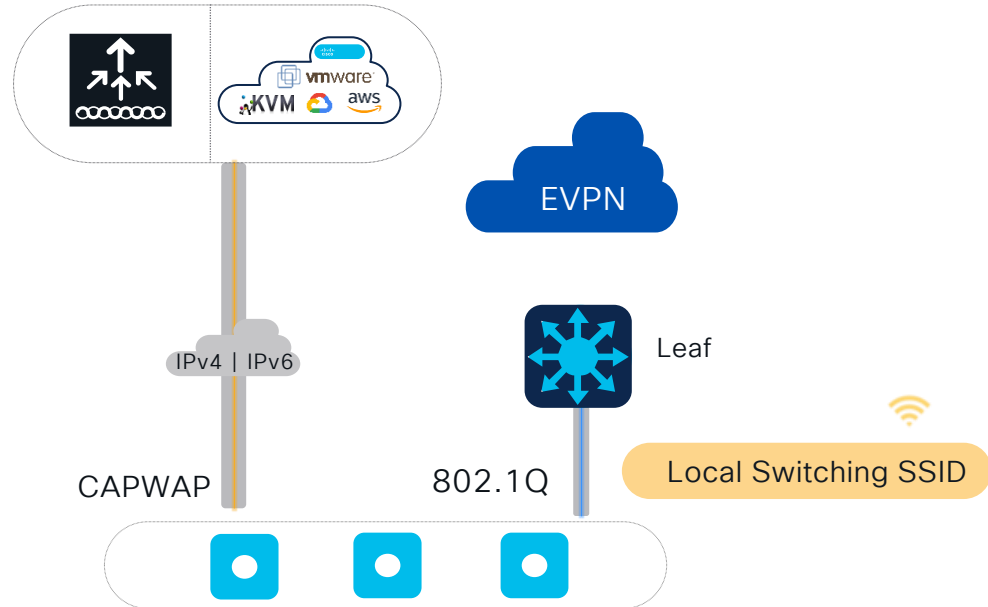
Local Mode Wireless

Central Switching



FlexConnect Mode Wireless

Local Switching



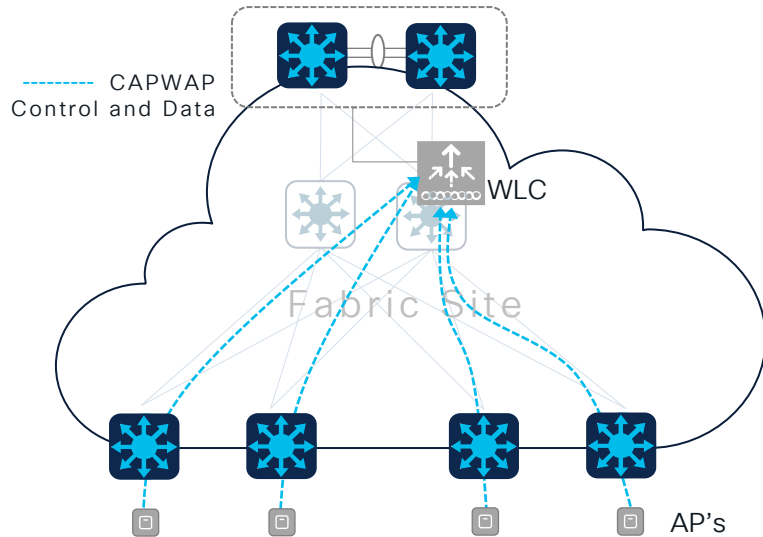
Wireless

Over the Top Wireless. Intact WLC and AP communication in Underlay
Flexible SSID alternatives - Central Switching, Local Switching, Central + Local Switching
Fabric boundary initiates from Wireless Client IP gateway.

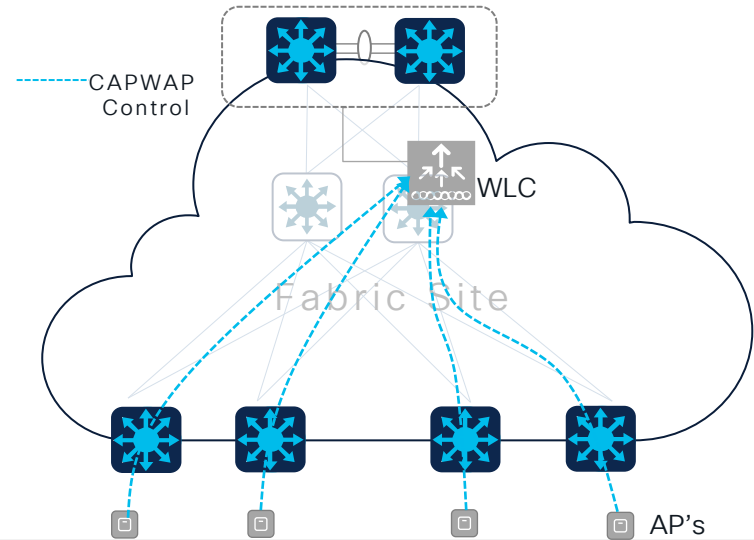
Vxlan with BGP EVPN

Wireless Options

Central Switching




Flex Connect Mode



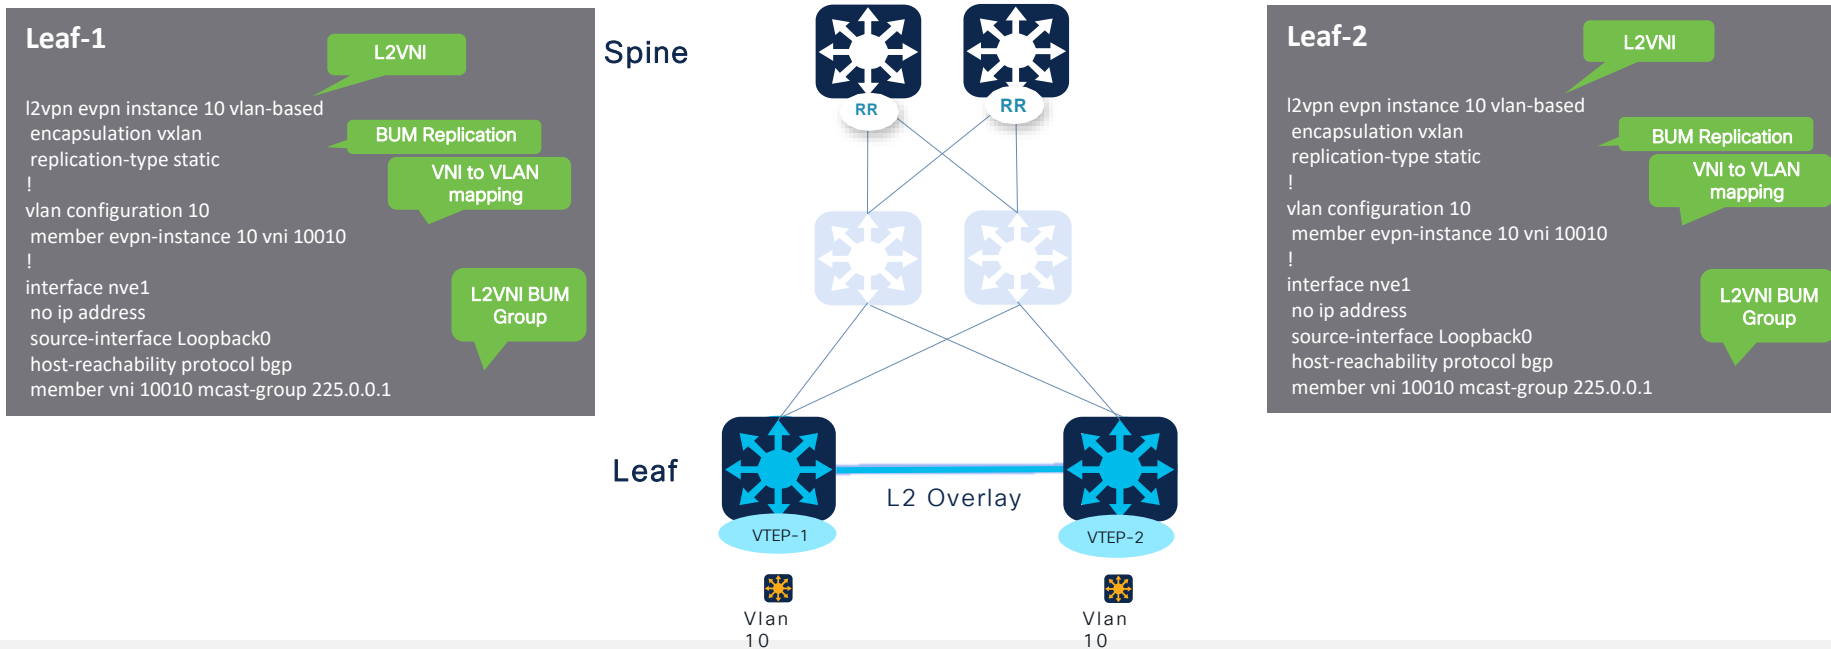
Wireless

Over the Top Wireless. Intact WLC and AP communication in Underlay
Flexible SSID alternatives - Central Switching, Local Switching, Central + Local Switching
Fabric boundary initiates from Wireless Client IP gateway.

Overlay Network Design



Flexible Routing and Bridging Overlay Types



- L2 only stretch across the EVPN domain
- Flexible workload placement, host mobility, and optimal traffic forwarding across the BGP EVPN VXLAN fabric.

Flexible Routing and Bridging Overlay Types

Leaf-1

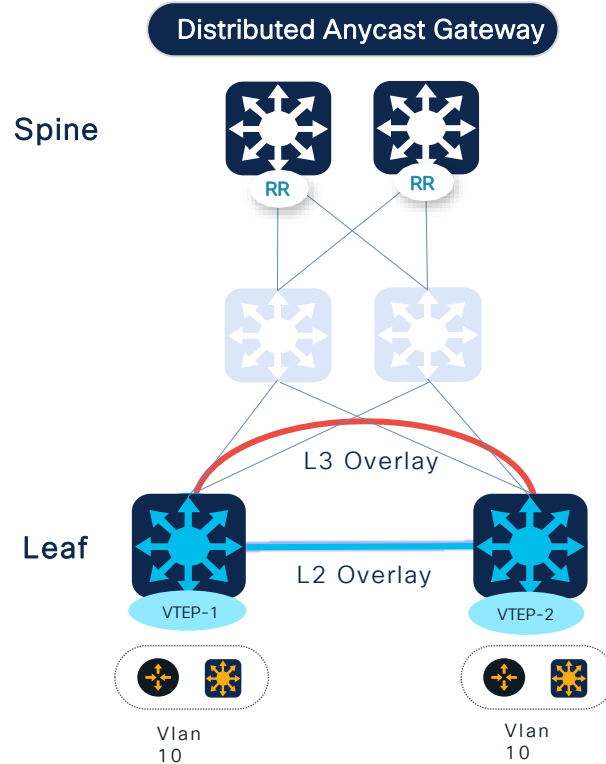
```
l2vpn evpn instance 10 vlan-based
encapsulation vxlan
replication-type ingress
!
vlan configuration 10
member evpn-instance 10 vni 101000
!
vlan configuration 102
member vni 102102
!
interface Vlan102
vrf forwarding S1-EVPN
ip unnumbered Loopback0
!
interface Vlan10
mac-address 0001.0001.0001
vrf forwarding S1-EVPN
ip address 10.10.10.1 255.255.255.0
!
interface nve1
no ip address
source-interface Loopback0
host-reachability protocol bgp
member vni 102102 vrf S1-EVPN
member vni 101000 ingress-replication
```

L2VN

L3VN

DAG

NVE



Leaf-2

```
l2vpn evpn instance 10 vlan-based
encapsulation vxlan
replication-type ingress
!
vlan configuration 10
member evpn-instance 10 vni 101000
!
vlan configuration 102
member vni 102102
!
interface Vlan102
vrf forwarding S1-EVPN
ip unnumbered Loopback0
!
interface Vlan10
mac-address 0001.0001.0001
vrf forwarding S1-EVPN
ip address 10.10.10.1 255.255.255.0
!
interface nve1
no ip address
source-interface Loopback0
host-reachability protocol bgp
member vni 102102 vrf S1-EVPN
member vni 101000 ingress-replication
```

L2VN

L3VN

DAG

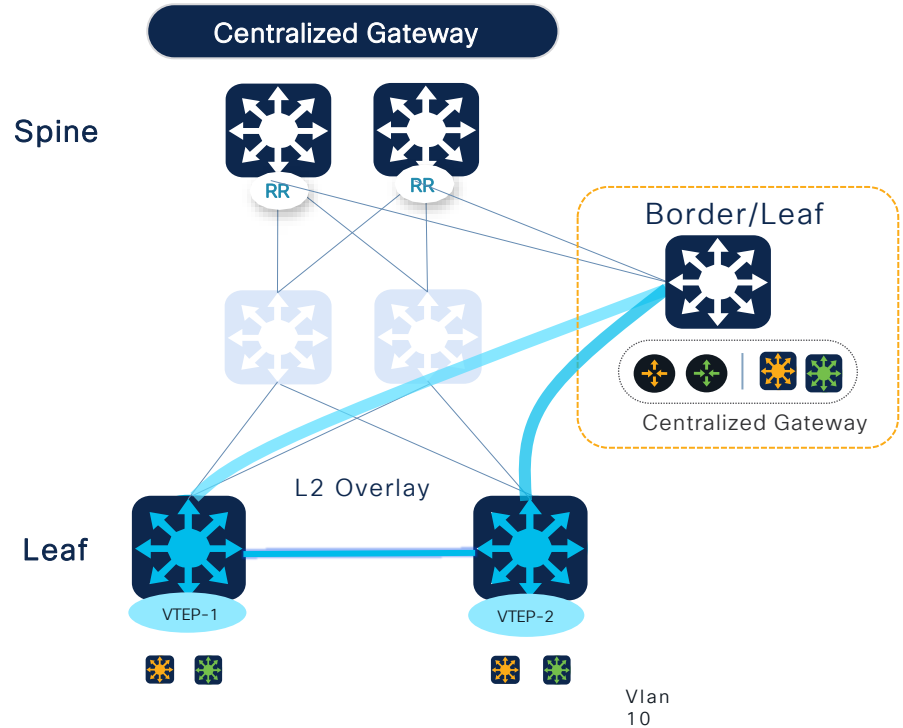
NVE

- The same anycast gateway virtual IP address and MAC address are configured on all VTEPs.
- Flexible workload placement, host mobility, and optimal traffic forwarding across the BGP EVPN VXLAN fabric.

Flexible Routing and Bridging Overlay Types

Inter-Subnet Traffic routed through Centralized Gateway

Intra-Subnet Traffic Bridged through Layer 2 VNI



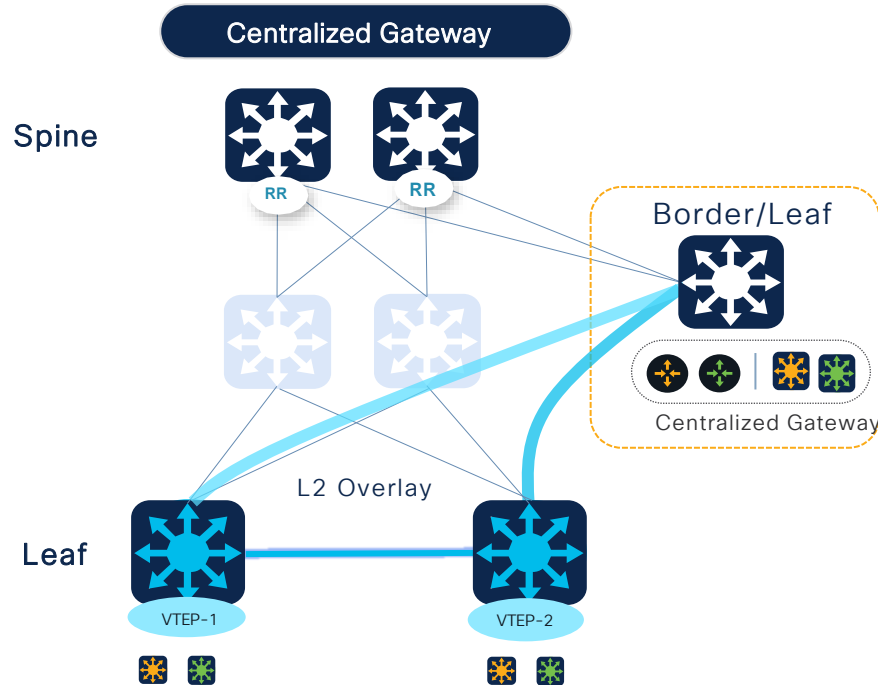
Flexible Routing and Bridging Overlay Types

Centralized Gateway

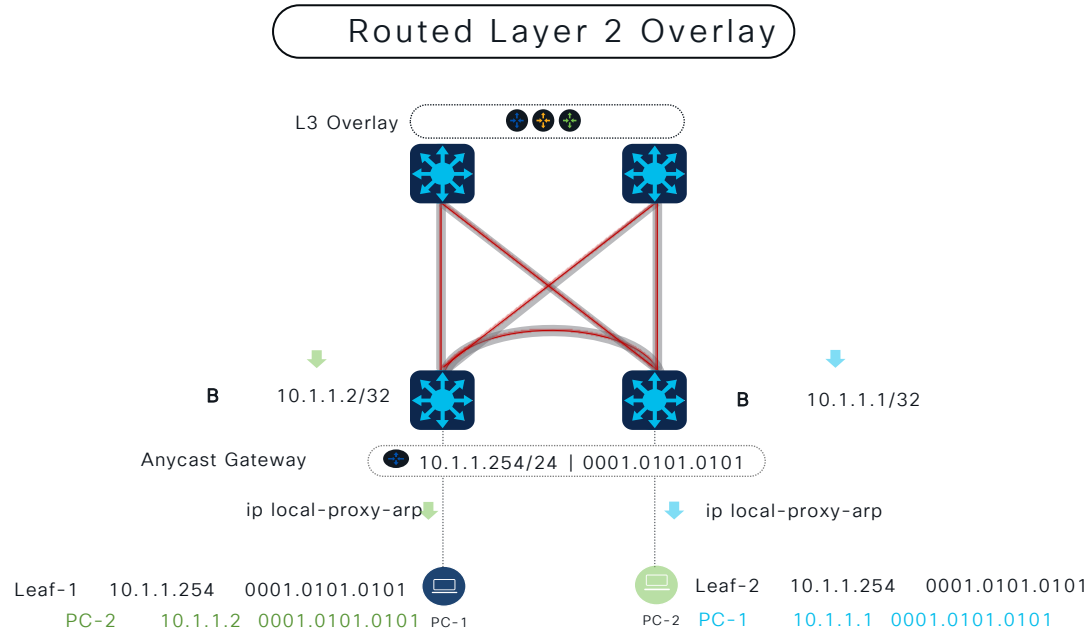
```
l2vpn evpn instance 101 vlan-based
encapsulation vxlan
replication-type ingress
default-gateway advertise enable
!
l2vpn evpn instance 102 vlan-based
encapsulation vxlan
replication-type ingress
default-gateway advertise enable
!
vlan configuration 101
member evpn-instance 101 vni 10101
vlan configuration 102
member evpn-instance 102 vni 10102
```

```
interface Vlan101
vrf forwarding green
ip address 10.1.101.1 255.255.255.0
!
interface Vlan102
vrf forwarding green
ip address 10.1.102.1 255.255.255.0
```

```
interface nve1
no ip address
source-interface Loopback0
host-reachability protocol bgp
member vni 10101 ingress-replication
member vni 10102 ingress-replication
!
```



Proxy – Layer 2 without flooding



Routed
Layer 2

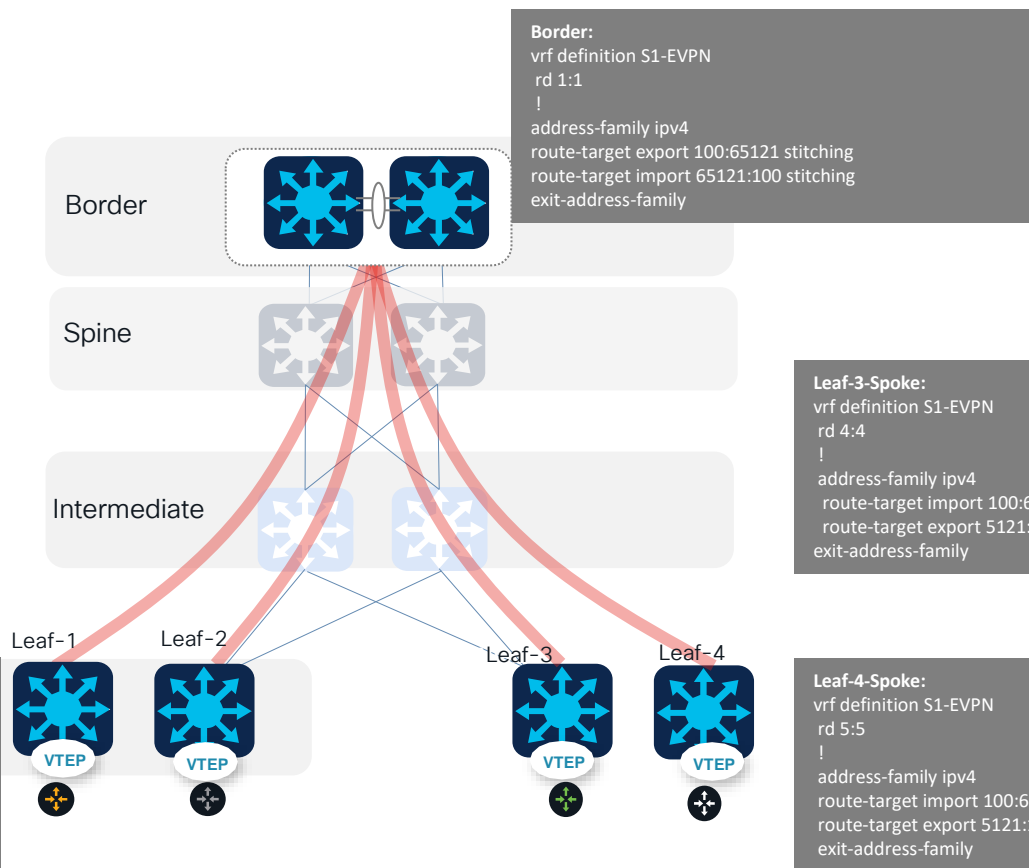
Layer 2 network stretch without Layer 2 overlay extensions (Bridging)

Manual override building dynamic Layer 2 VXLAN tunnel peer across fabric

Dynamic local host discovery and MAC/IP (RT-2) prefix advertisement following IP VRF routing policy

Local proxy function for remote IPv4 ARP and IPv6 ND discovery Leaf local AnyCast MAC address

Overlay Topologies – Hub-n-Spoke



```
Border:
vrf definition S1-EVPN
rd 1:1
!
address-family ipv4
route-target export 100:65121 stitching
route-target import 65121:100 stitching
exit-address-family
```

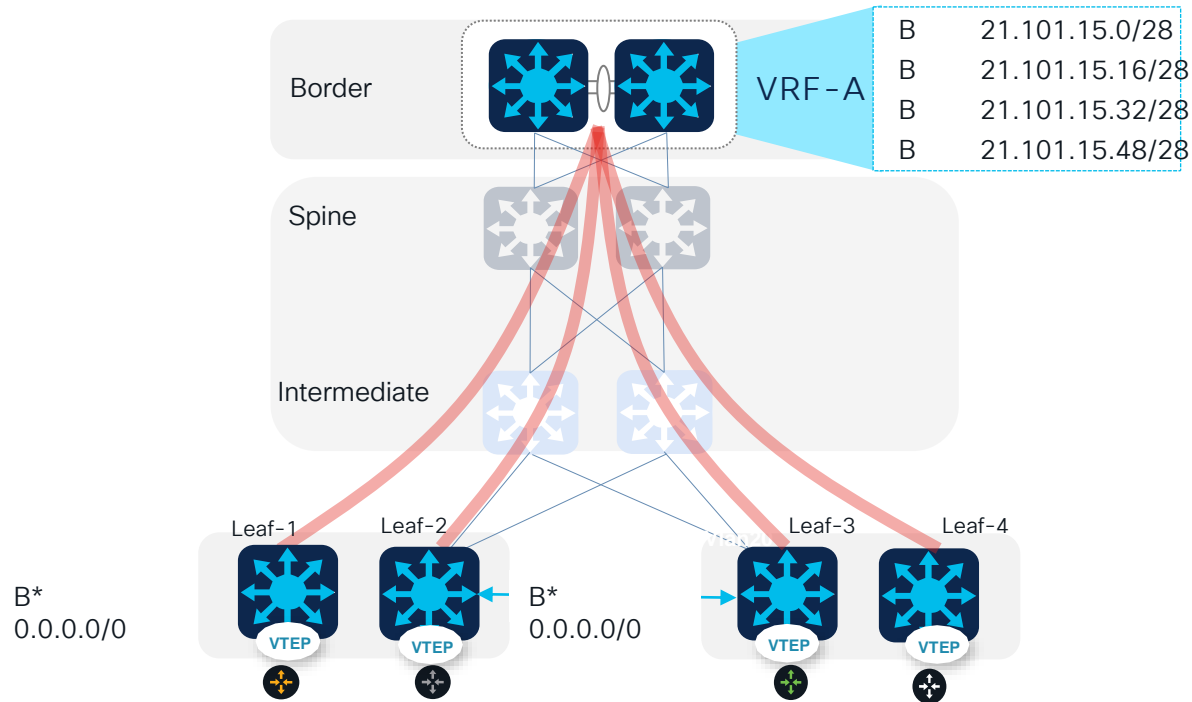
```
Leaf-1-Spoke:
vrf definition S1-EVPN
rd 2:2
!
address-family ipv4
route-target import 100:65121 stitching
route-target export 65121:100 stitching
exit-address-family
```

```
Leaf-3-Spoke:
vrf definition S1-EVPN
rd 4:4
!
address-family ipv4
route-target import 100:65121stitching
route-target export 5121:100 stitching
exit-address-family
```

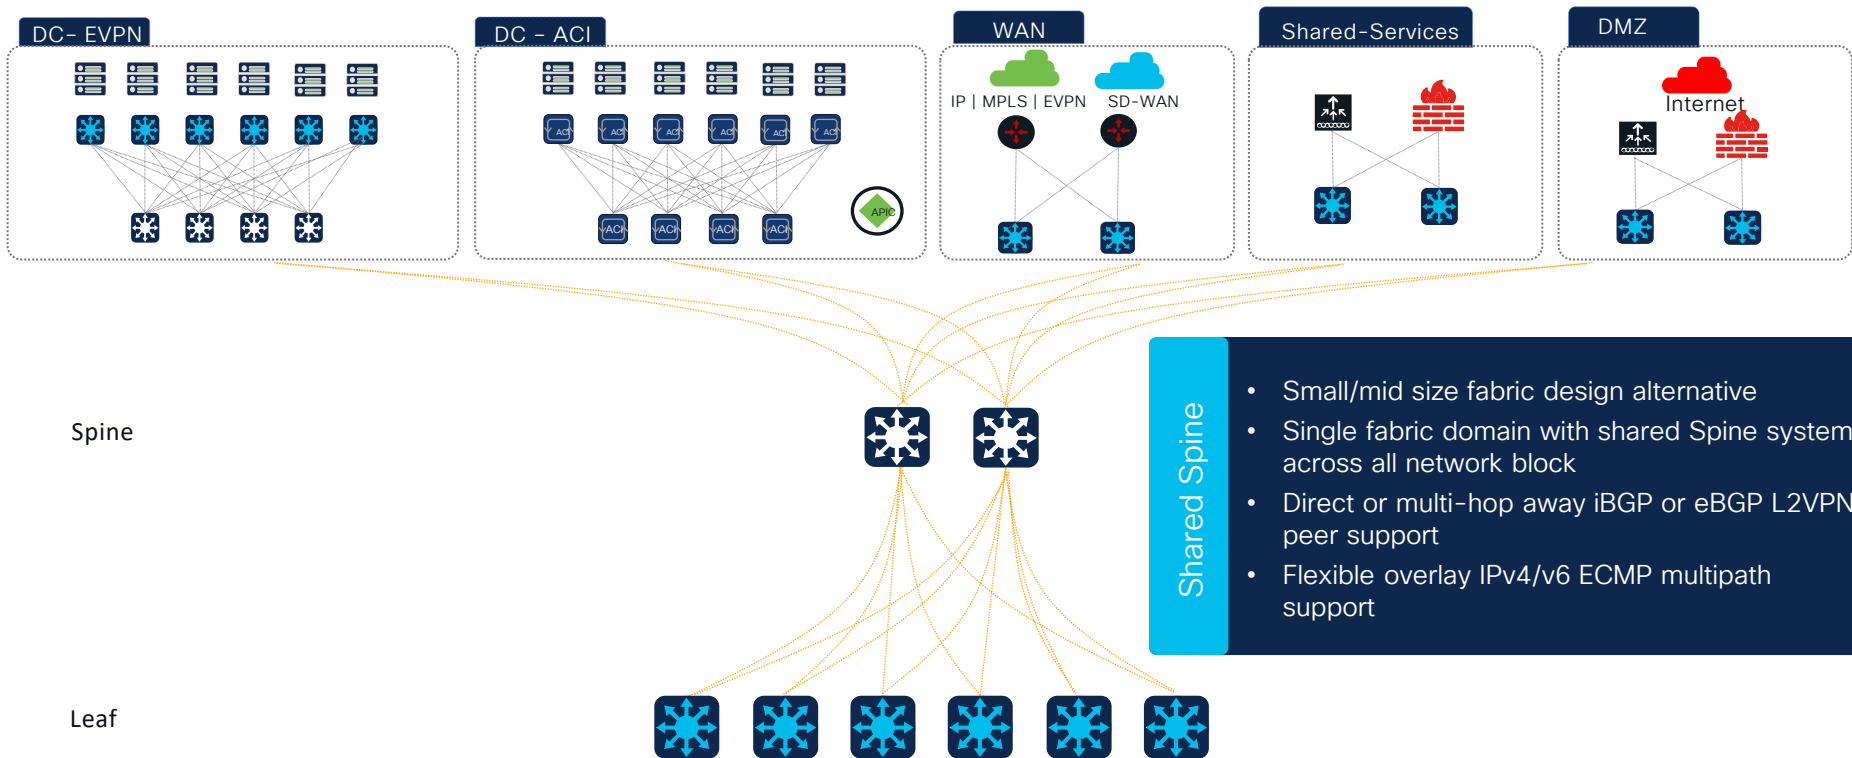
```
Leaf-2-Spoke:
vrf definition S1-EVPN
rd 3:3
!
address-family ipv4
route-target import 100:65121 stitching
route-target export 5121:100 stitching
exit-address-family
```

```
Leaf-4-Spoke:
vrf definition S1-EVPN
rd 5:5
!
address-family ipv4
route-target import 100:65121 stitching
route-target export 5121:100 stitching
exit-address-family
```

Overlay Topologies - Hub-n-Spoke

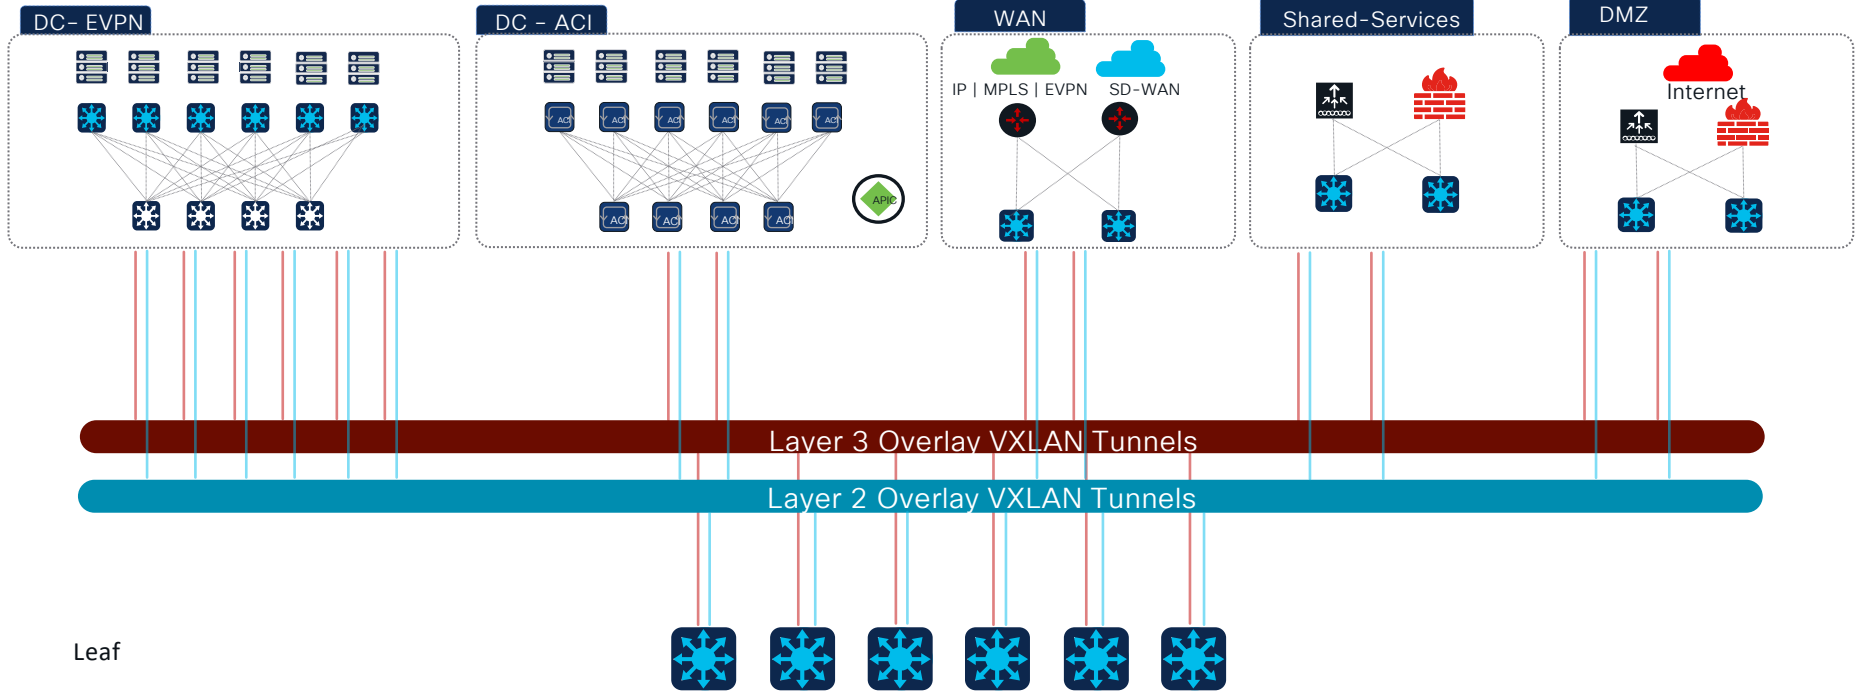


Single Cluster Fabric Architecture



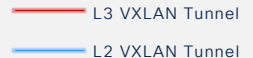
- Shared Spine
- Small/mid size fabric design alternative
 - Single fabric domain with shared Spine system across all network block
 - Direct or multi-hop away iBGP or eBGP L2VPN peer support
 - Flexible overlay IPv4/v6 ECMP multipath support

Non-Hierarchical Fabric Design

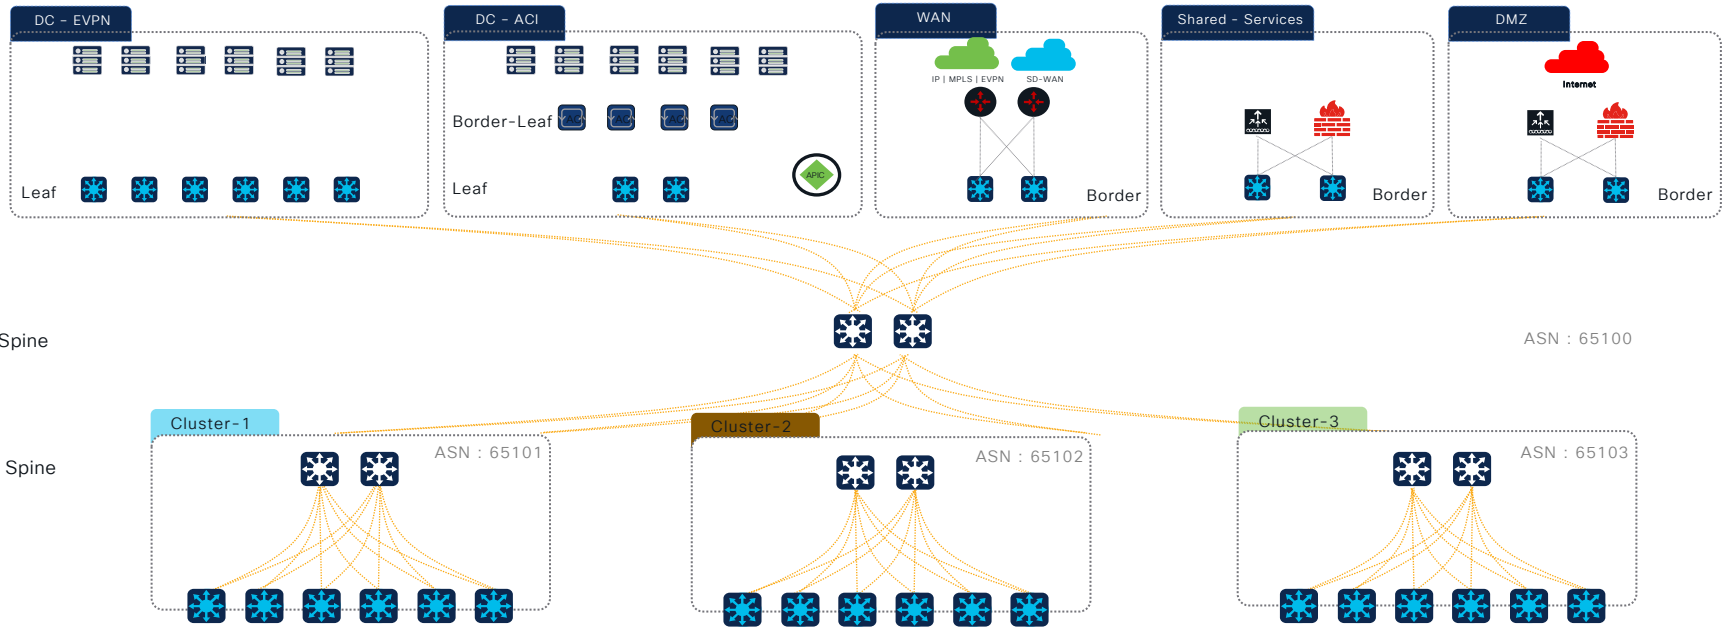


Non-Hierarchical Fabric

- Non-hierarchical dynamic overlay VXLAN tunnels
- Layer 2 / 3 overlay topologies based on route-target policies
- Linear VN & Leaf growth may impact overall fabric domain scale
- Limited Layer 2 flood control support



Multi Cluster Fabric Architecture



Mid to large size fabric design alternative

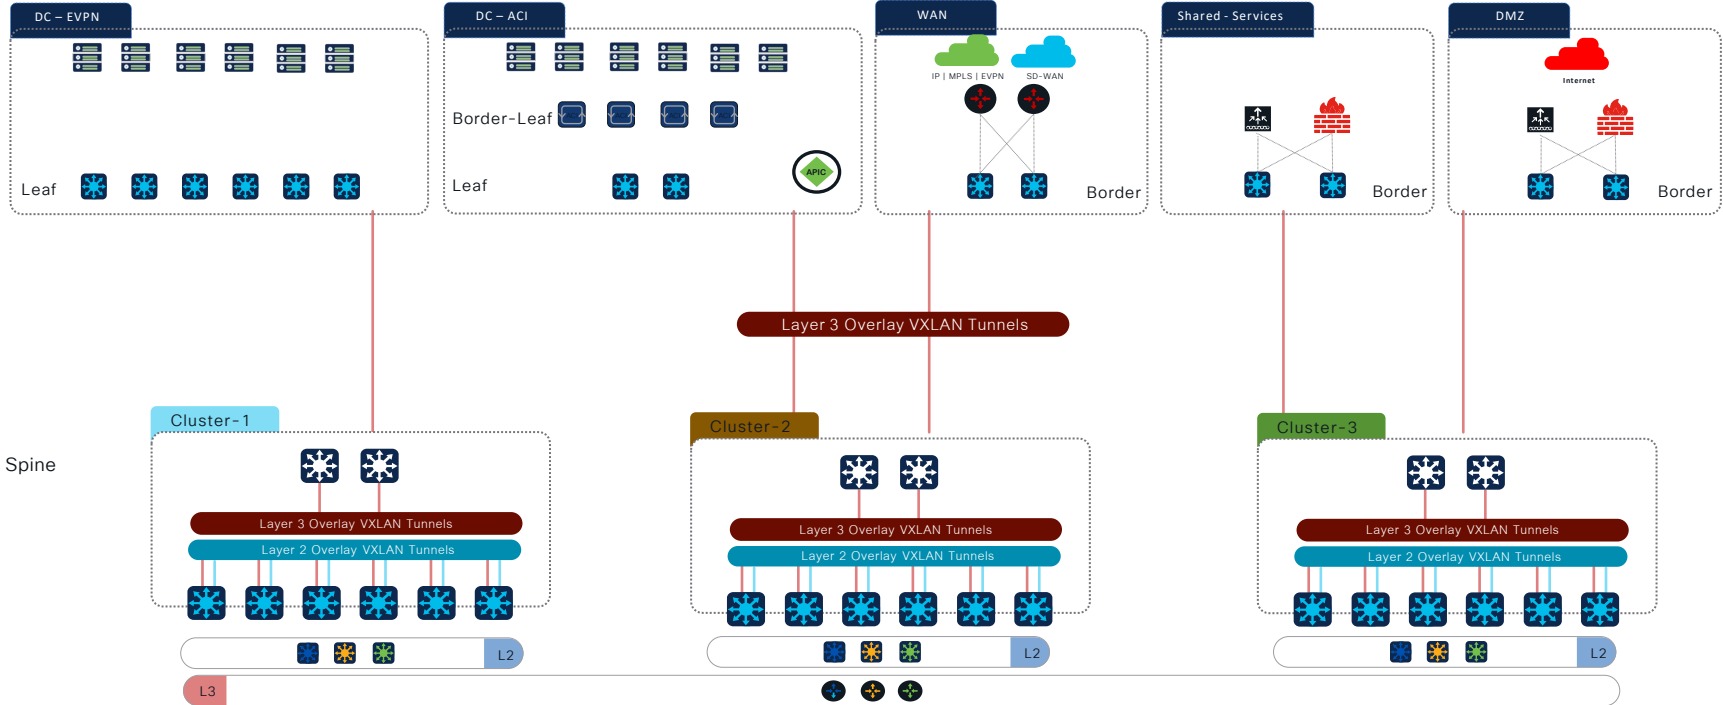
Reduce fault-domain with distributed RR & RRC clusters for high scale fabric

Structured fabric control-plane architecture between distributed Spine and Super-Spine

Optimized fabric and system scale, performance and resiliency

Distributed Spine

Scalable Multi Cluster Fabric Architecture

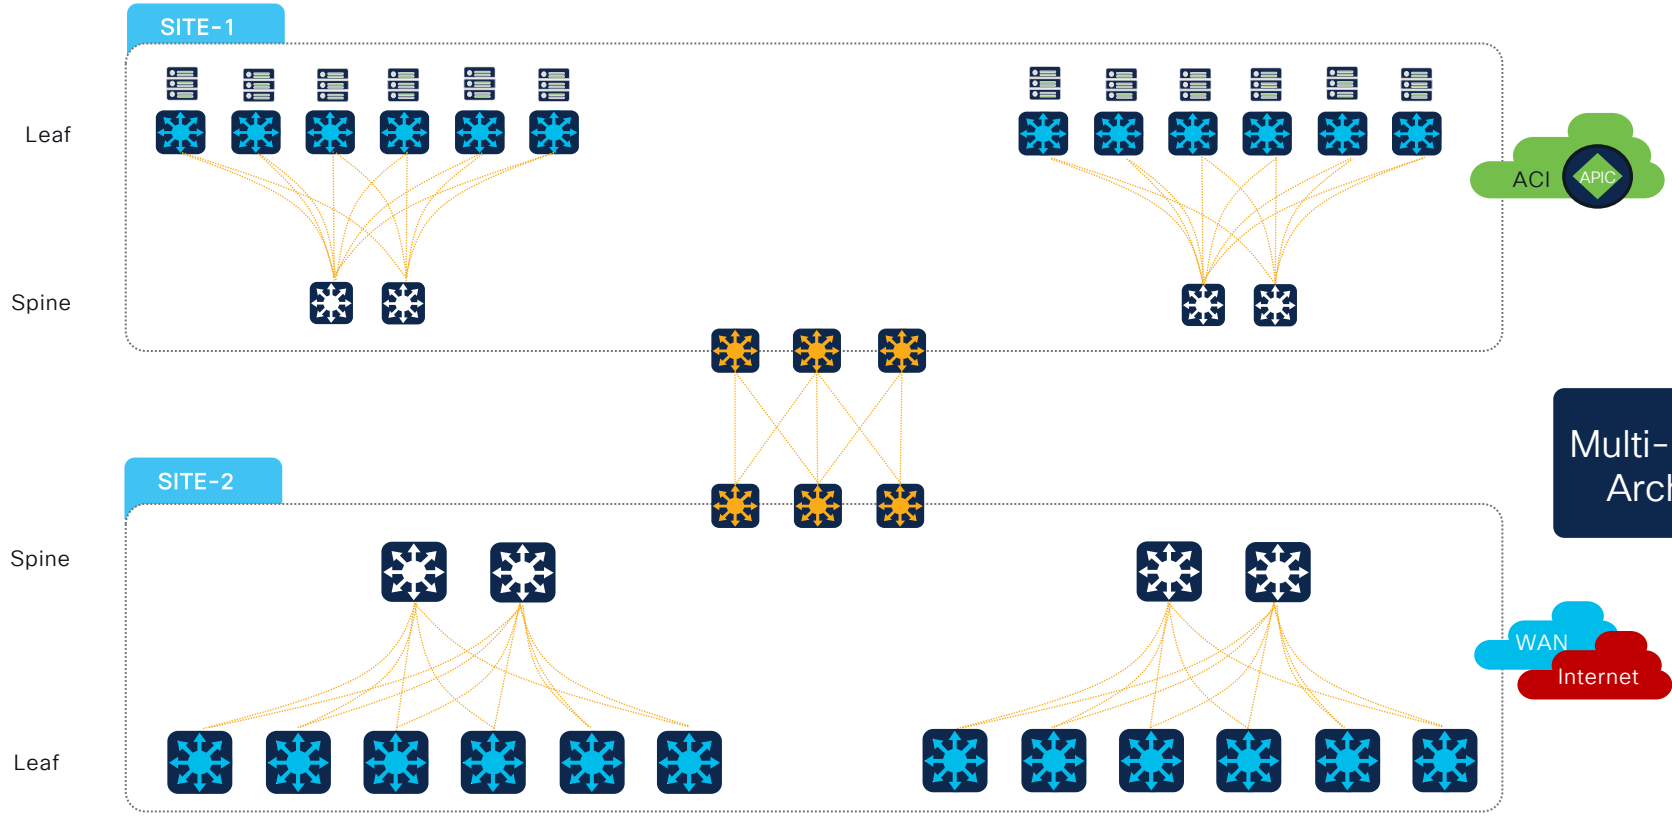


Increase fabric domain scale with hierarchical dynamic overlay VXLAN tunnels per fabric cluster

Consistent Layer 2 domain scale size as traditional non-fabric networks

Scalable overlay routing with per-VN prefix summarization and re-origination by each Border Spine cluster

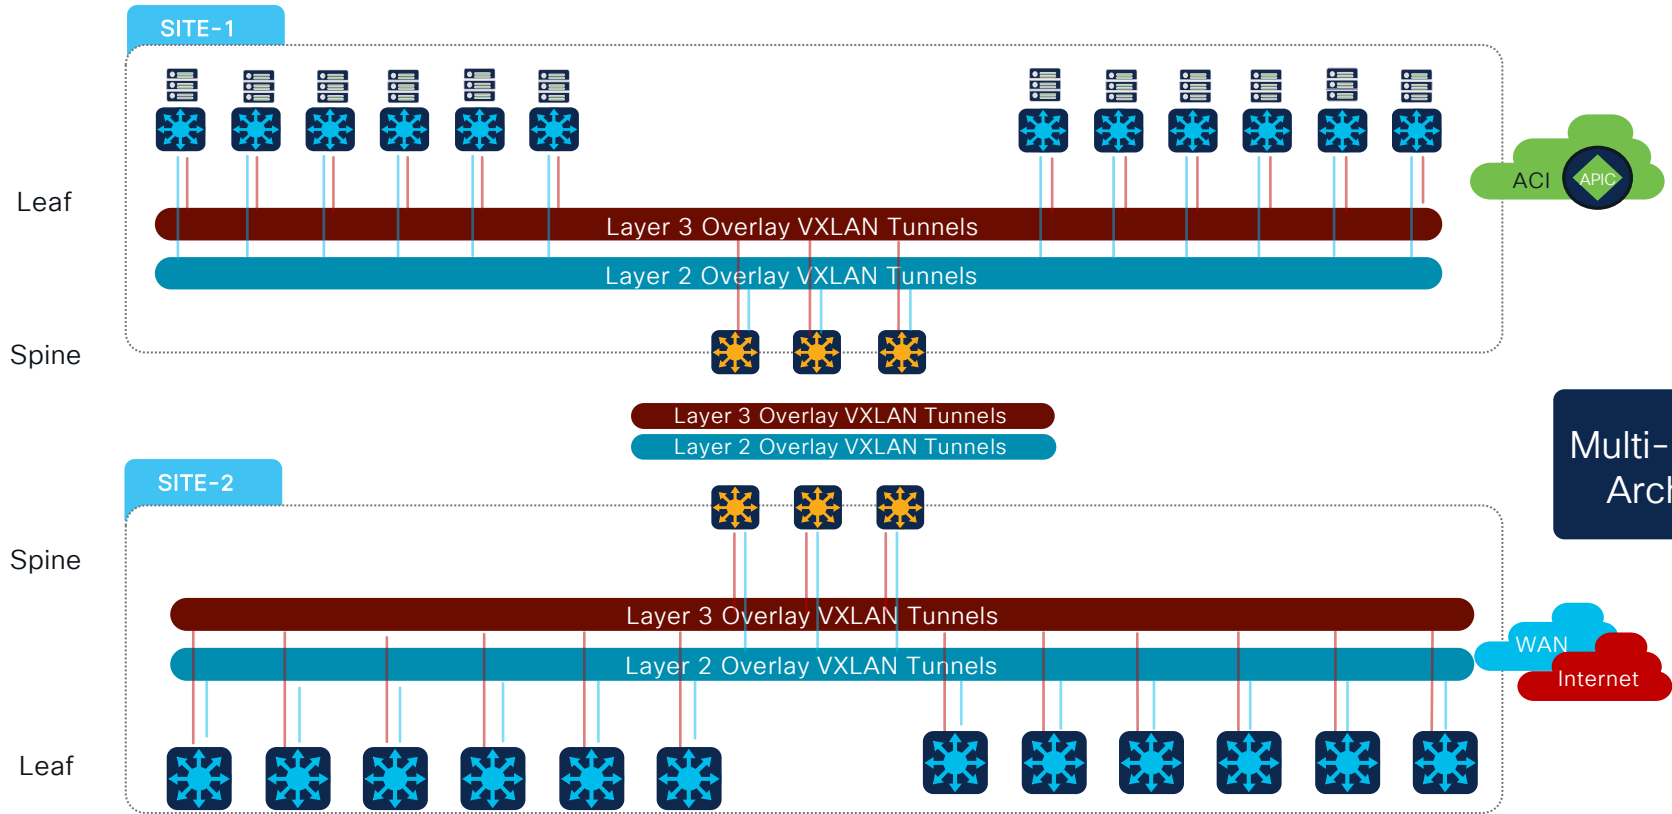
End-to-End Unicast IPv4/IPv6 support. Layer 2 Mobility and overlay Multicast limited per cluster



Multi-Site Fabric Architecture

Multisite Fabric

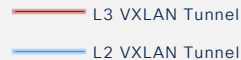
- Well-structured fabric overlay solution for large EN/DC networks
- Single fabric site representation enables scalable overlay network hierarchy
- Granular control of Layer 2 and Layer 3 overlay flood and routing control
- Seamless integration between Catalyst and Nexus 9K (Border-GW)



Multi-Site Fabric Architecture

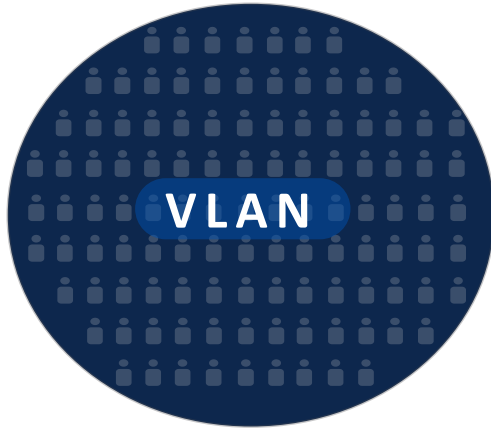
Multisite Fabric

- Well-structured fabric overlay solution for large EN/DC networks
- Single fabric site representation enables scalable overlay network hierarchy
- Granular control of Layer 2 and Layer 3 overlay flood and routing control
- Seamless integration between Catalyst and Nexus 9K (Border-GW)



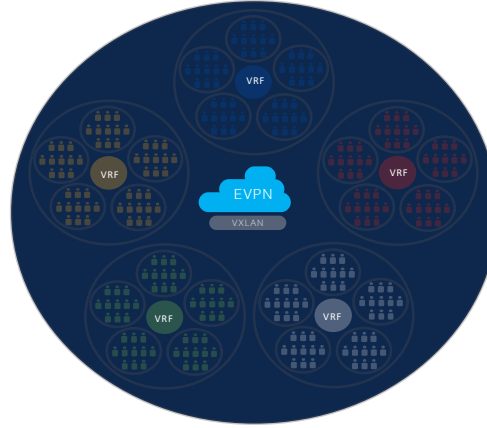
Segmentation

Logical Local Grouping



Macro-Segmentation

Extended Group with EVPN

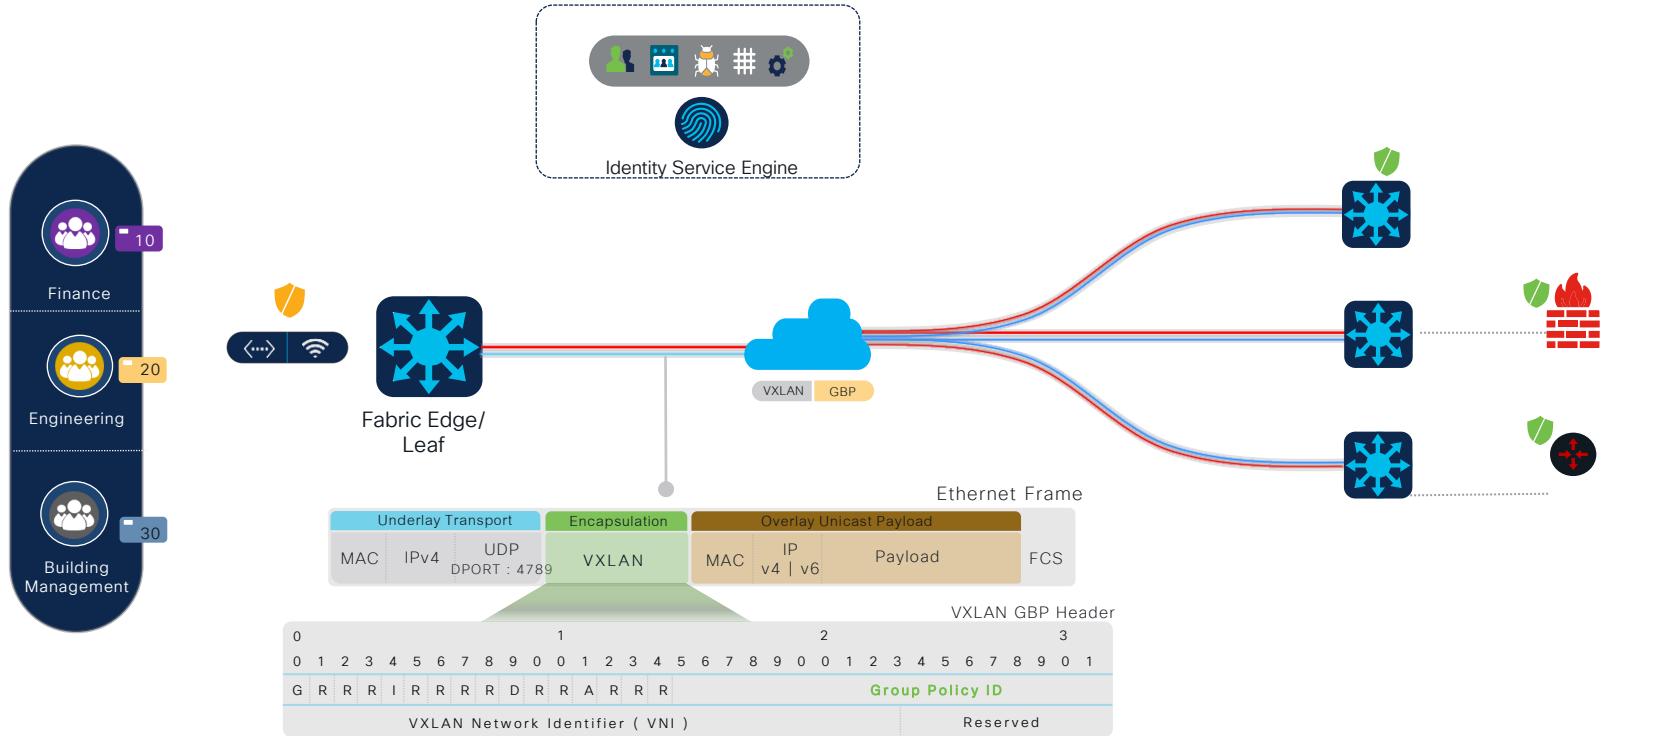


Micro-Segmentation

Policy-Plane enforced fabric



Micro Segmentation

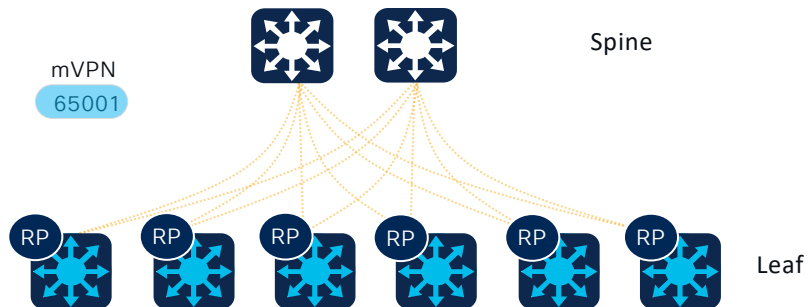


 Classification
 Enforcement

Multicast over VXLAN

Layer 3 Overlay

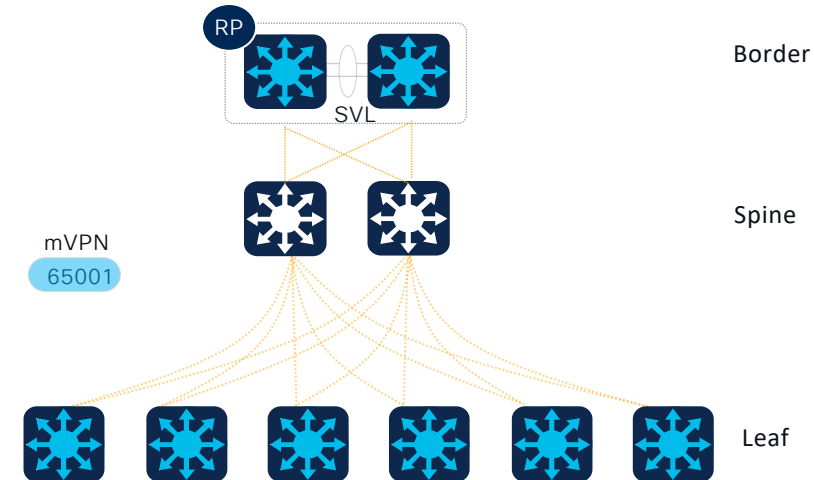
Distributed Anycast RP



VRF	RP	IP	MDT
Blue	Anycast	Lo1: 10.1.1.101	239.1.1.101
Yellow	Anycast	Lo2: 10.2.1.101	239.2.1.101
Green	Anycast	Lo3: 10.3.1.101	239.3.1.101

Layer 3 Overlay

Fabric Border RP



VRF	RP	IP	MDT
Blue	Anycast	Lo1: 10.1.1.101	239.1.1.101
Yellow	Anycast	Lo2: 10.2.1.101	239.2.1.101
Green	Anycast	Lo3: 10.3.1.101	239.3.1.101

Overlay RP Design

Standard-based Multicast overlay network design support

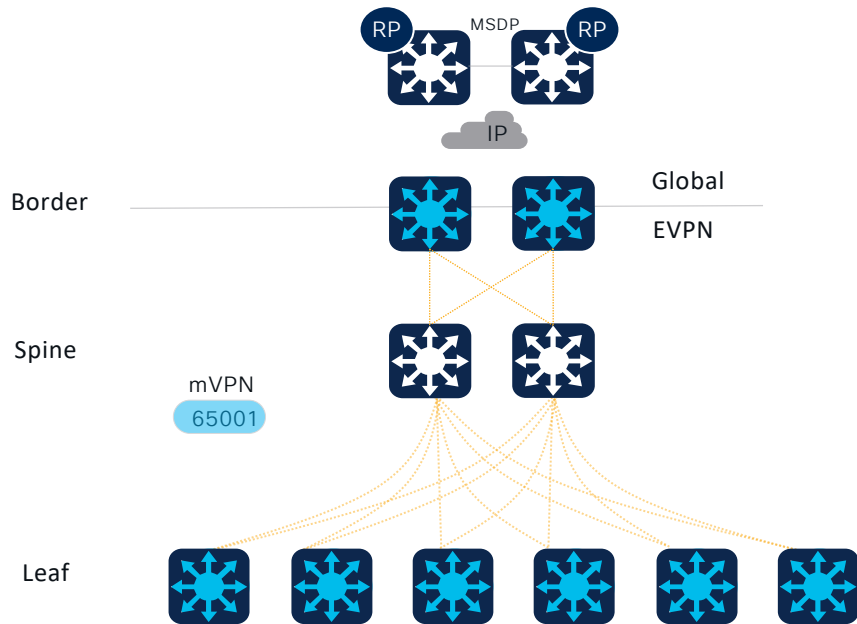
Flexible Multicast RP design alternatives to address scale, performance, resiliency

AnyCast RP at Leaf or Border enables distributed Multicast administrative domains supporting unified routing policies

Unified Multicast RP between Underlay and Overlay RP supporting existing brownfield deployment models

Layer 3 Overlay

External Domain RP



VRF	RP	IP	MDT
Blue	Anycast	Lo1: 10.1.1.101	239.1.1.101
Yellow	Anycast	Lo2: 10.2.1.101	239.2.1.101
Green	Anycast	Lo3: 10.3.1.101	239.3.1.101

Overlay RP Design

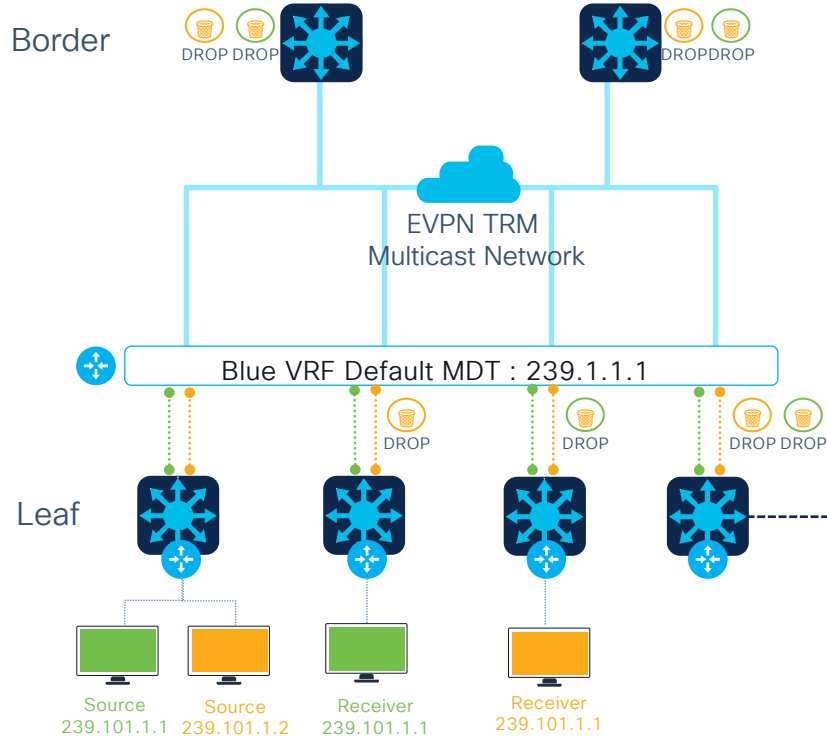
Standard-based Multicast overlay network design support

Flexible Multicast RP design alternatives to address scale, performance, resiliency

AnyCast RP at Leaf or Border enables distributed Multicast administrative domains supporting unified routing policies

Unified Multicast RP between Underlay and Overlay RP supporting existing brownfield deployment models

TRM Default MDT

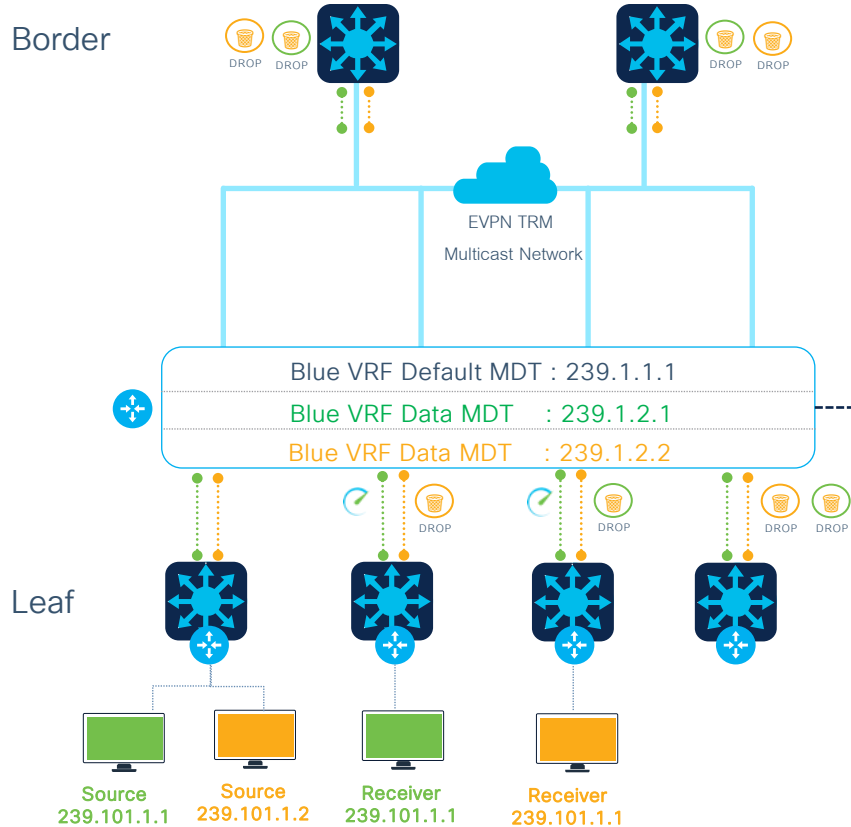


```
vrf definition S1-EVPN
rd 10:10
!
address-family ipv4
mdt auto-discovery vxlan inter-as
mdt default vxlan 239.1.1.1 → MDT Default
mdt overlay use-bgp spt-only
```

Challenges

- Non-selective overlay Multicast replication
- Inessential core network bandwidth utilization
- Redundant system resources utilization
- Limited scale for dense network environment

TRM Data MDT




```
vrf definition S1-EVPN
rd 10:10
!
address-family ipv4
mdt auto-discovery vxlan inter-as
mdt default vxlan 239.1.1.1
mdt data vxlan 239.1.2.0 0.0.0.255 → MDT Data
mdt data threshold 1
mdt overlay use-bgp spt-only
```

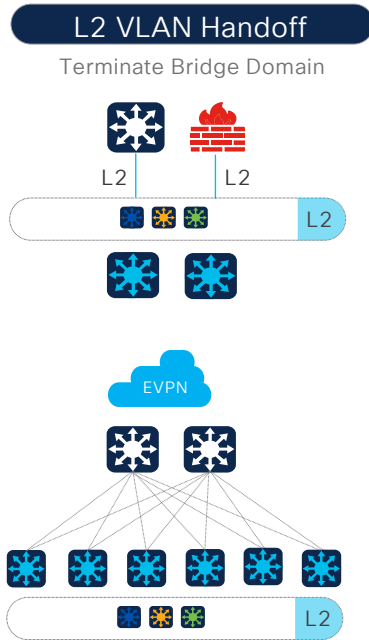
Key Benefits

- Stateful L2 Multicast Overlay network
- Industry-standard based control-plane
- Applicable to Centralized Gateway or Cross-Connect Overlay networks
- Scale. Performance. Security.

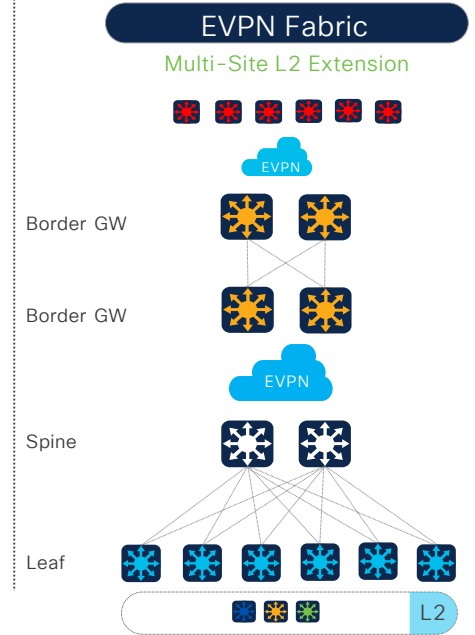
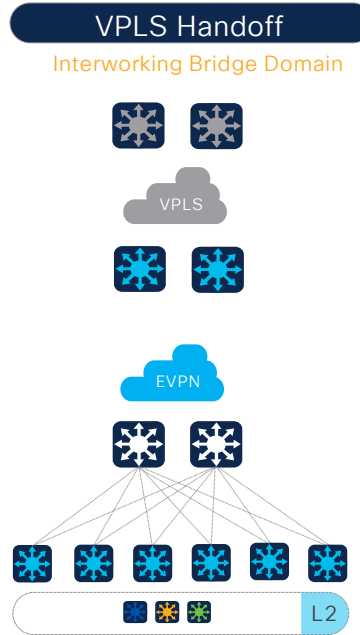
EVPN Fabric Interworking



Layer – 2 Handoffs Options



PE



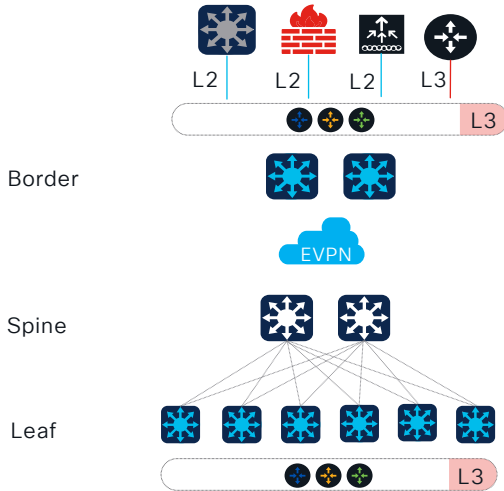
Seamless
Layer 2
Handoff

Multiple end-to-end seamless Layer 2 extensions supports across fabric and beyond
Terminate L2 overlays and perform simple Layer 2 trunk handoff to non-fabric devices, i.e., Firewalls
Integrated EVPN Border and VPLS PE function to extend multi-domain L2 for seamless migrations
Extendable Layer 2 EVPN domains with highly scalable Catalyst and Nexus 9000 Multisite Border Gateway

Layer – 3 Handoffs Options

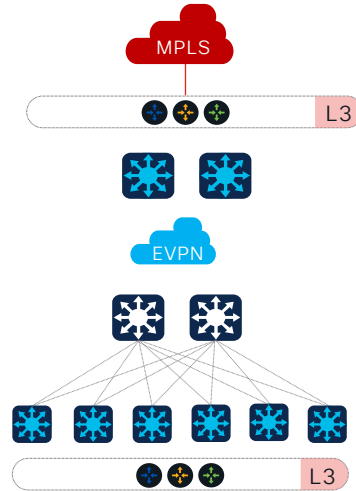
IP VRF Handoff

Terminating Routing Domain



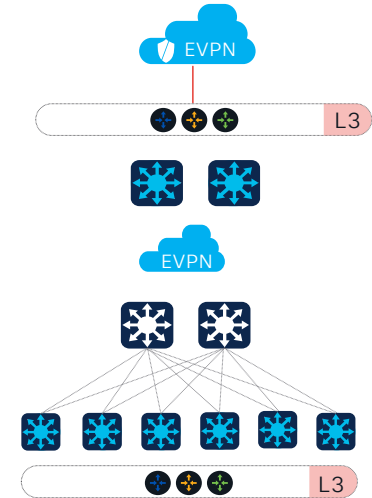
MPLS VPN Handoff

Interworking Overlay Domain



EVPN Fabric


Re-originating Fabric Domain



Integrated
Extranet

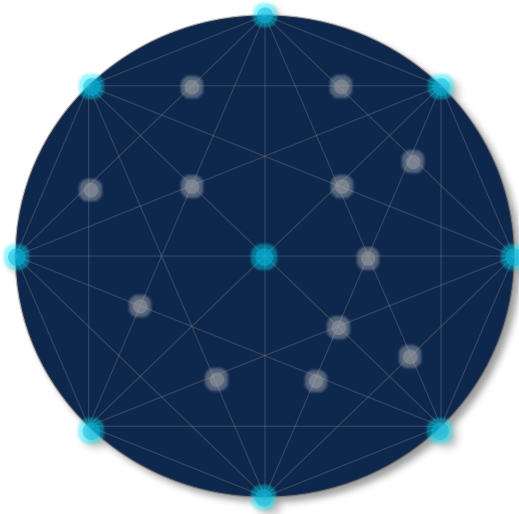
- Transparent EVPN handoff to Layer 2 or Layer 3 to traditional underlay segmented networks
- Seamless multi-domain interworking at Border – IP, MPLS VPN, EoMPLS/VPLS, SD-WAN, etc.
- Extendable Unicast | Multicast support for IPv4 and IPv6 between EVPN to external domain
- Dedicated or collapsed system-role – Leaf, Spine, Border, Border-Leaf, Border-Spine

Fabric Deployment Options

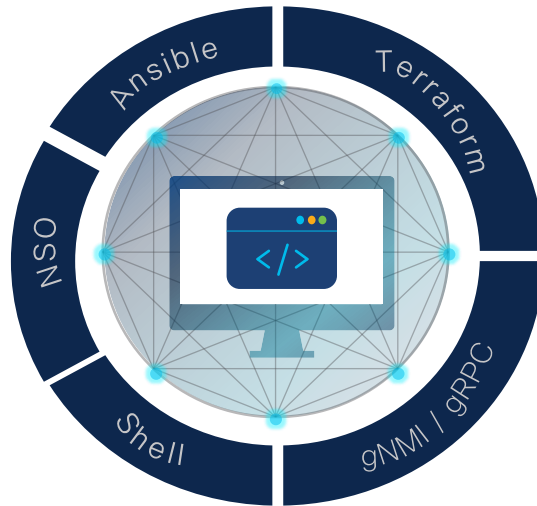


Cisco Enterprise BGP EVPN Solution

Do-It-Yourself

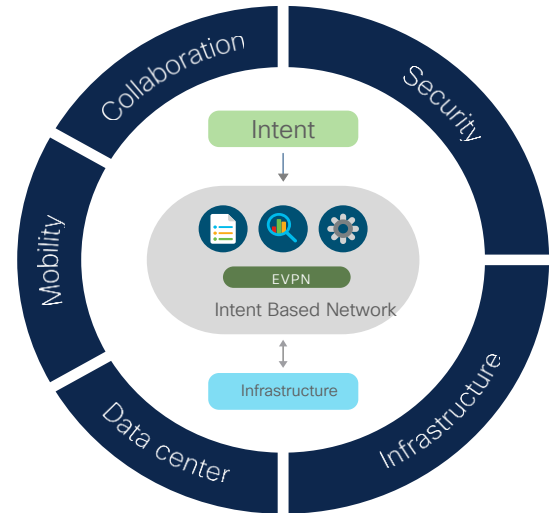


Programmable



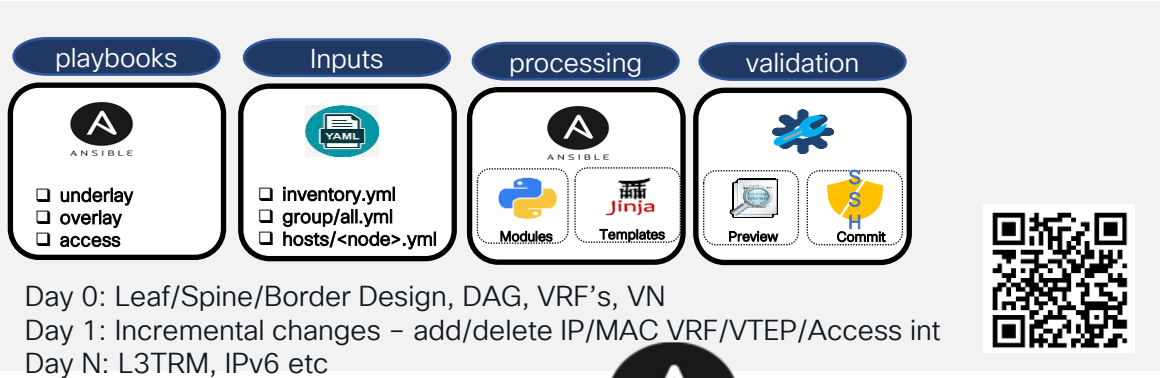
Intent-Based

Q1CY25

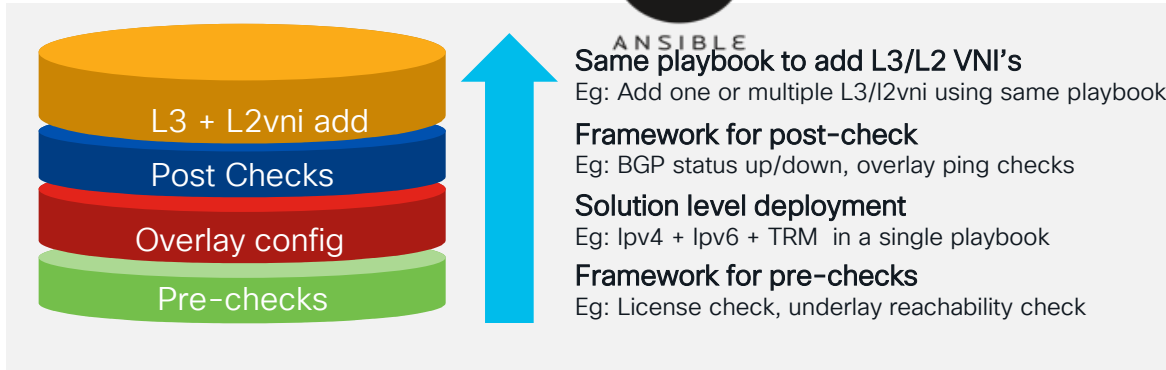


BGP EVPN Automation – Ansible & Terraform

Feature + Solution Playbooks



Solution Playbooks

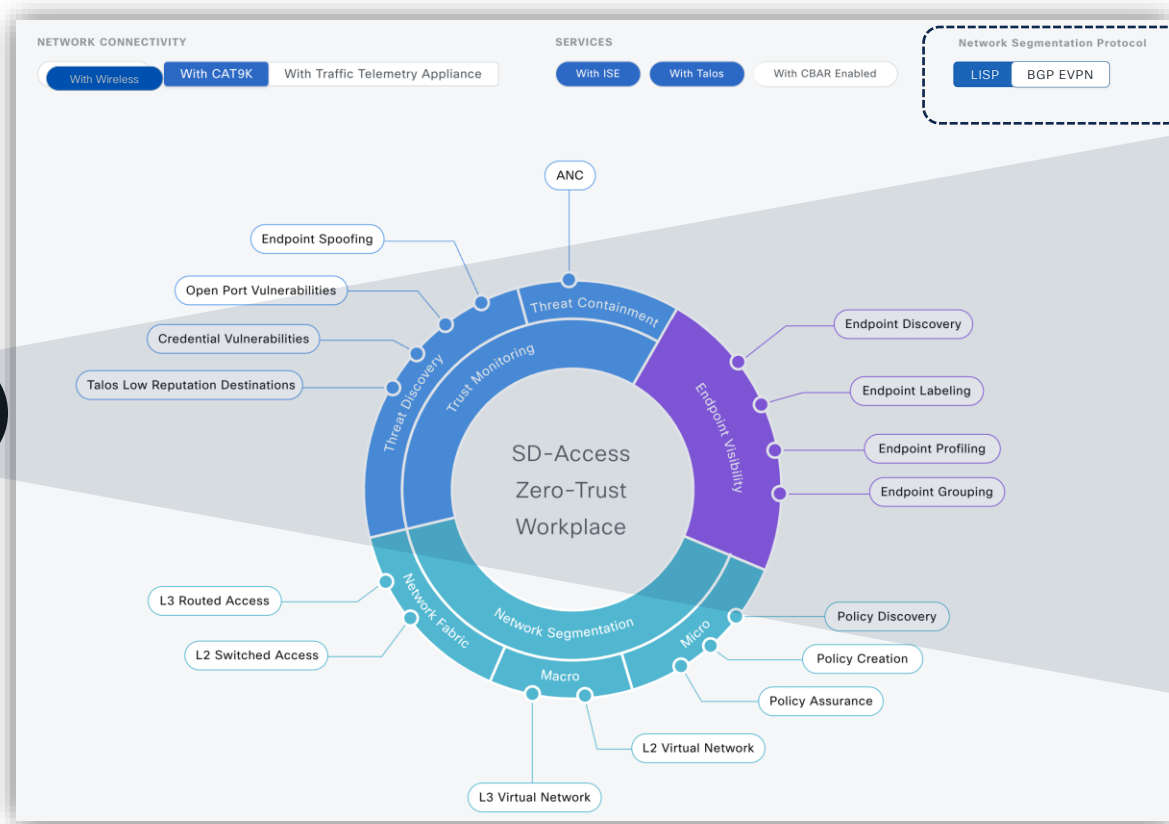


EVPN Terraform Provider

```

1 # EVPN Settings
2 resource "ciscoevpn_evpe" "evpe" {
3   roles = ["leafs"]
4   replication_type = "static"
5   mac_duplication_limit = 20
6   mac_duplication_time = 10
7   ip_duplication_limit = 20
8   ip_duplication_time = 10
9   router_id = local.lookup_interface
10  default_gateway = "advertise"
11  logging_peer_state = true
12  route_target_auto = "vni"
13 }
14
15 # EVPN Multicast
16 resource "ciscoevpn_evpn_instance" "instance_101" {
17   roles = ["leafs"]
18   instance_id = 101
19   vlan_based = true
20   encapsulation = "vxlan"
21   replication_type = "static"
22   rt = "101:101"
23   rt = "101:101"
24   rt_type = "both"
25   ip_learning = true
26   default_gateway_advertise = false
27   re_originate = "route-type5"
28 }
29
30 # Network Virtual Interface
31 resource "ciscoevpn_vni" "leafs" {
32   vni_id = 1
33   ciscoevpn_vlan_vlan_101,
34   ciscoevpn_vlan_vlan_102,
35   ciscoevpn_vlan_vlan_103,
36   ciscoevpn_vlan_vlan_104,
37   ciscoevpn_vlan_vlan_105,
38 }
39
40 roles = ["leafs"]
41 source_interface = local.lookup_interface
42 vni = [
43   "${ciscoevpn_vrf_green.name} = "${ciscoevpn_vlan_vlan_103.vni}",
44   "${ciscoevpn_vrf_blue.name} = "${ciscoevpn_vlan_vlan_104.vni}",
45   "${ciscoevpn_vrf_red.name} = "${ciscoevpn_vlan_vlan_105.vni}"
46 ]
    
```

Choose Your Fabric Control-Plane (LISP or BGP EVPN)



IBN based workflows

Automation

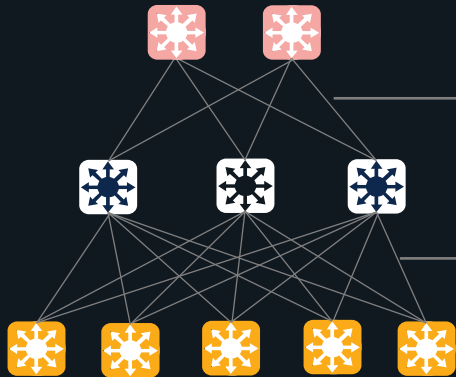
Assurance

Macro/Micro Segmentation

Single Data-Plane Vxlan

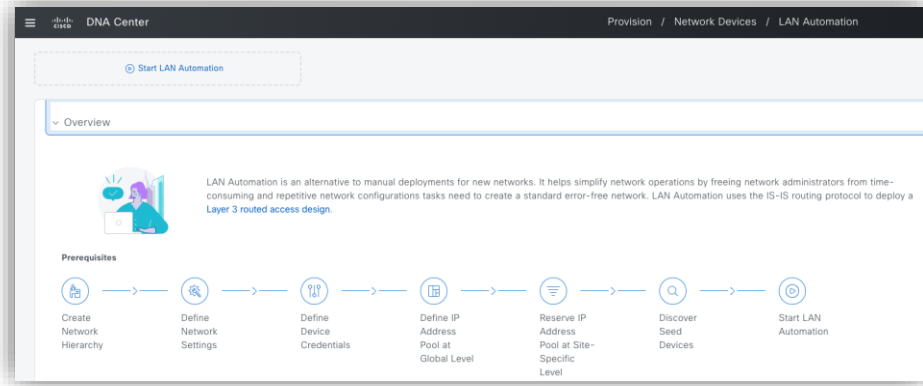
Seamless experience irrespective of choice of protocol

Underlay Provisioning



Build your Underlay

LAN Automation



DIY



DNAC Template Editor



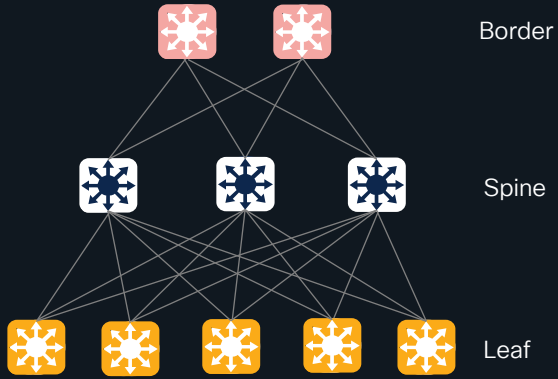
Programmable



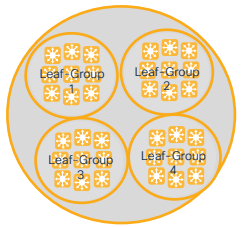
CLI

EIGRP | OSPF | ISIS


Fabric Roles & Fabric Groups



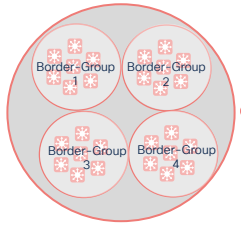
Fabric Groups




Leaf Group



Spine Group



Border Group



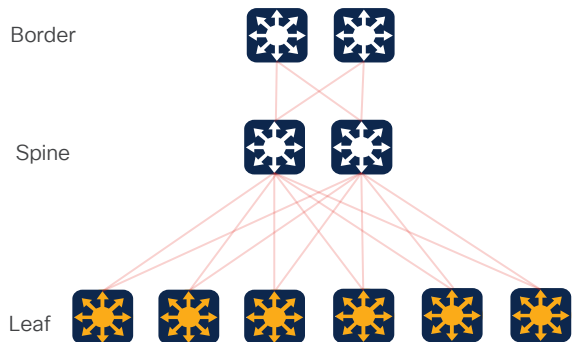
Border Spine Group

Spine Group: 1 per fabric site

SD-Access with BGP EVPN: Supported Underlay's

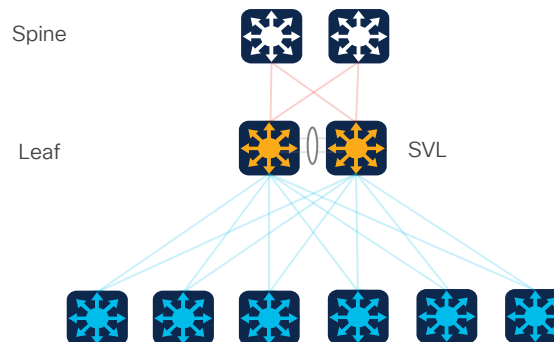


Layer 3 Access



- Leaf Layer – Access
- Spine/RR – Direct | Multi-hop
- Overlay IP gateway
- L2 | L3 Overlay support | Multicast

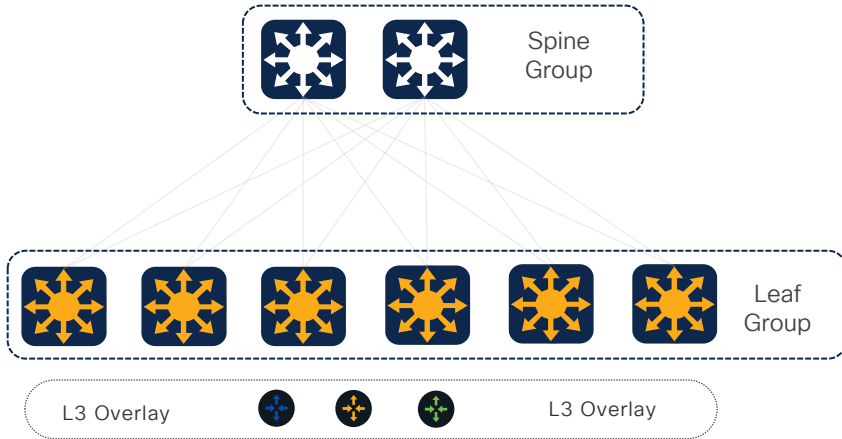
Cisco StackWise-Virtual



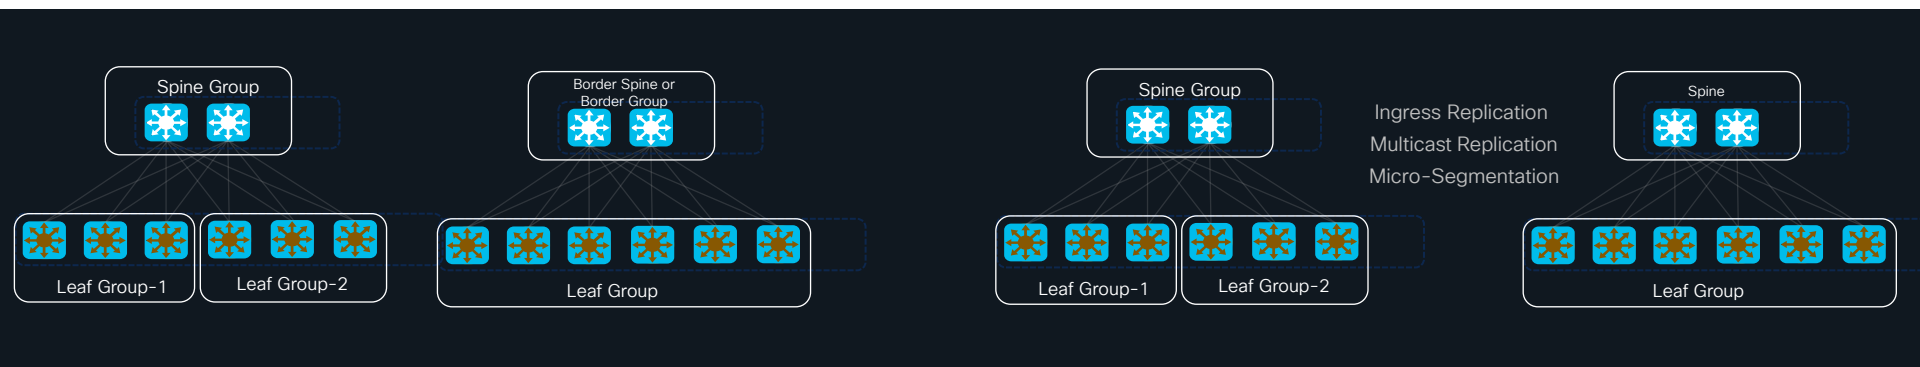
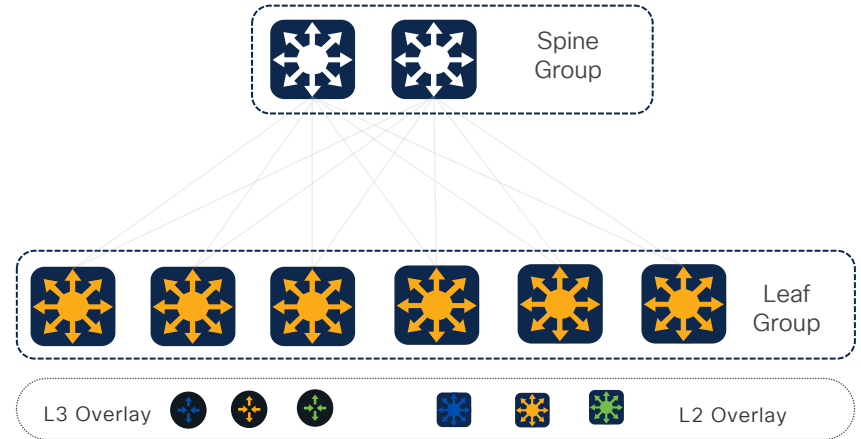
- Leaf Layer – Access
- Spine/RR – Direct | Multi-hop
- Overlay IP gateway
- L2 | L3 Overlay support | Multicast

SD-Access with BGP EVPN: Overlay Topologies

Layer 3 Overlays

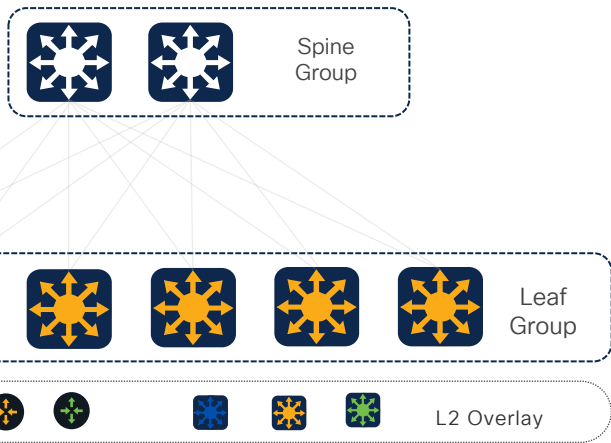


Distributed Anycast Gateway

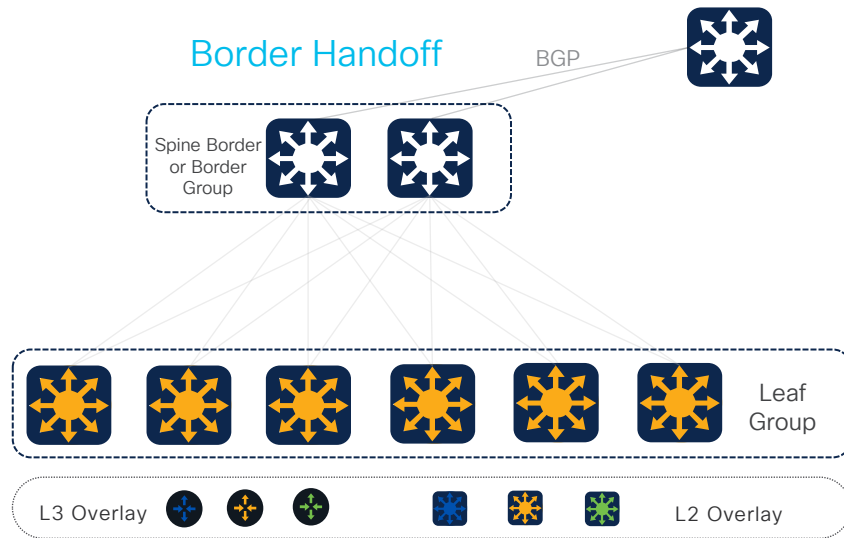


SD-Access with BGP EVPN: TRM & Border Handoff

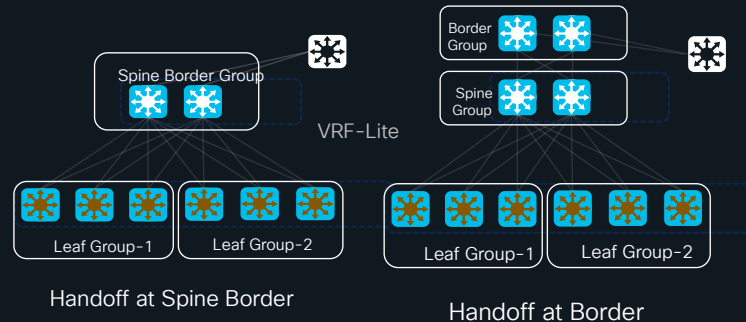
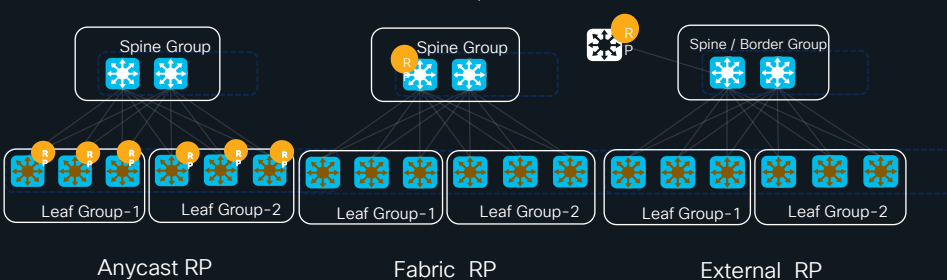
Tenant Routed Multicast



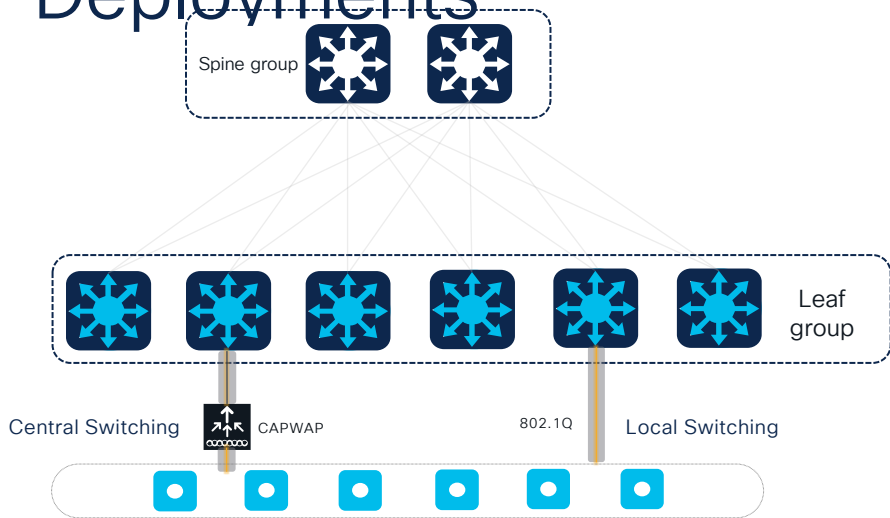
Border Handoff



Default MDT | Data MDT

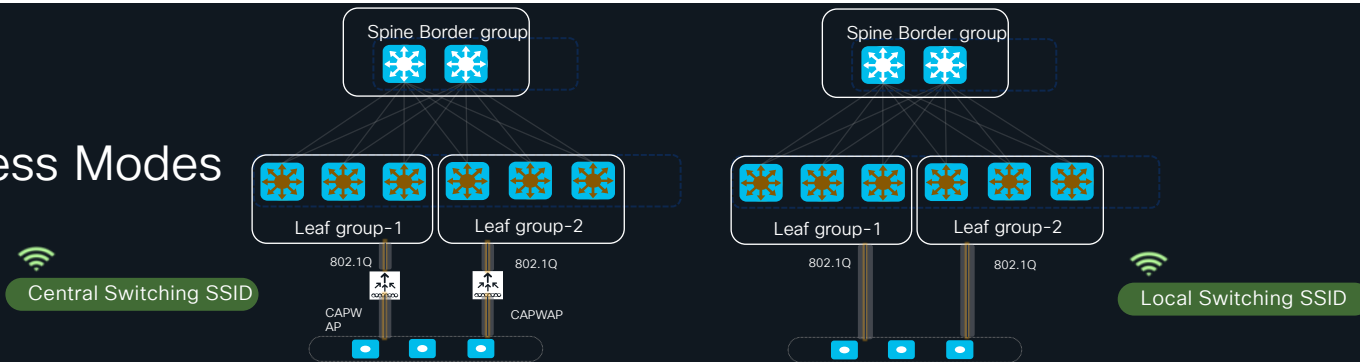


SD-Access with BGP EVPN: Wireless Deployments



- Transparent Wireless support in fabric. Intact WLC and AP communication in Underlay
- Flexible SSID alternatives – Central Switching, Local Switching, Central + Local Switching
- Fabric boundary initiates from Wireless Client IP gateway. Flex Local seamless roaming up to 300 AP.
- Consistent Wired and Wireless network access control policy enforcement

Supported Wireless Modes

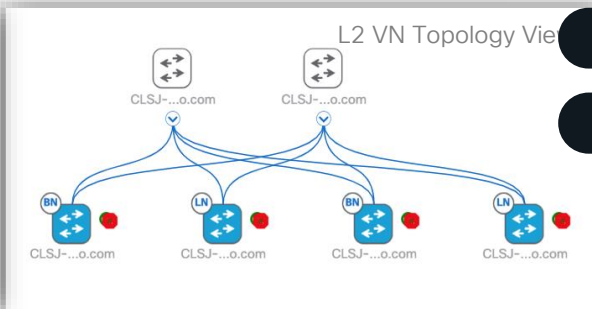
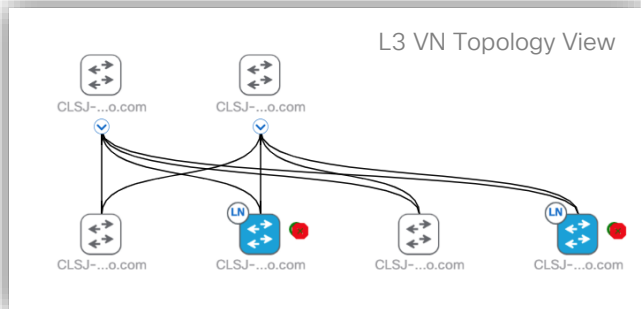
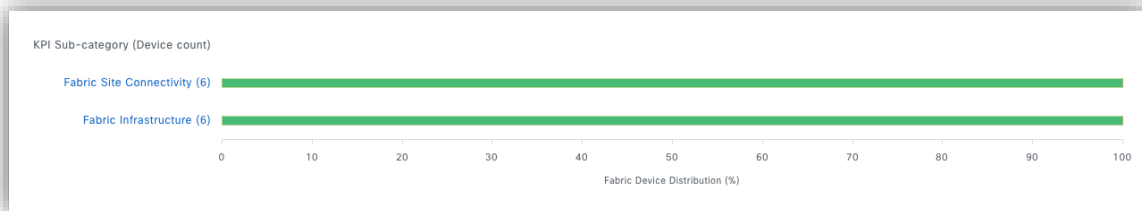


SD-Access with BGP EVPN: Assurance

Cisco Catalyst Center



Device Mgmt | 360° | Topology | Fabric Infrastructure | Fabric Site Connectivity | Base Automation | Network Profiles | SWIM



VNI Status

NVE peer detection

AAA Server Status

BGP peer monitoring to Spine

BGP peer monitoring to Non Fabric peer

Webex App

Questions?

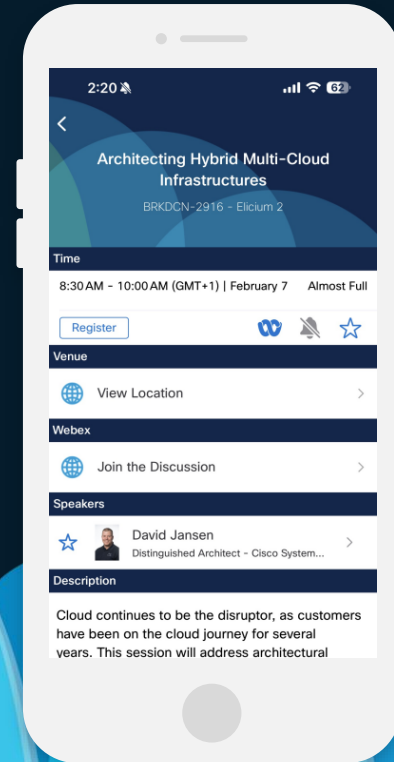
Use the Webex app to chat with the speaker after the session

How

- 1 Find this session in the Cisco Events 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 February 28, 2025.

CISCO *Live!*



Fill Out Your Session Surveys



Participants who fill out a minimum of 4 session surveys and the overall event survey will get a unique Cisco Live t-shirt.

(from 11:30 on Thursday, while supplies last)



All surveys can be taken in the Cisco Events mobile app or by logging in to the Session Catalog and clicking the 'Participant Dashboard'



Content Catalog

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at ciscolive.com/on-demand. Sessions from this event will be available from March 3.

Contact me at: ragoli@cisco.com



Thank you

CISCO *Live!*

CISCO *Live!*

GO BEYOND

A series of overlapping, rounded, teardrop-shaped abstract forms in various shades of blue, ranging from light to dark, positioned on the right side of the image.