The bridge to possible

BGP EVPN in Enterprise Campus

Building Scalable Fabrics with Catalyst 9000 Switches

Raj Kumar Goli
Technical Marketing Engineer

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

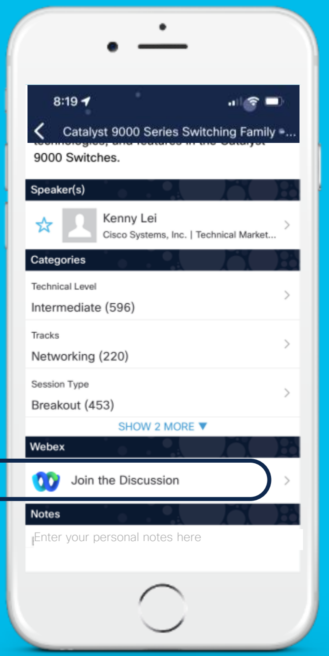

Webex spaces will be moderated until February 24, 2023.

Agenda

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

Education | Healthcare | Financial | Manufacturing | Media/Ent | Public Sector | Retail | Service Provider | Telecommunication | Utilities

BGP EVPN

16.9.1 (7/18)
- L3VNI
- L2VNI
- Multicast Replication
- IPv4
- VXLAN QoS

16.10.1 (10/18)
- L2VNI ARP/ND Suppression
- Ingress Replication
- MPLS VPN Border Support

16.11.1 (3/19)
- IPv6 Host Support
- IPv6 DHCP Relay
- L2VNI Auto RD/RT
- L2 Leaf support

16.12.1 (6/19)
- VPLS+EVPN L2 Interworking
- MPLS+EVPN L3 Interworking
- Catalyst + Nexus + XR Interop
- Border + Spine + Leaf Role
- EVPN Multi-homing with SVL
- EVPN SNMP MIB support

17.1.1 (10/19)
- Tenant Routed Multicast v4
- TRM + mVPN L3 Interworking
- Centralized Gateway
- L2 Leaf ARP/ND Suppression
- Per-VNI BUM Rate-limiter

17.2.1 (3/20)
- RT2-to-RT5 Originate

17.3.1 (6/20)
- 512 VNI Support
- Tenant Routed Multicast v6
- Wide Area Bonjour over EVPN
- mDNS over VXLAN

17.4.1 (11/20)
- Private VLAN over EVPN

17.6.1 (7/21)
- Isolated PVLAN Support
- ESI-Based Multihoming
- 512 VNI Scale
- Micro-location mDNS over VXLAN
- 802.1Q Tunnel Mode

17.7.1 (11/21)
- L3 TRM with MDT Data

17.8.1 (04/22)
- Optimizing L2 Multicast with IGMP/mld Proxy
- Protect Vlan over EVPN Fabric

17.9.1 (07/22)
- BGP EVPN over IPsec

17.10.1 (11/22)
- IPv6 Undelay
- L3TRM support with BGP EVPN over IPsec
- BGP EVPN Dynamic peering

VXLAN Overview

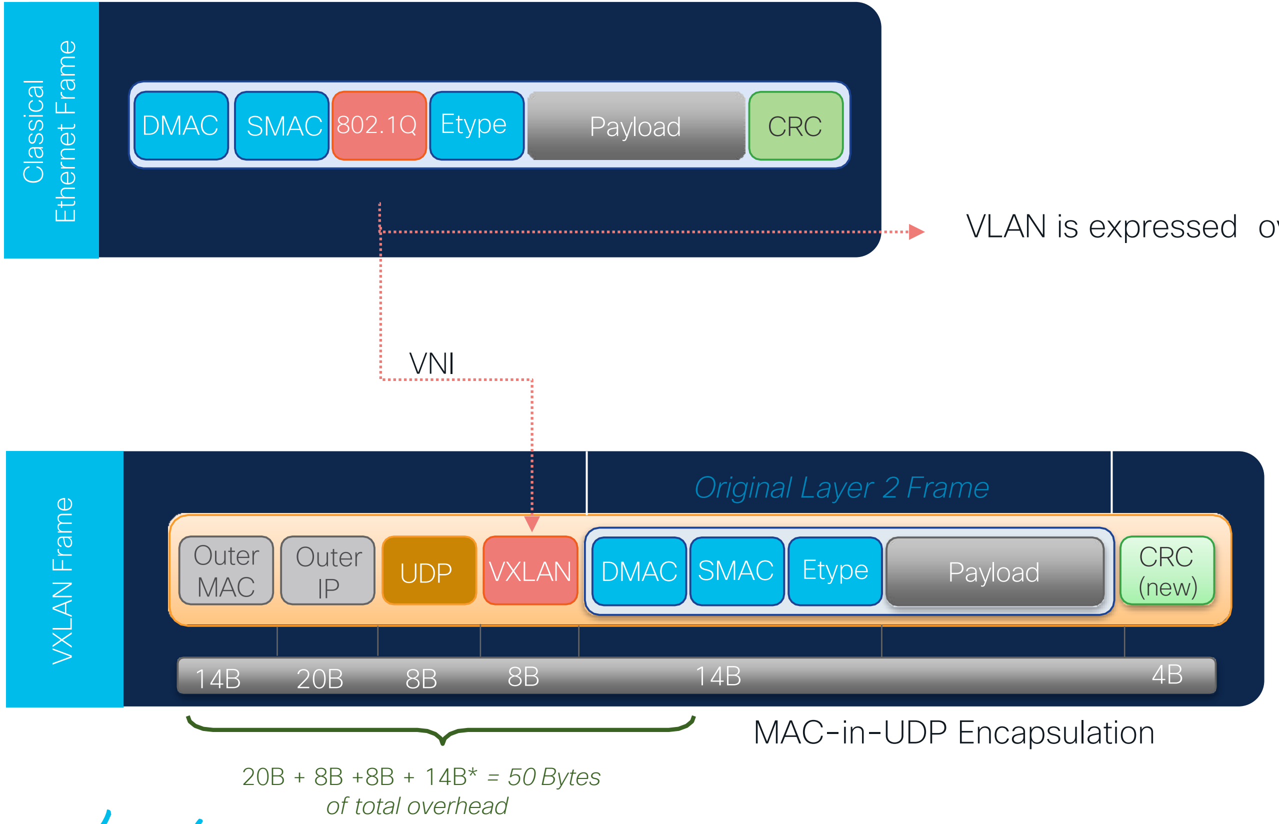

VLAN is expressed over 12 bits (802.1Q tag)

MAC-in-UDP Encapsulation

VXLAN leverages the VNI field with a total address space of 24 bits

The VXLAN Network Identifier (VNI/VNID) is part of the VXLAN Header

VXLAN with BGP EVPN

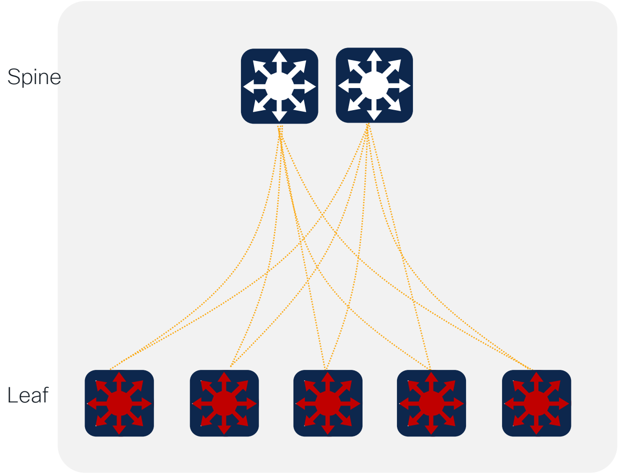

- Standards based Overlay (VXLAN) with Standards based Control-Plane (BGP)
- Layer-2 MAC and Layer-3 IP information distribution by Control-Plane (BGP)
- Forwarding decision based on Control-Plane (minimizes flooding)
- Integrated Routing/Bridging (IRB) for Optimized Forwarding in the Overlay
- Multi-Tenancy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

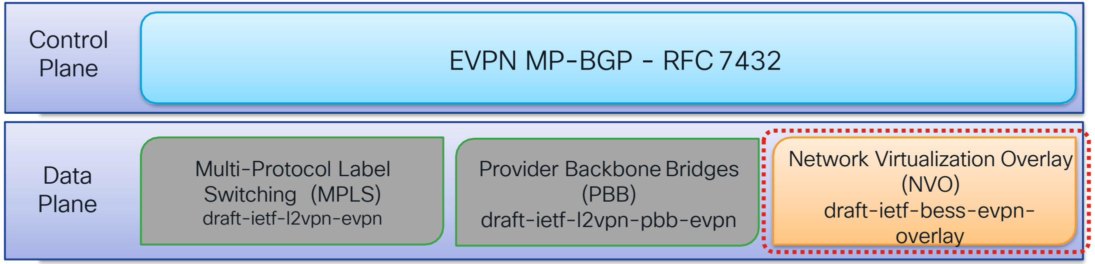

EVPN over NVO Tunnels (VXLAN)

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.

* - Roadmap
★ - Recommended

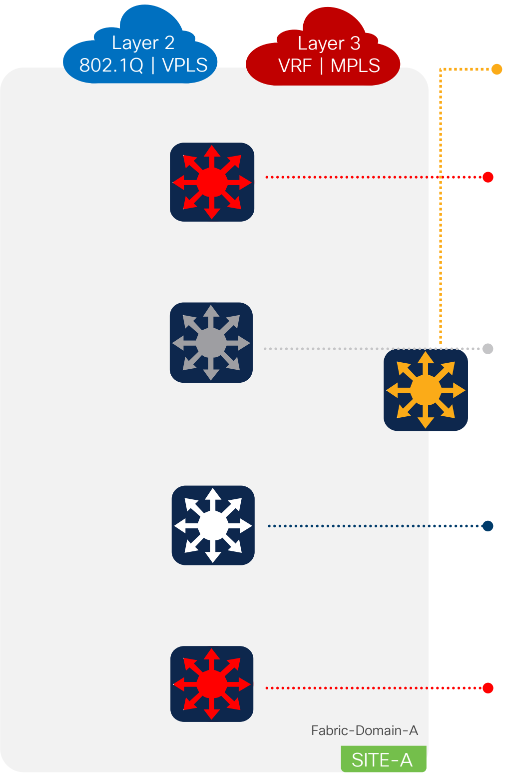

System Support	Mode
Nexus 9000	Standalone

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	Multi-vendor IT strategy
One Fabric Architecture	Unified operation across – Campus | DC | WAN
Proven and Scalable	BGP Protocol History. Minimum new learning curve
Hierarchical Fabric Domain	Multi-tier Overlay network architecture
Flexible Overlay	Use-case driven customize Overlay networks Types and Topologies

Enterprise BGP EVPN Reference Architecture

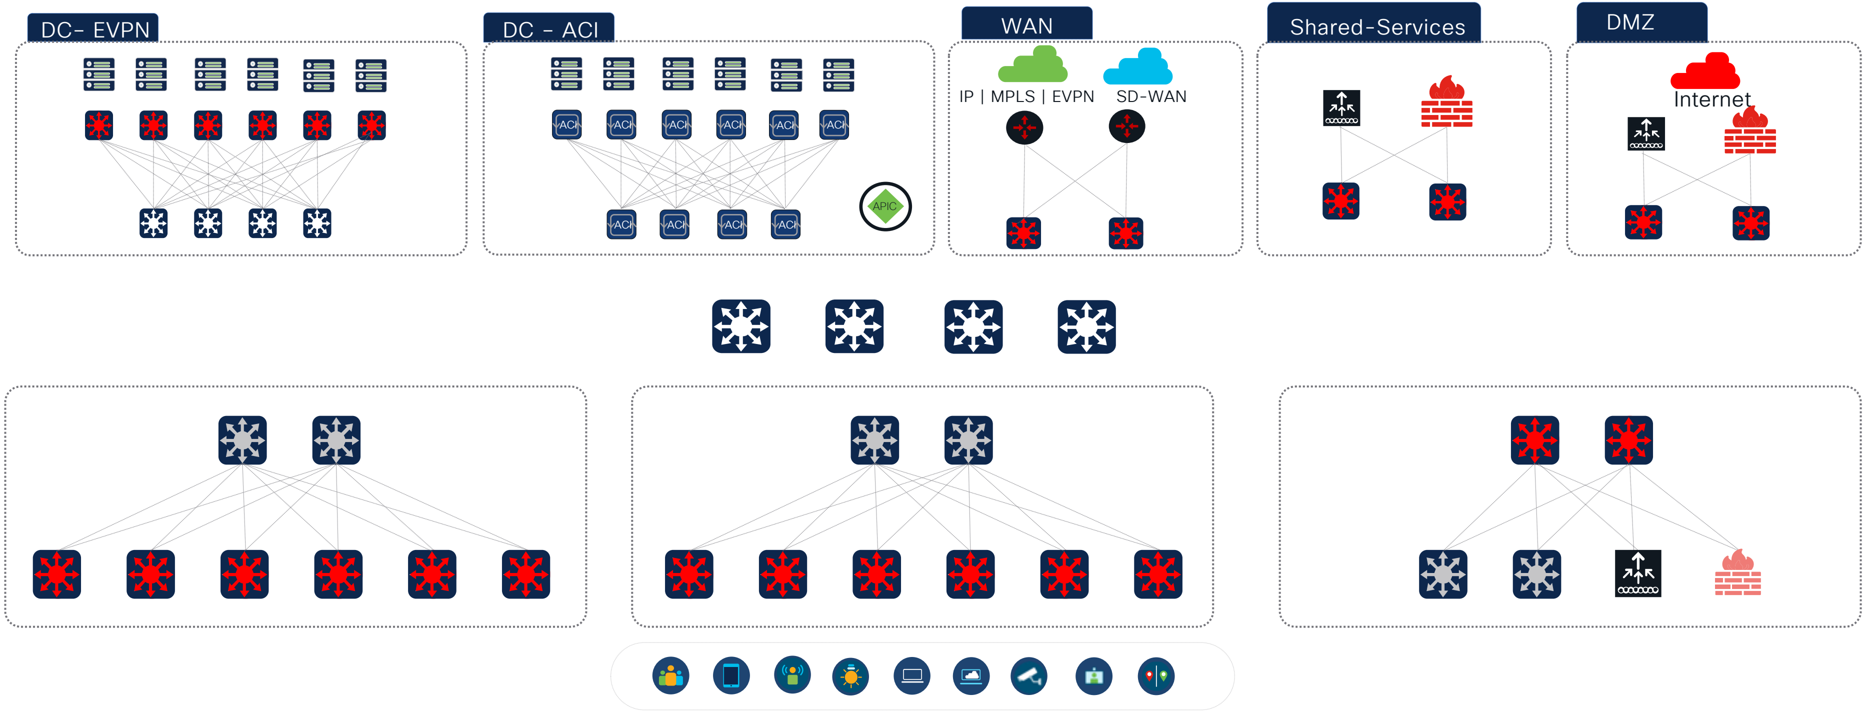

Industry Standard

Standard-based Fabric
Multi-vendor interoperable
Broad innovation adoption

Unified Fabric

Cross-PIN single fabric
Extensible beyond site
Simplified Management

Proven

Reliable control-plane
Multi-protocol capabilities
Less new learning-curve

Hierarchical

Non-blocking architecture
Structured & Scalable fabric
Hybrid system role support

Flexible

Complex network solution
Tailored L2/L3 overlays
Deep eco-system integration

EVPN Basics

VxLAN Packet Structure

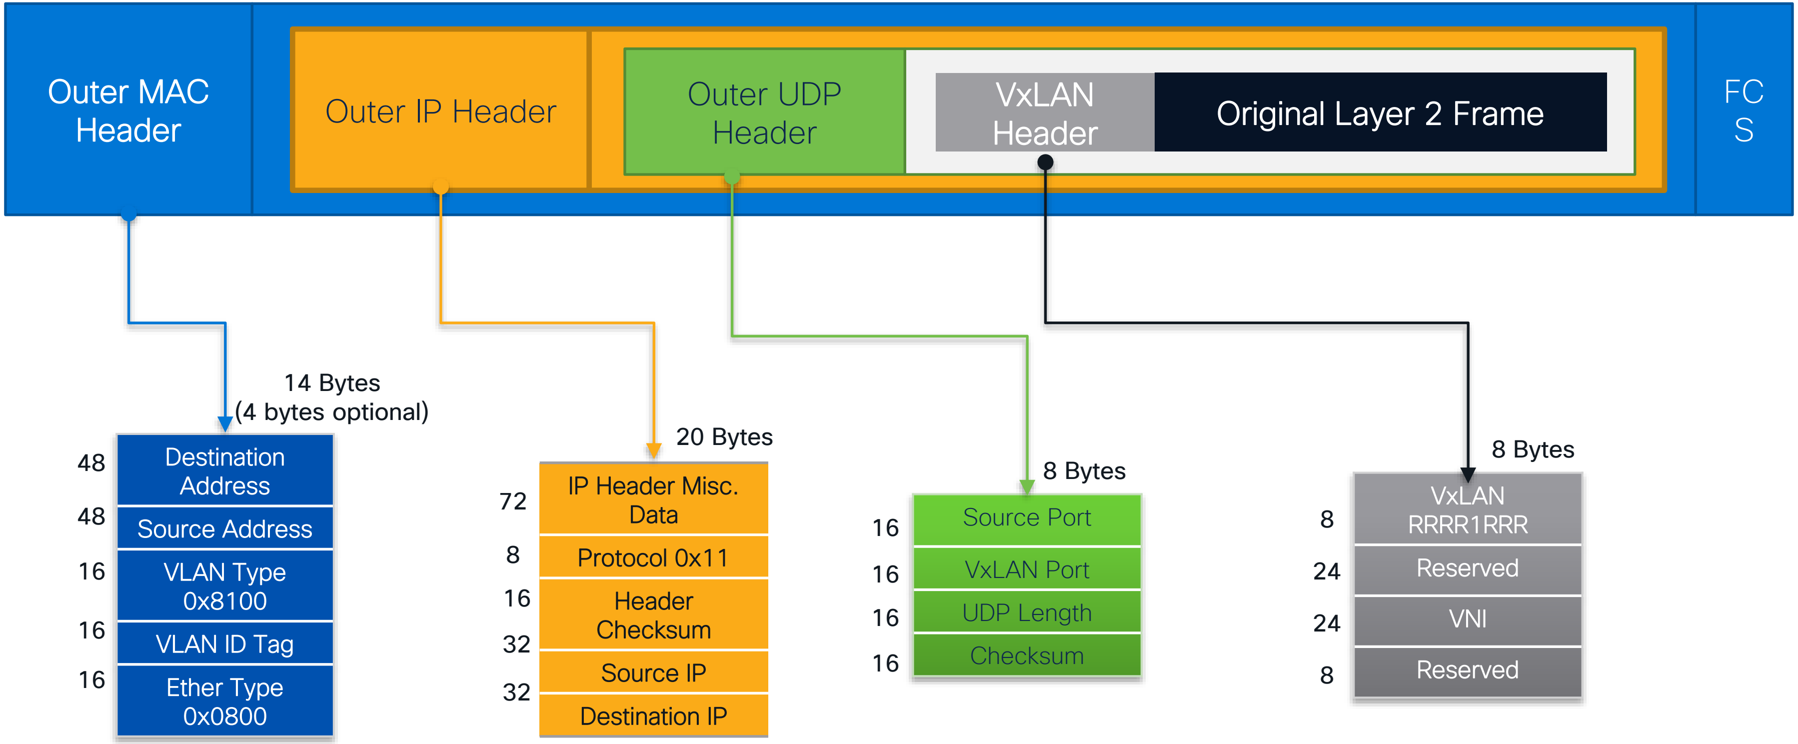

50(54) Bytes of overhead

MP-BGP EVPN Route Type(s)

RFC	Type	Description
RFC 7432	Type 1	Ethernet Auto-Discovery (A-D) route
RFC 7432	Type 2	MAC/IP advertisement route
RFC 7432	Type 3	Inclusive Multicast Route EVPN Ingress Replication (IR) (unicast mode for BUM)
RFC 7432	Type 4	Ethernet Segment Route
RFC 9136	Type 5	IP Prefix Route → Layer-3 VNI Route
RFC 9251	Type 6	Selective Multicast Ethernet Tag Route
RFC 9251	Type 7	IGMP Sync routes (Join/Leave)
RFC 9251	Type 8	IGMP Sync routes (Join/Leave)

VXLAN Overview

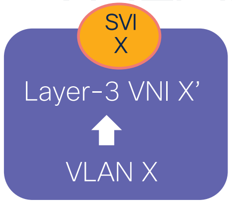

- 1 Layer-3 VNI per Tenant (VRF) for routing
- VNI X’ is used for routed packets

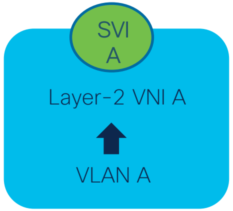

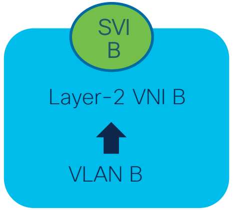

- 1 Layer-2 VNI per Layer-2 segment
- Multiple Layer-2 VNIs per tenant
- VNI A’ and B’ are used for bridged packets

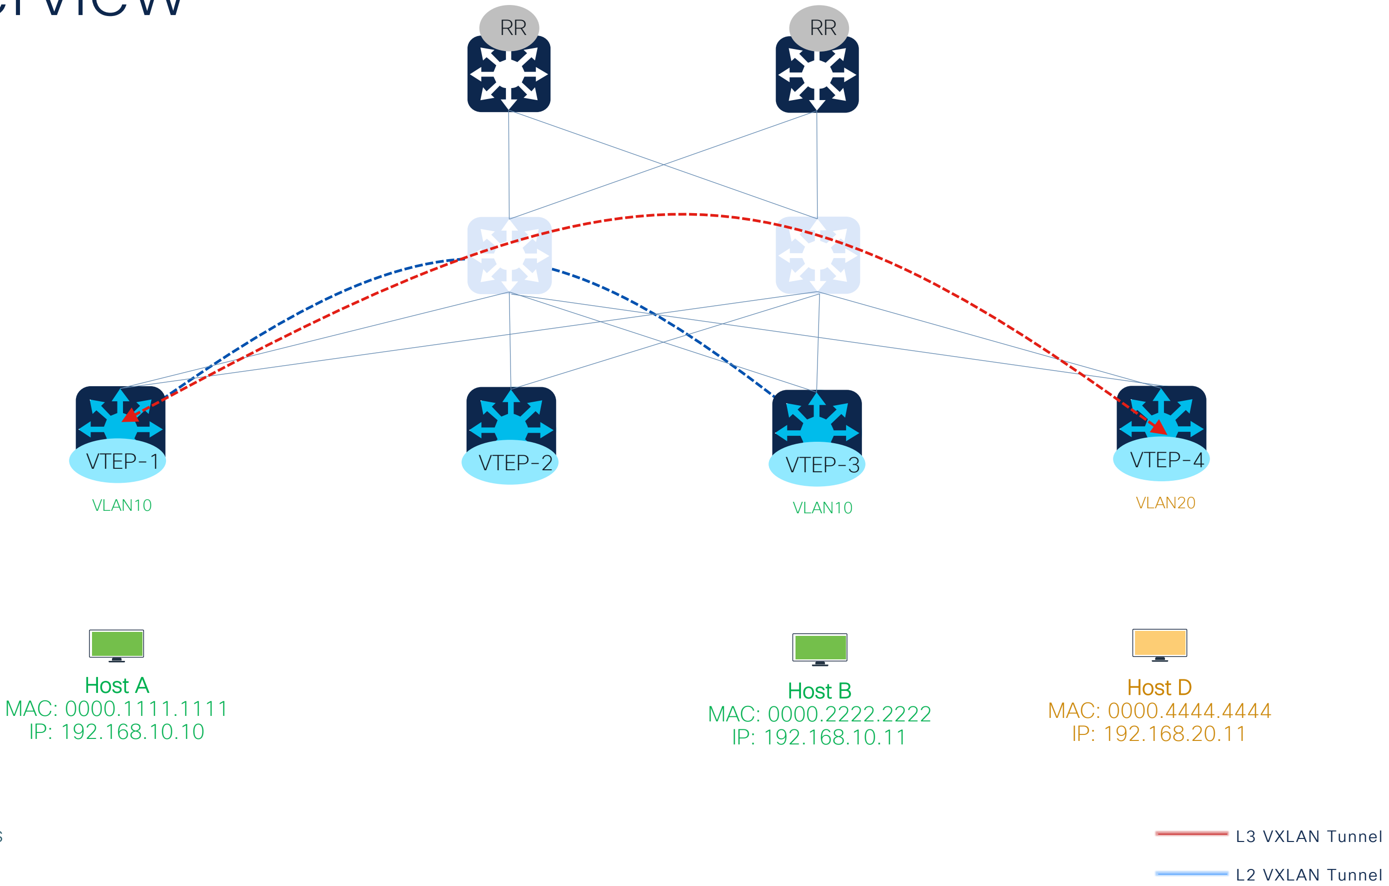

Packet Walk – ARP Request

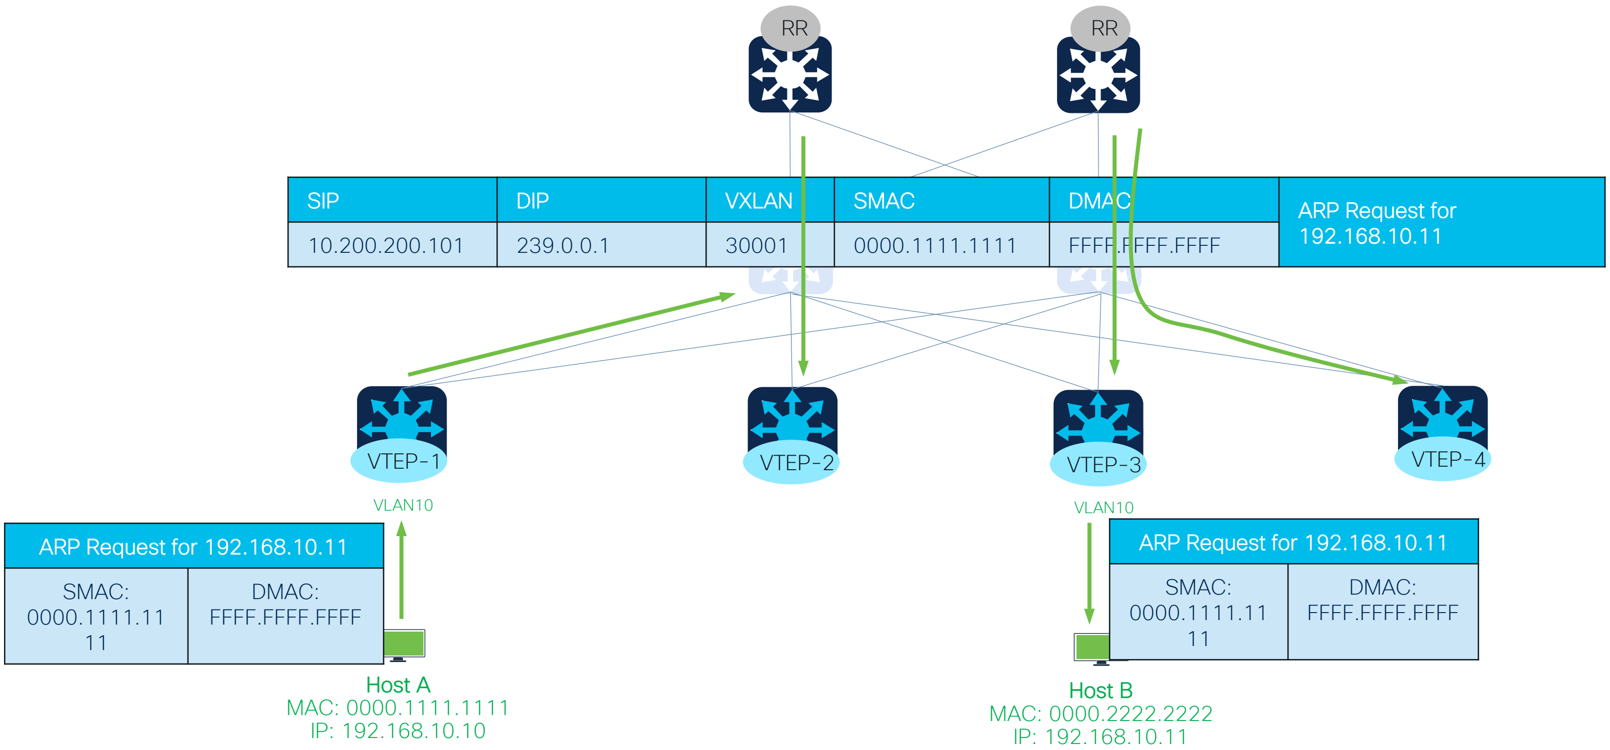

Packet Walk – ARP Response

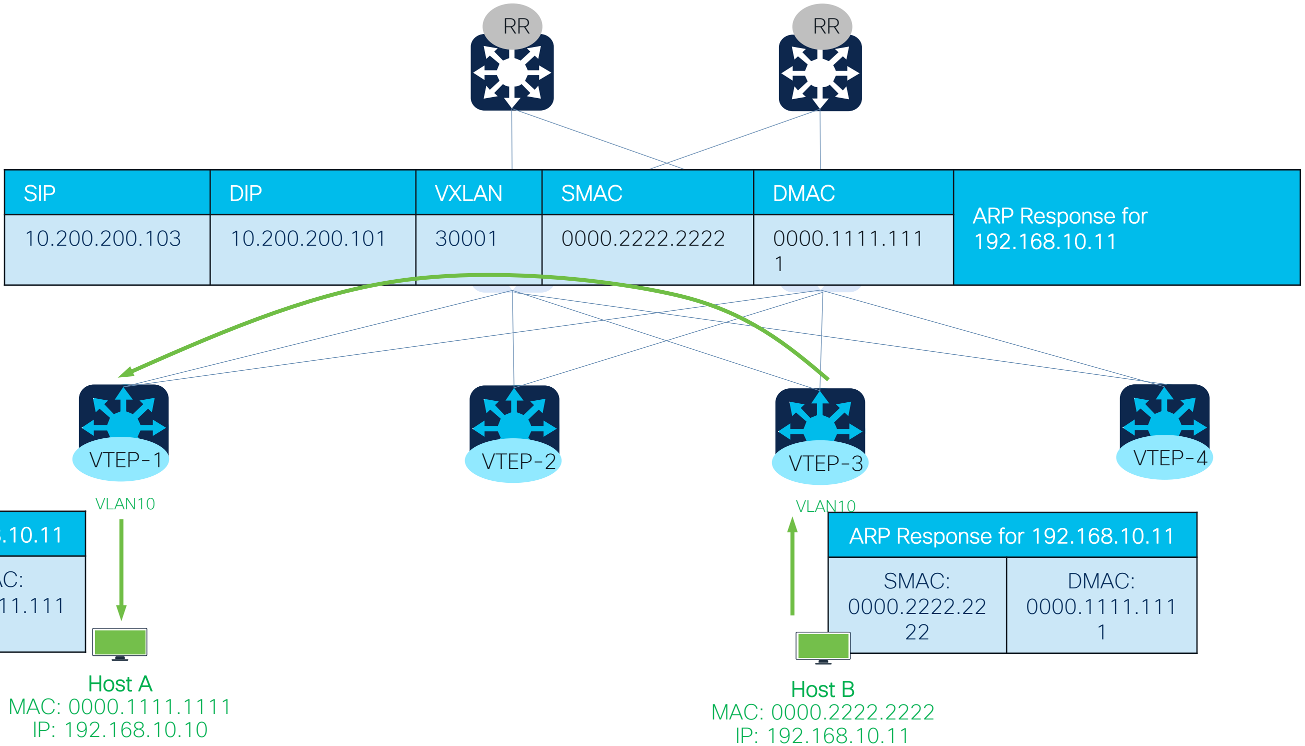

SIP	DIP	VXLAN	SMAC	DMAC	ARP Response for 192.168.10.11
10.200.200.103	10.200.200.101	30001	0000.2222.2222	0000.1111.111 1	

ARP Response for 192.168.10.11	
SMAC: 0000.2222.22 22	DMAC: 0000.1111.111 1

ARP Response for 192.168.10.11	
SMAC: 0000.2222.22 22	DMAC: 0000.1111.111 1

MAC/IP Advertisement route

"MAC or MAC/IP host Advertisement (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

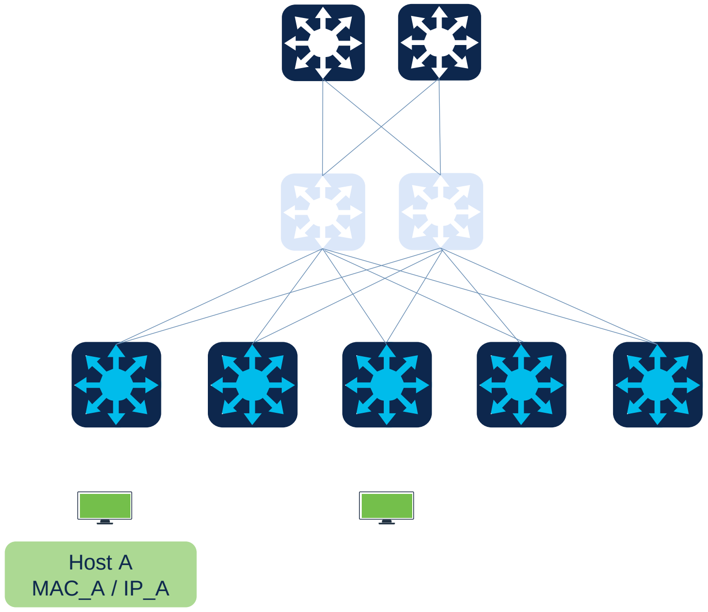

Protocol Learning & Distribution

“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

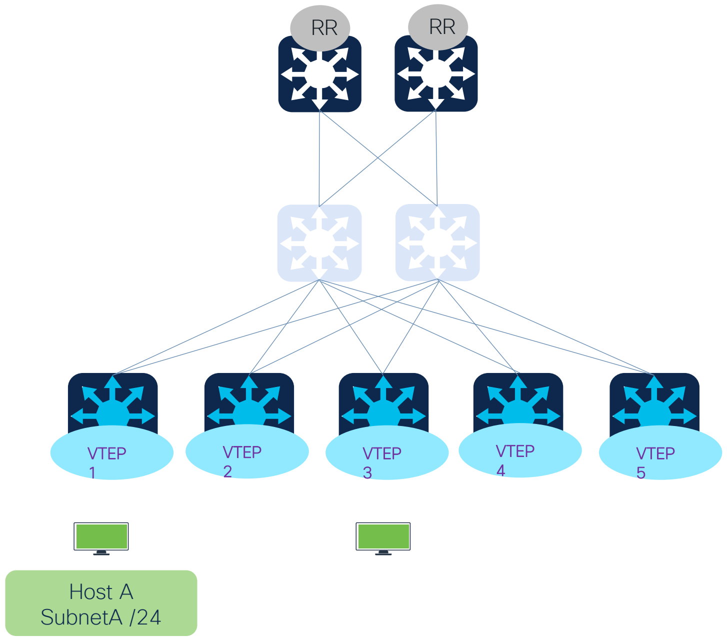

Underlay Network

Underlay Design Considerations

Layer 3 Access

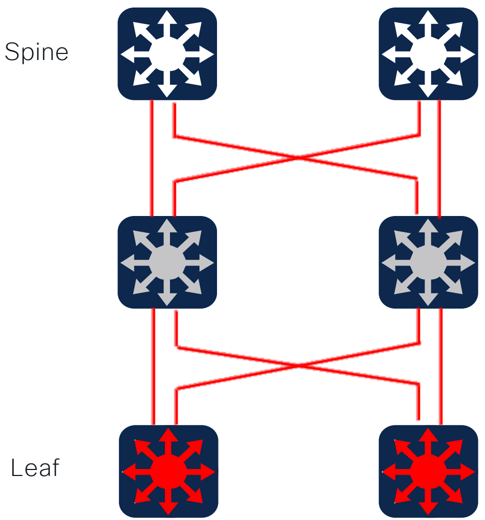

Leaf Layer – Access
Spine/RR – Direct | Multi-hop
Underlay | Overlay IP gateway
ECMP | Multicast
L2 | L3 Overlay support

Cisco StackWise-Virtual

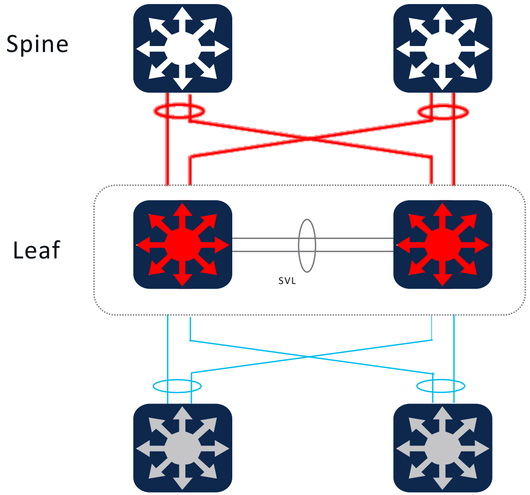

Leaf Layer – Distribution
Spine/RR – Direct | Multi-hop
Underlay | Overlay IP gateway
MEC | ECMP | Multicast
L2 | L3 Overlay support

Layer 2 | 3 Distribution

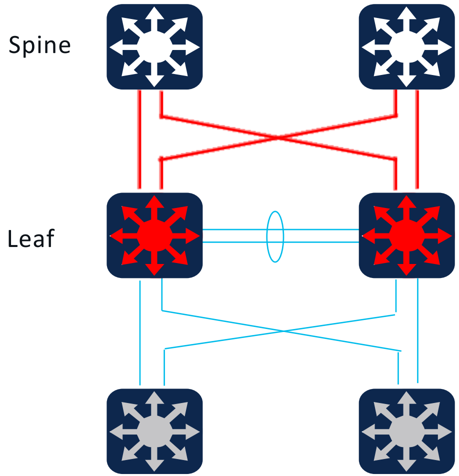

Leaf Layer – Distribution
Spine/RR – Direct | Multi-hop
Underlay | Overlay IP gateway
FHRP | ECMP | Multicast
L3 Overlay. No L2 Extension

Layer 2 Layer 3

Underlay Design Considerations

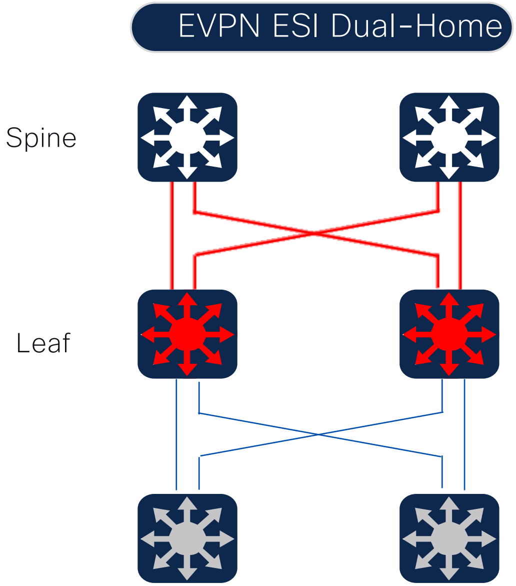

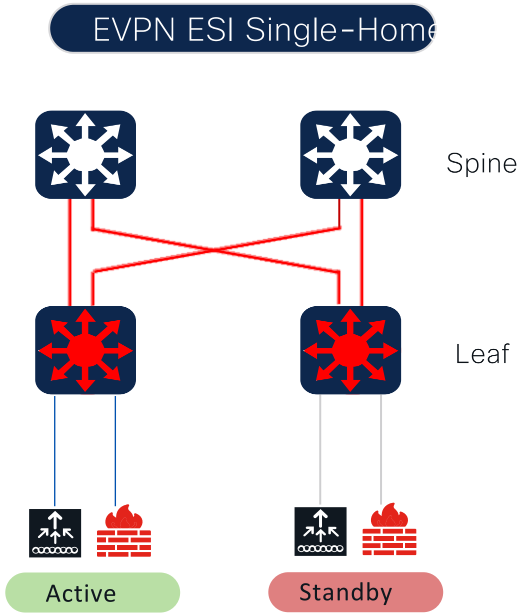

Leaf Layer – Distribution
Spine/RR – Direct | Multi-hop
Per-ESI AnyCast Gateway
Per-VLAN | ECMP | Multicast
L2 | L3 Overlay support

Leaf Layer – Distribution
Spine/RR – Direct | Multi-hop
Per-ESI AnyCast Gateway
Per-Leaf | ECMP | Multicast
L2 | L3 Overlay support

ESI – Ethernet Segment Identifier
Layer 2 Layer 3

Underlay Unicast Routing Design Alternatives

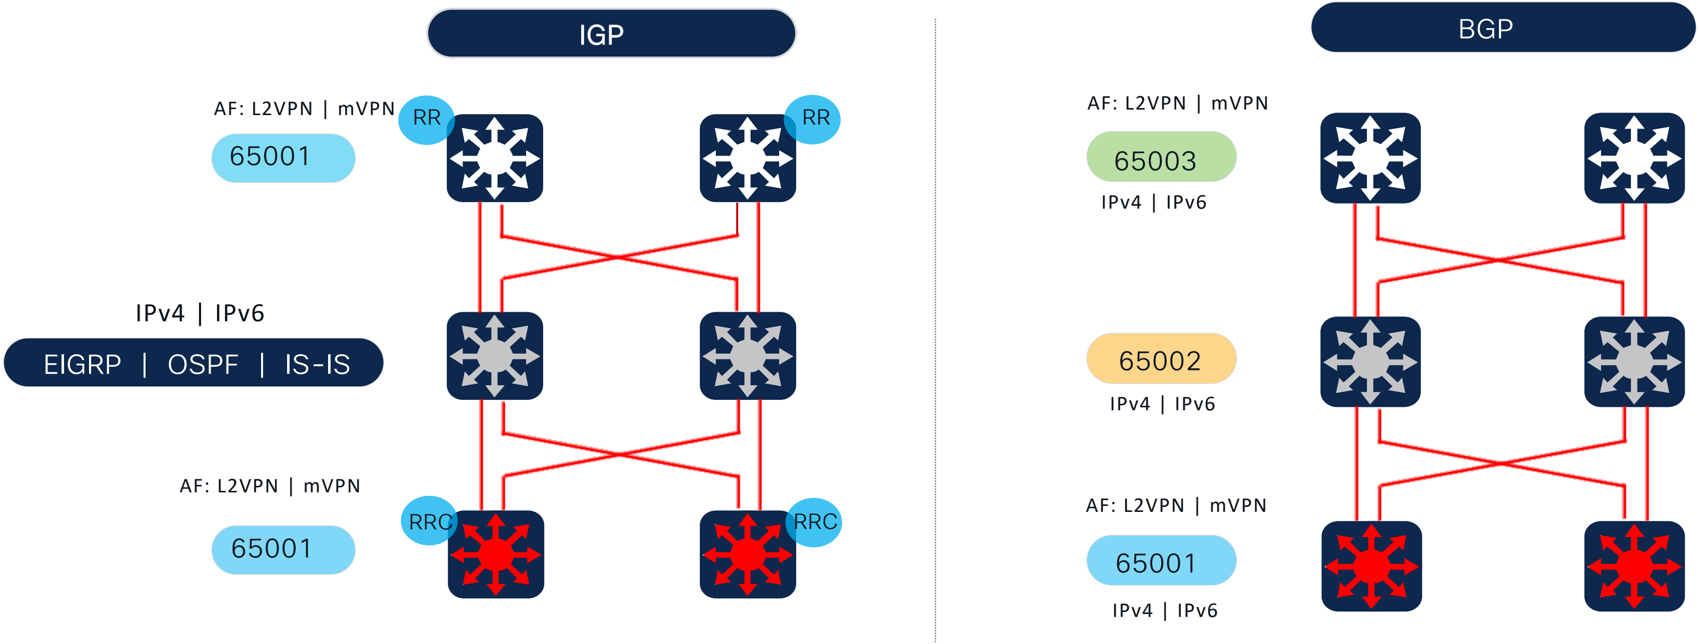

Underlay Unicast

Flexible Underlay Unicast alternatives – IGP | BGP | MPLS | IPSec
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.

Underlay Unicast Routing Design Alternatives

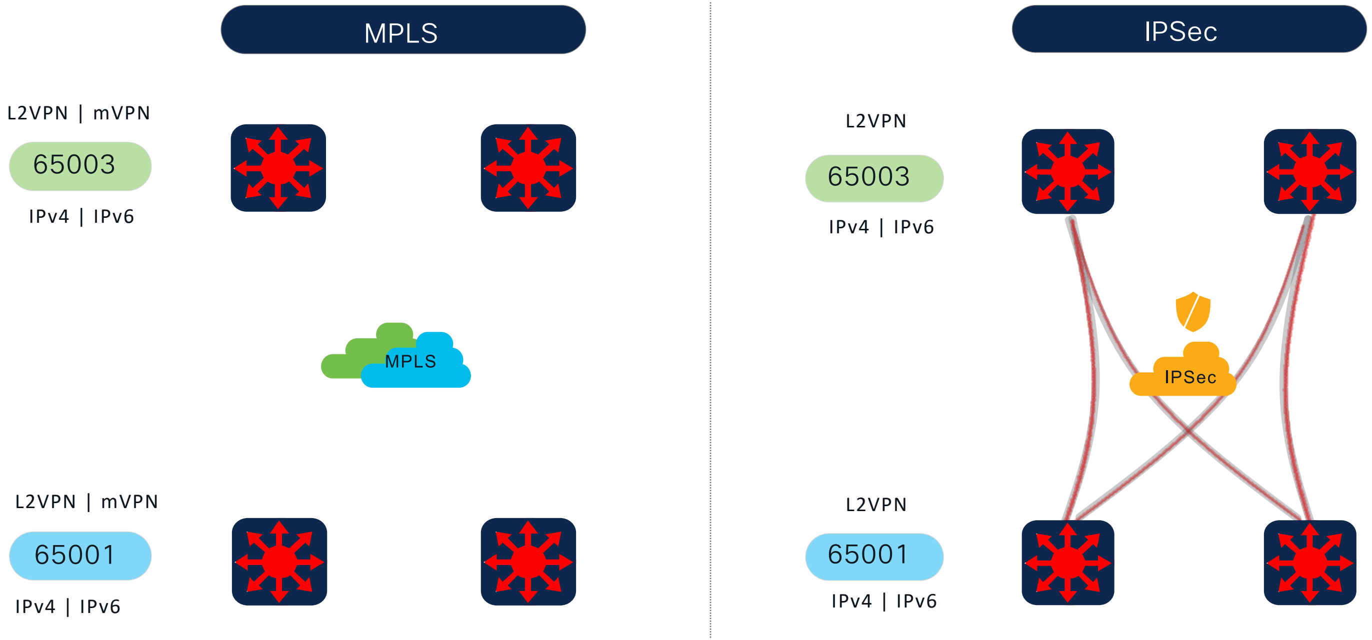

Underlay Unicast

Flexible Underlay Unicast alternatives – IGP | BGP | MPLS | IPSec

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.

Overlay Network Topologies

Single Cluster Fabric Architecture

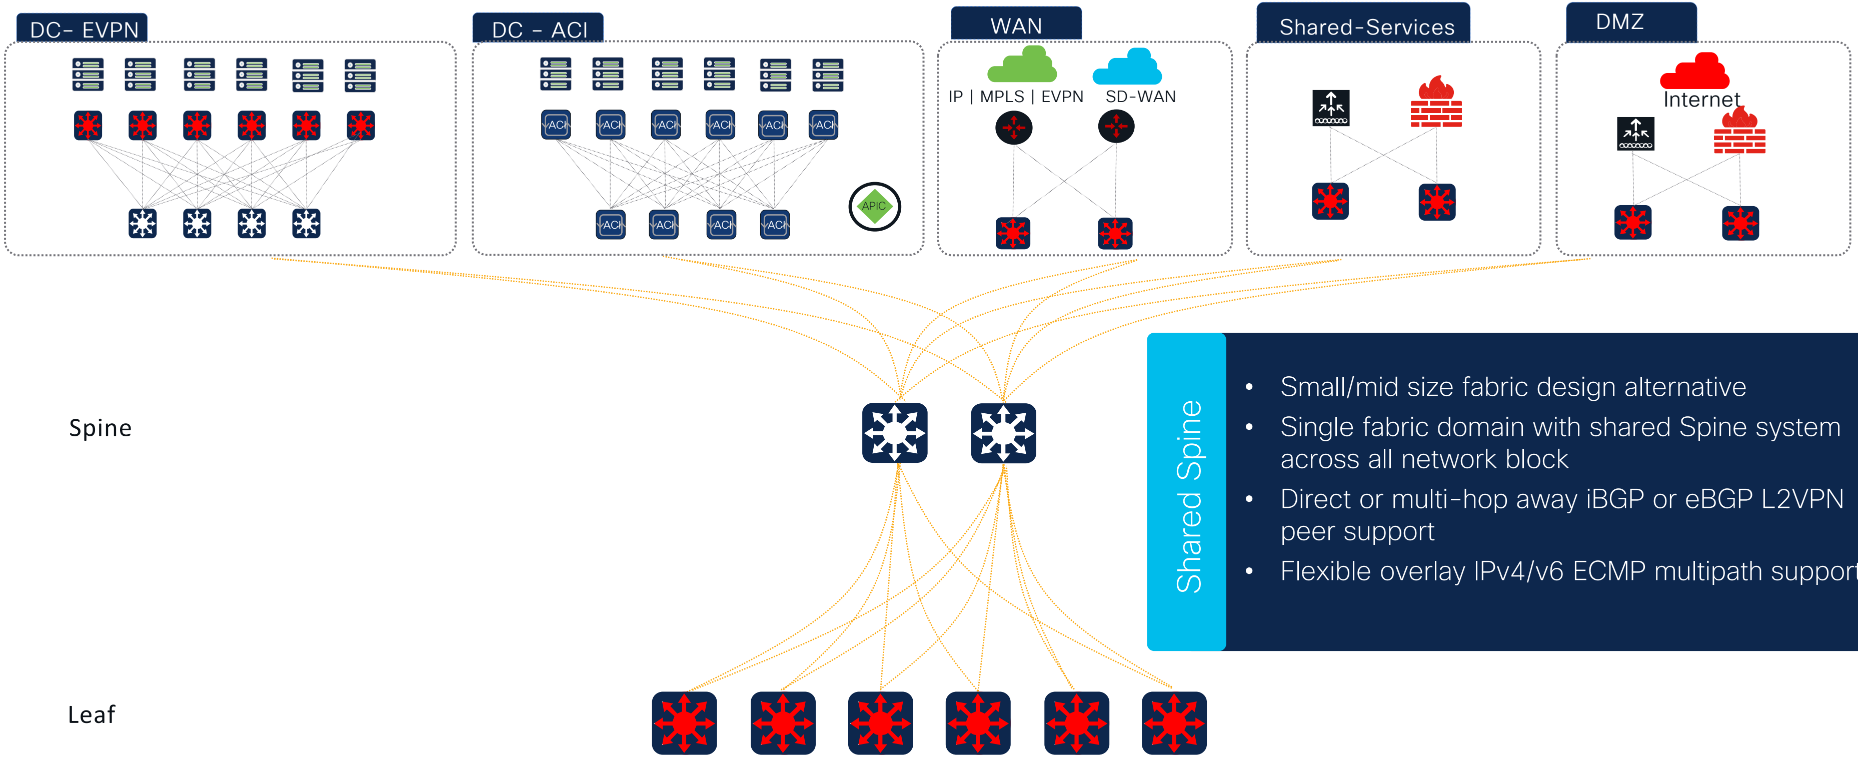

Shared Spine

- 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

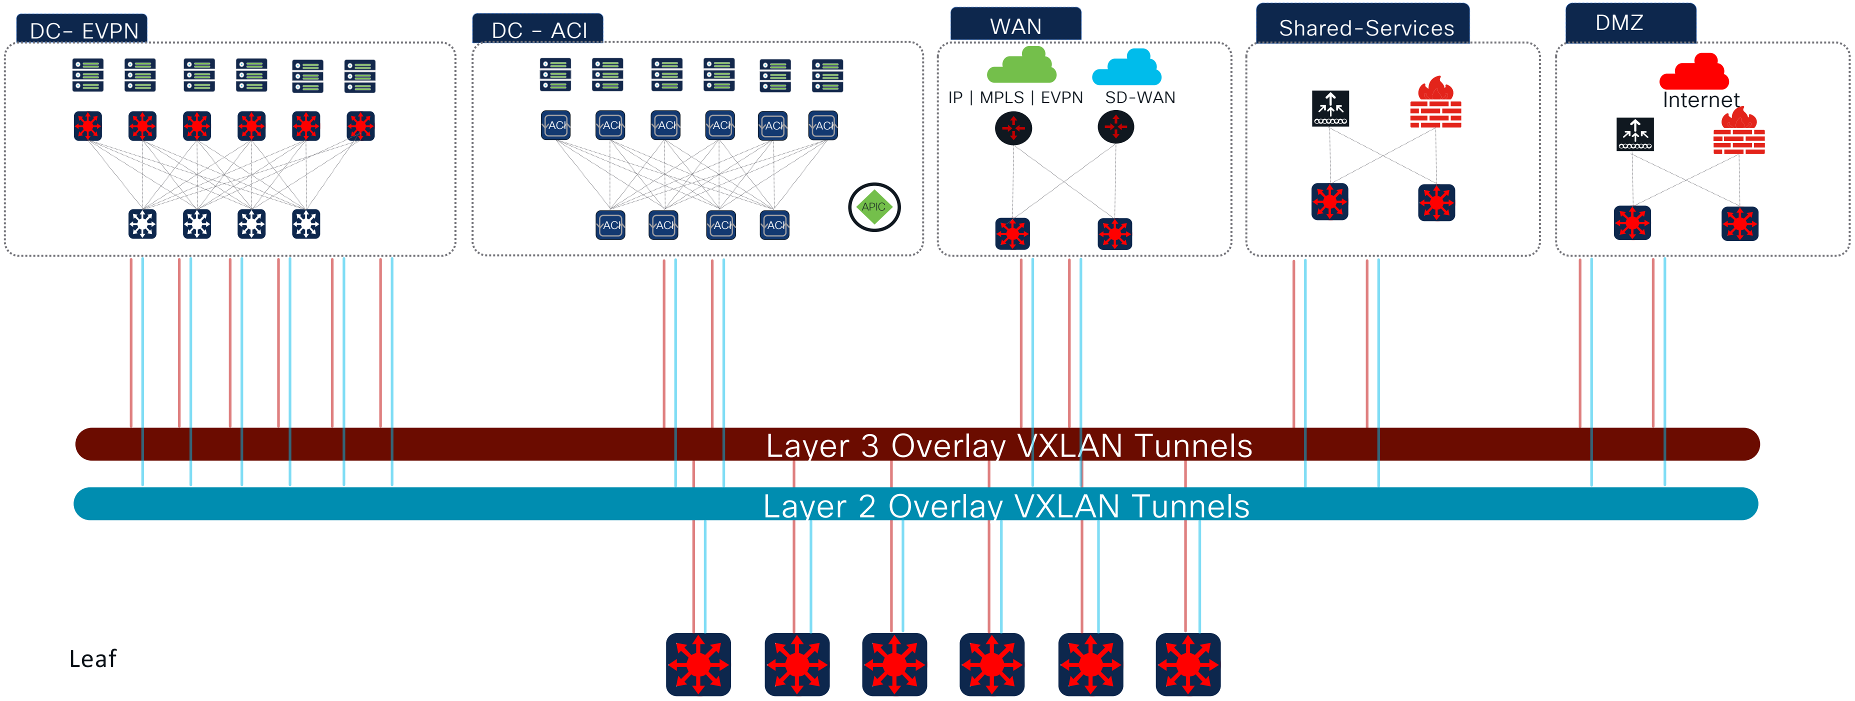

Non-Hierarchical Fabric

- 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

L3 VXLAN Tunnel
L2 VXLAN Tunnel

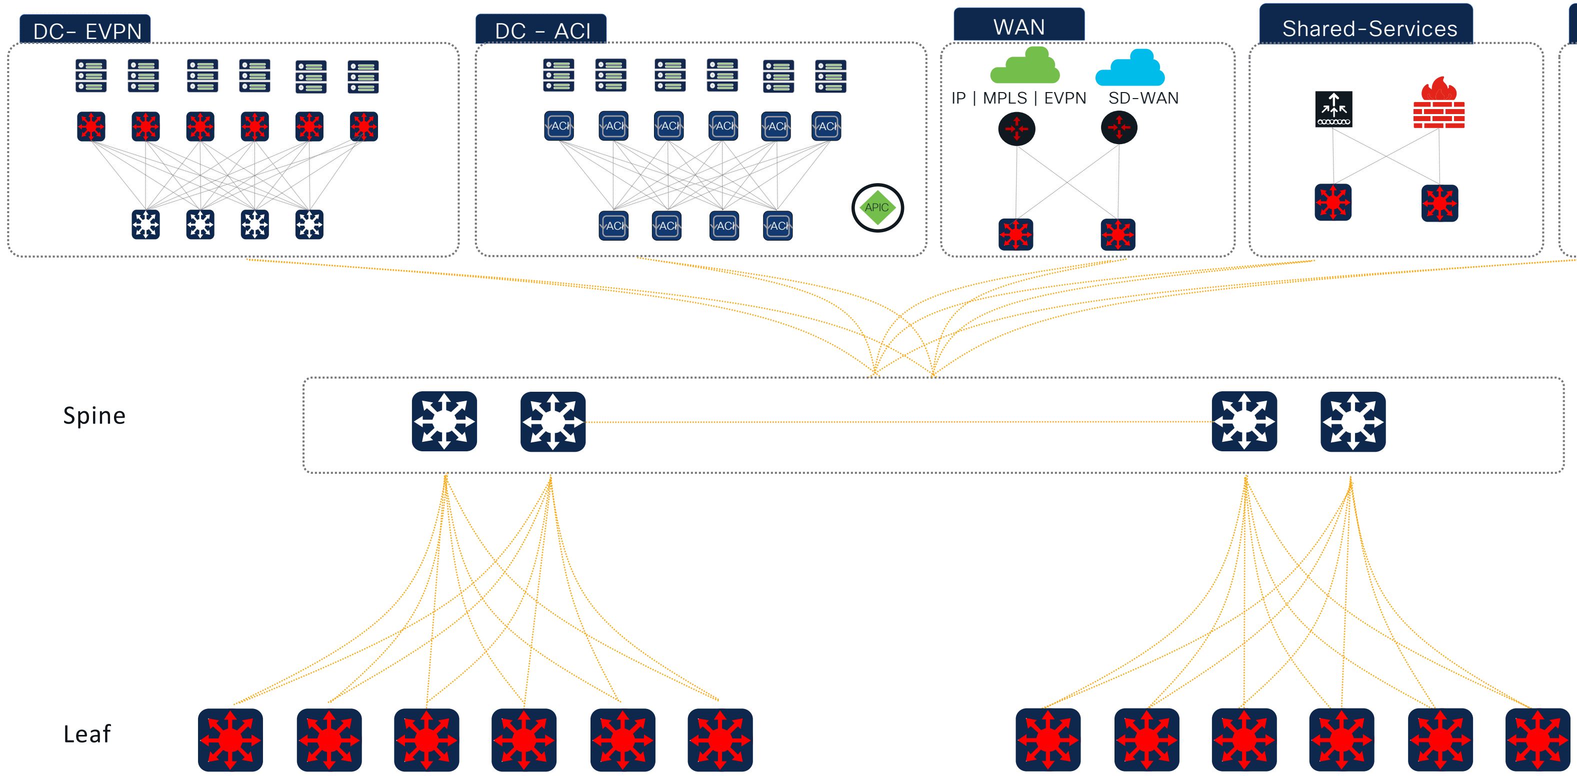

Multi Cluster Fabric Design

Distributed Spine

- Mid to large size fabric design alternative
- Single fabric domain with distributed RR clusters for high scale fabric
- RR cluster grouping for end-to-end simplified overlay fabric network
- Limited Layer 2 overlay support. Overlay Multicast (TRM) not supported.

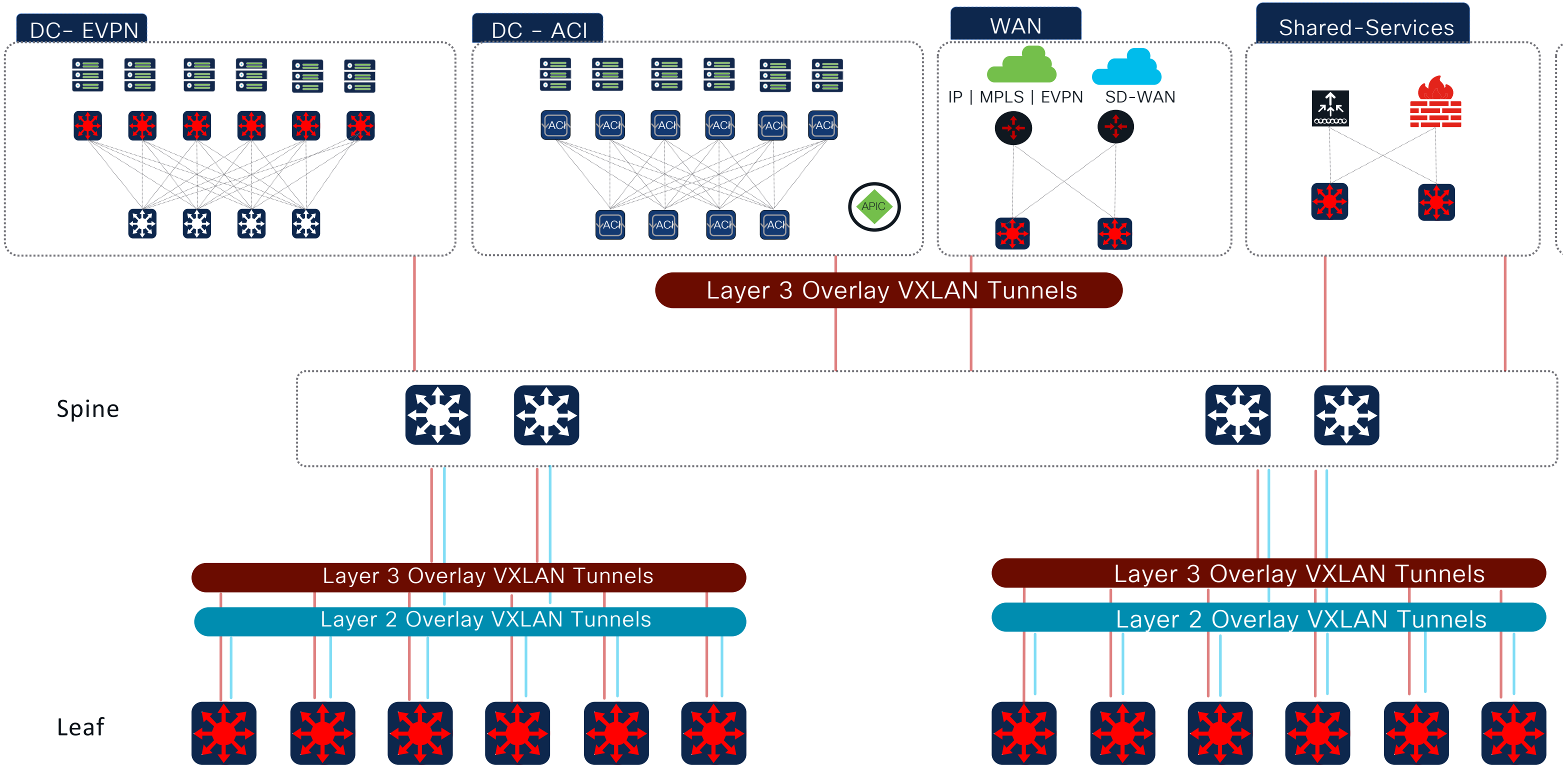

Hierarchical Fabric Design

Distributed Spine

- Mid to large size fabric design alternative
- Single fabric domain with distributed RR clusters for high scale fabric
- RR cluster grouping for end-to-end simplified overlay fabric network
- Limited Layer 2 overlay support. Overlay Multicast (TRM) not supported.

L3 VXLAN Tunnel
L2 VXLAN Tunnel

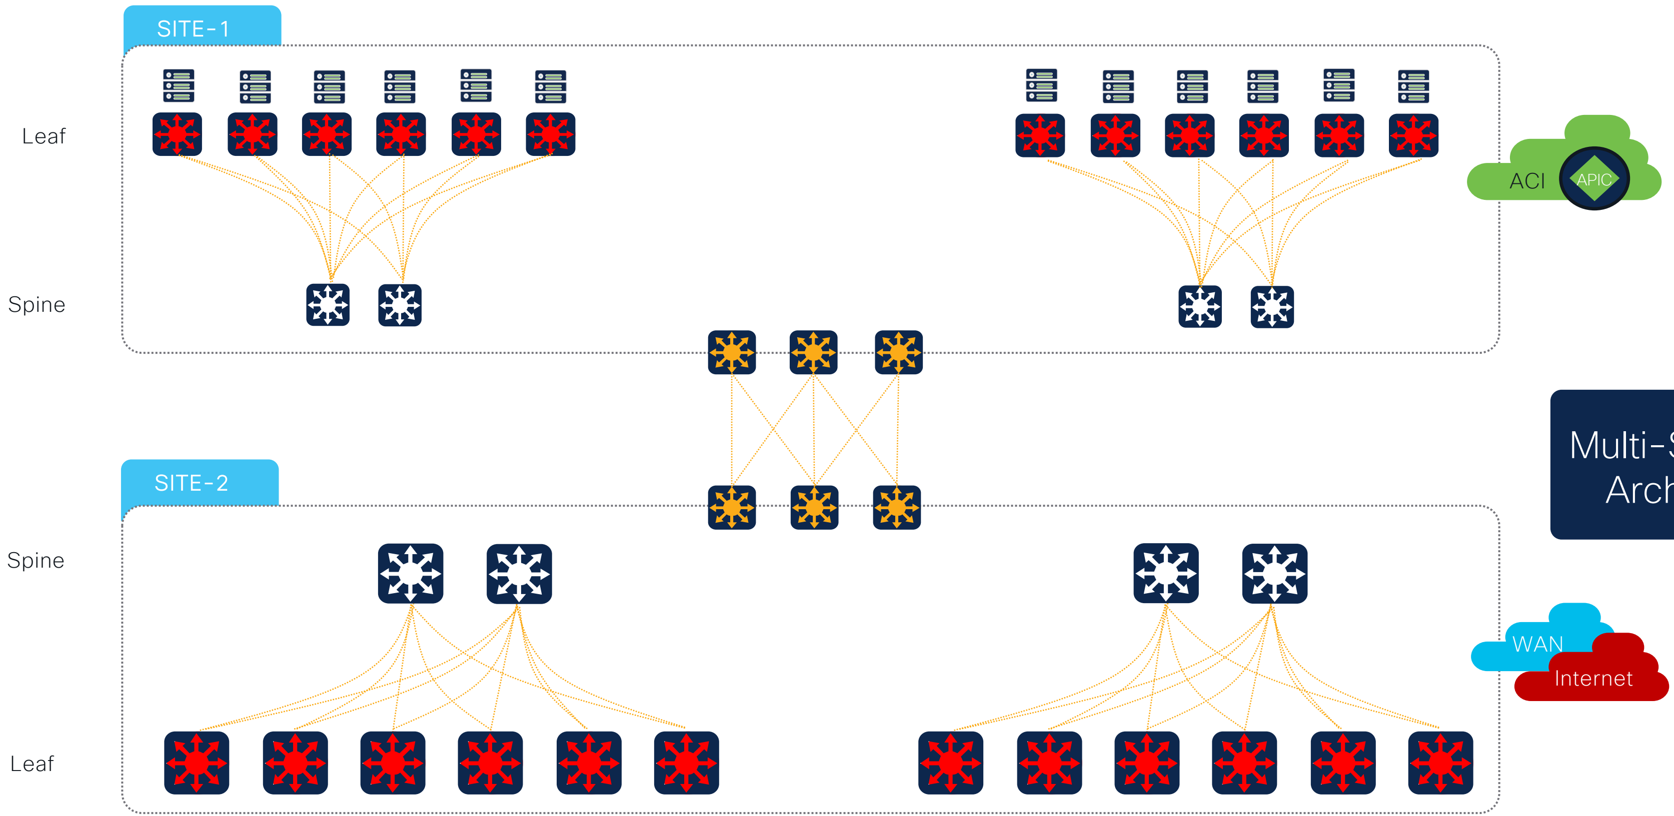

Multi-Site Fabric Architecture

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)

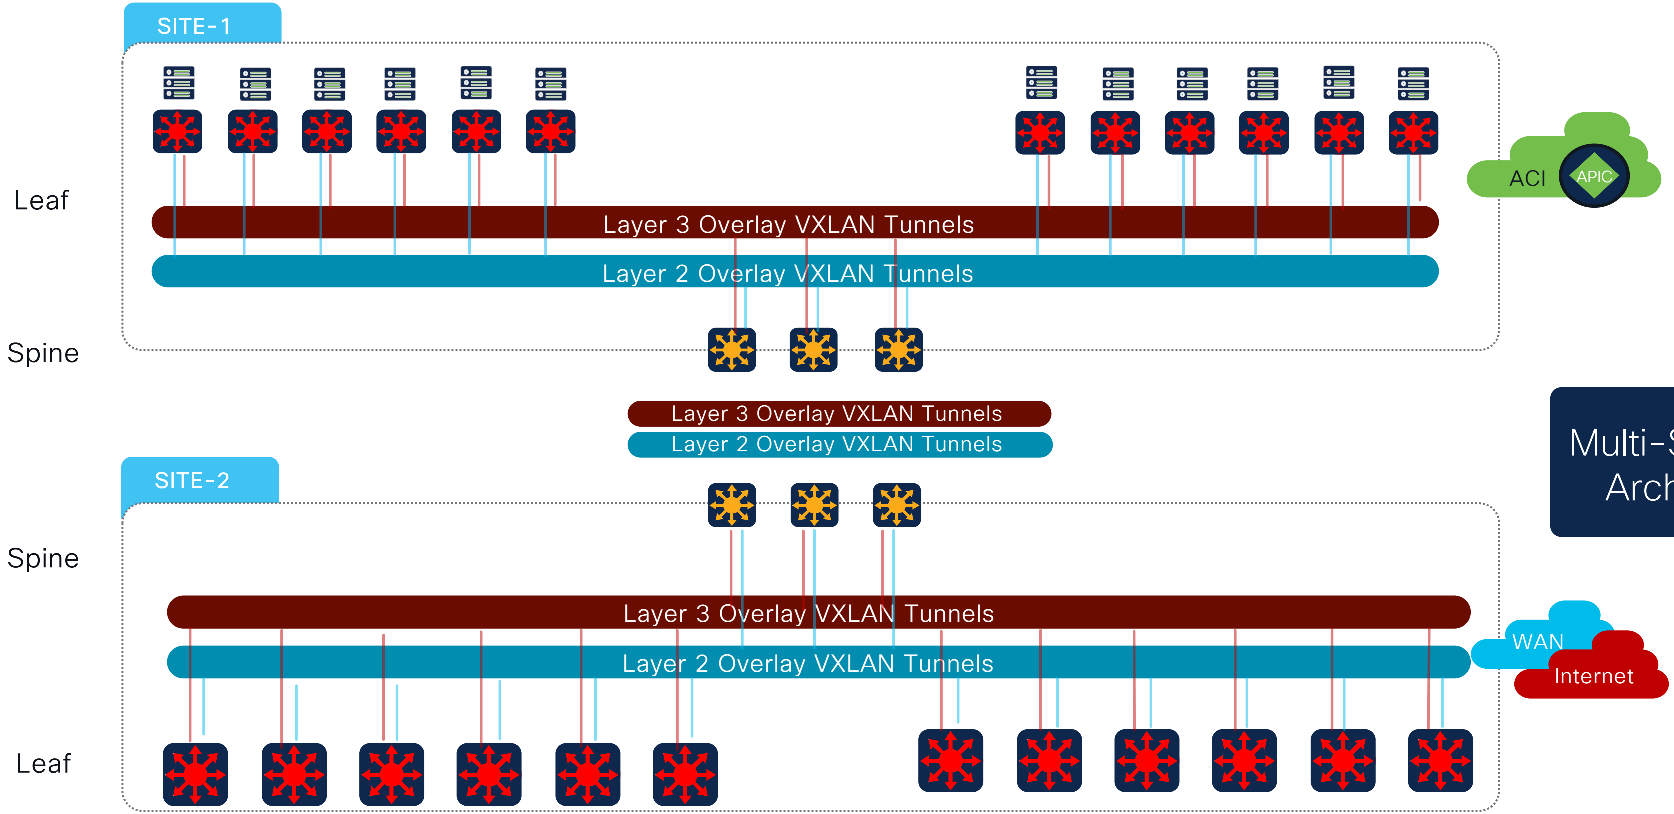

Multi-Site Fabric Architecture

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)

L3 VXLAN Tunnel
L2 VXLAN Tunnel

Flexible Routing and Bridging Overlay Types

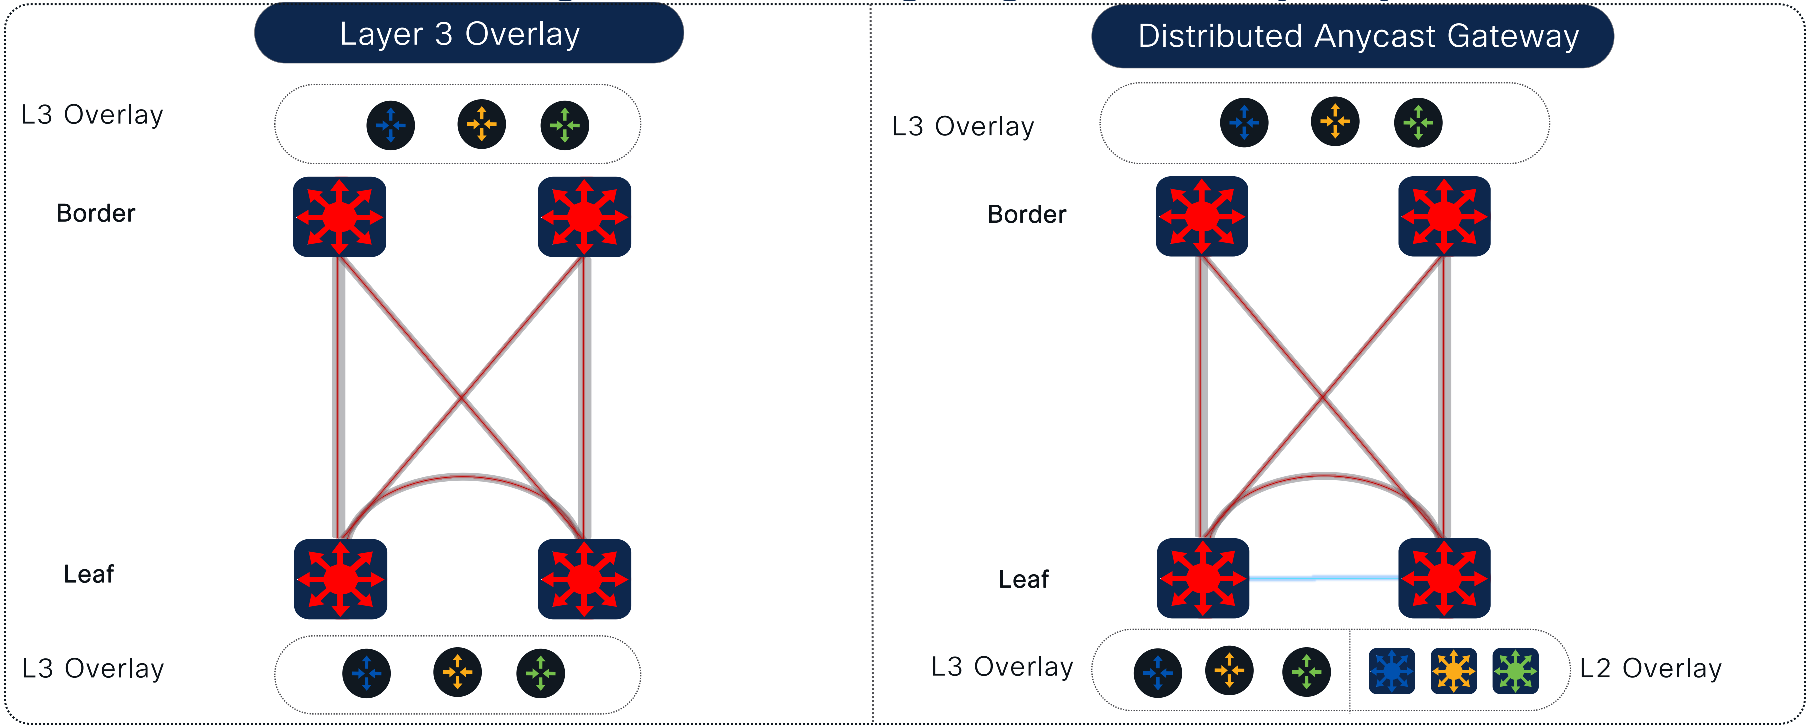

Overlay Types

Four overlay network types support at any network layer point

Route first. Bridge when-and-where need rule for scalable fabric architecture

Feature rich Layer 3 overlay network support – Unicast | Multicast – IPv4 | IPv6

Scalable Layer 2 overlay solution with suppression, flood management and more

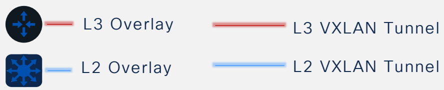

Flexible Routing and Bridging Overlay Types

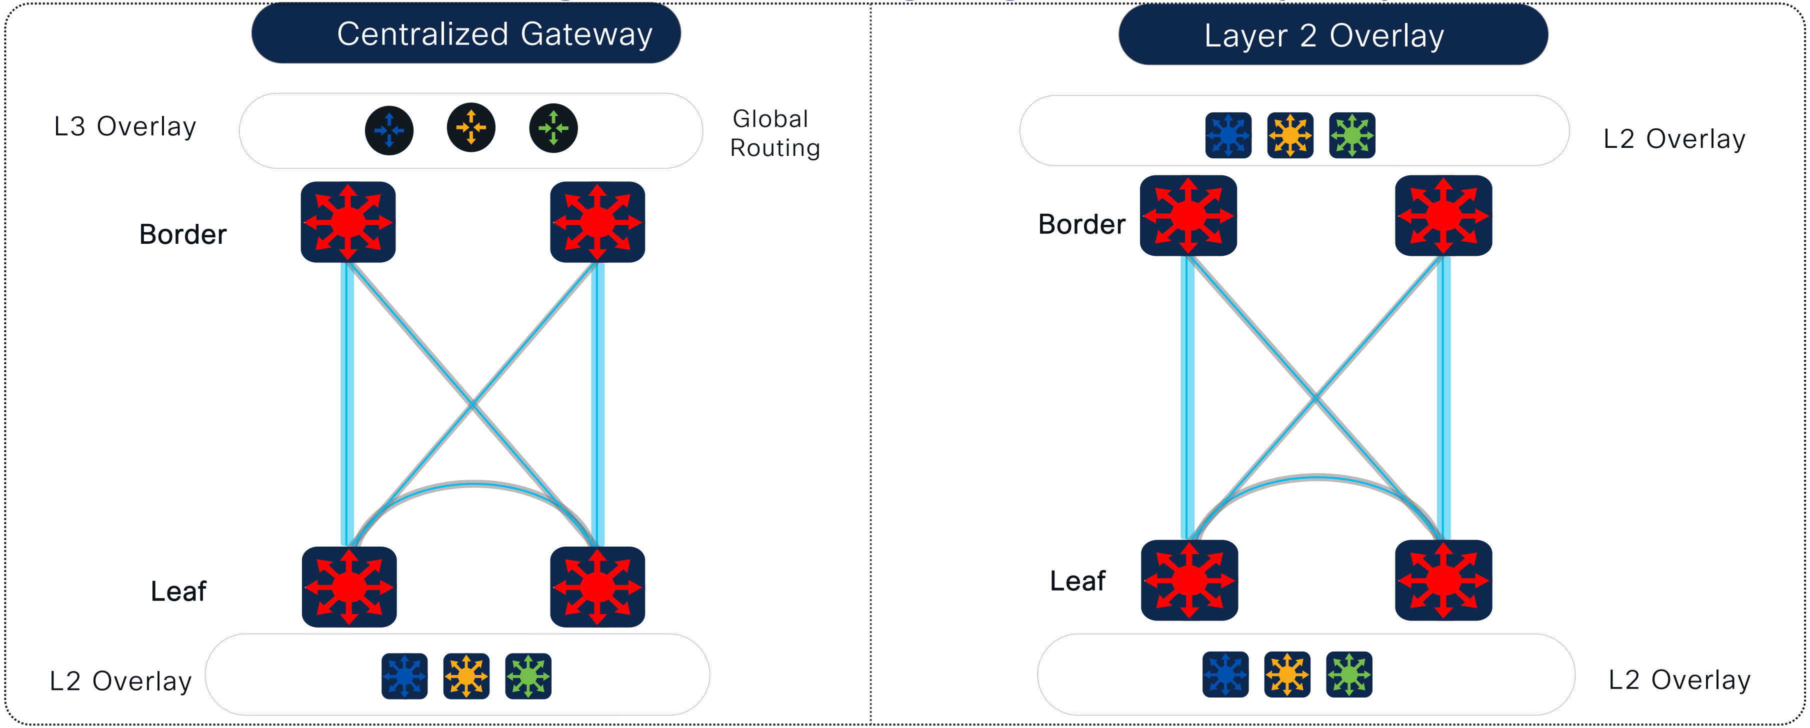

Overlay Types

- Four overlay network types support at any network layer point
- Route first. Bridge when-and-where need rule for scalable fabric architecture
- Feature rich Layer 3 overlay network support – Unicast | Multicast – IPv4 | IPv6
- Scalable Layer 2 overlay solution with suppression, flood management and more

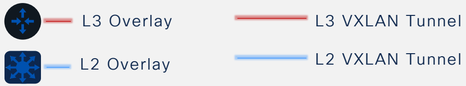

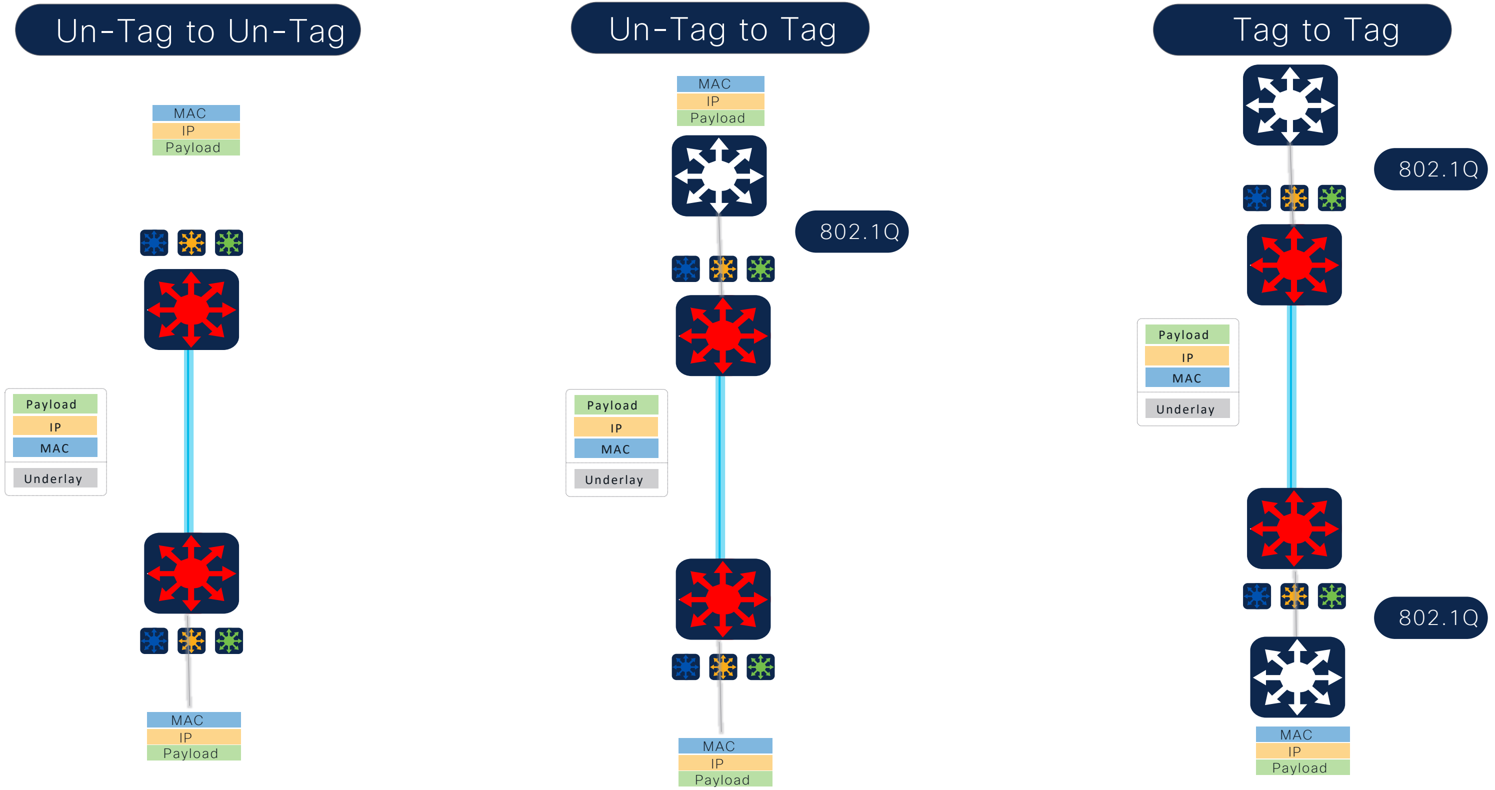

Adaptable Bridging

- Any to any Layer 2 overlay bridging across fabric
- Flexible solution to address Enterprise end-users to SP multi-tenant use-cases
- 1:1 or *n*:1 - Layer 2 VLAN to VNI mapping based on overlay transport requirements
- Maintains 802.1P QoS with option to optimize for enhanced application user-experience

Flexible Routing and Bridging Overlay Topologies

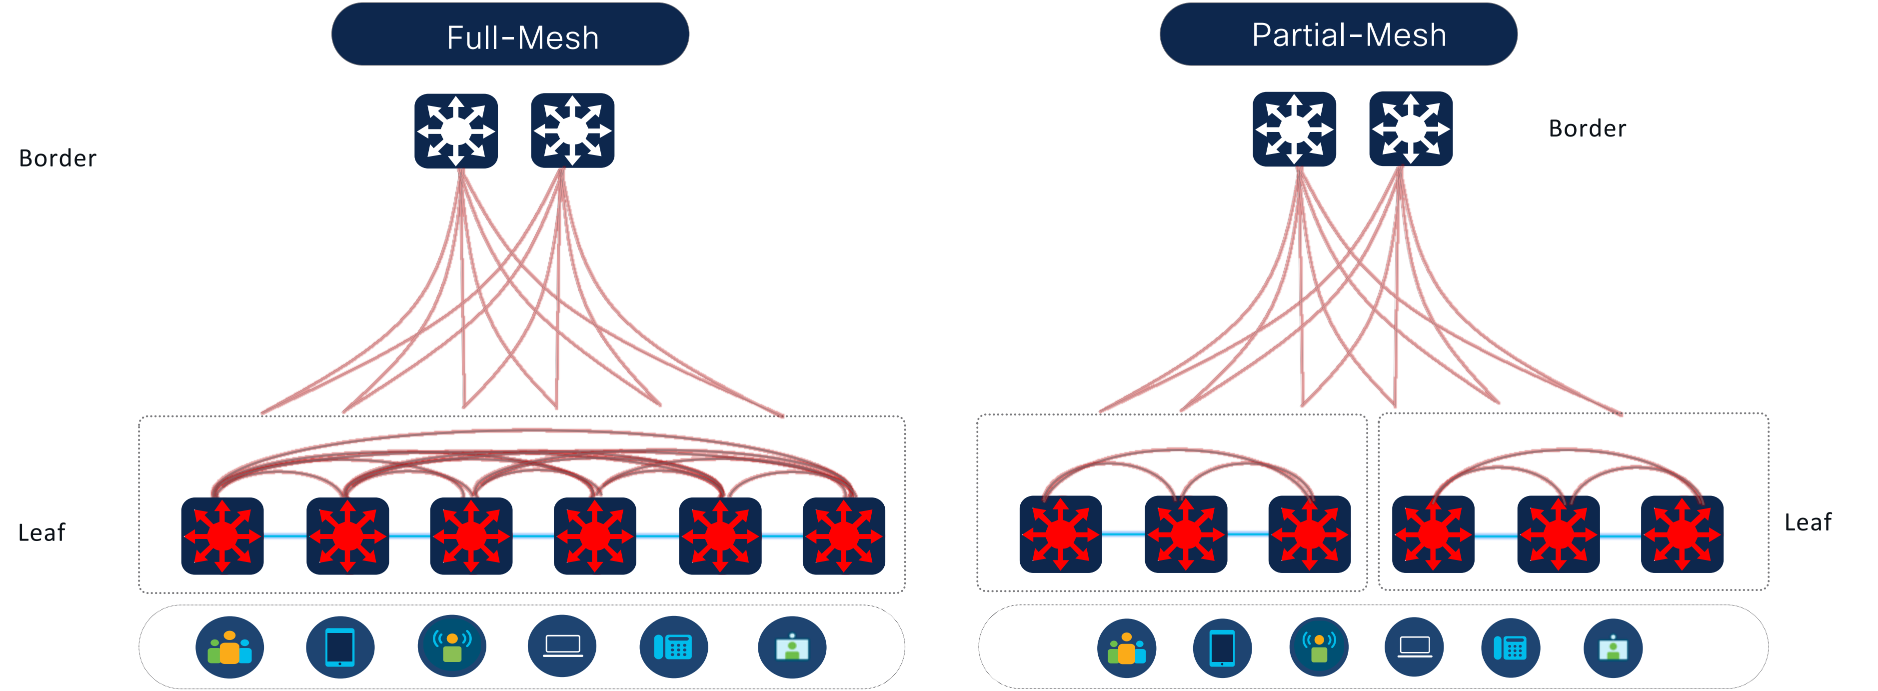

Segmentation Switch Group | Single Switch | VN | Subnet | VLAN | Port | App

Tailored Topologies

- Policy driven Layer 2 and Layer 3 overlay network topologies
- Simplified overlay network solution for broad Enterprise security use-cases
- Granular fabric overlay solution based on network access control policy
- Flexible central policy enforcement with external fabric domain

L3 VXLAN Tunnel
L2 VXLAN Tunnel

Flexible Routing and Bridging Overlay Topologies

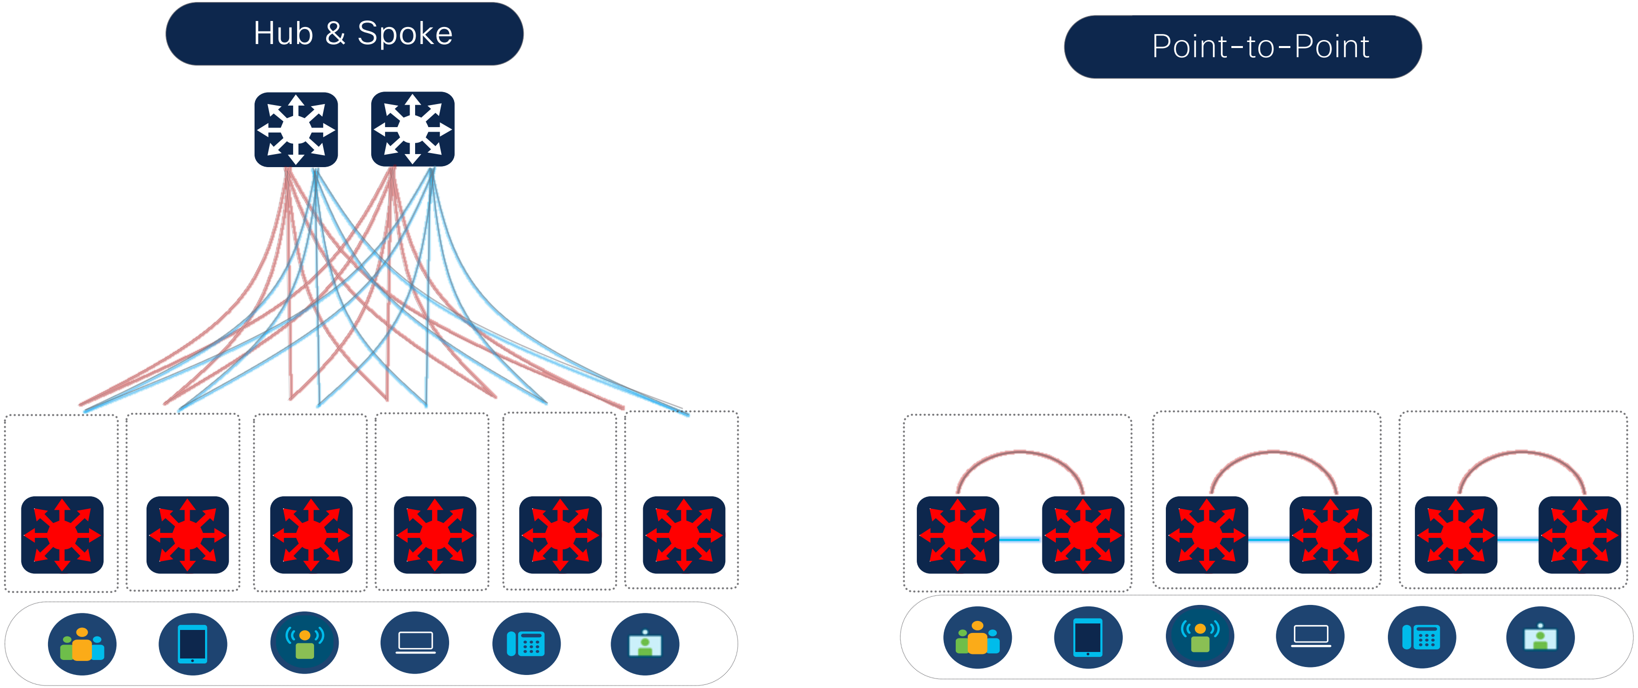

Tailored Topologies

- Policy driven Layer 2 and Layer 3 overlay network topologies
- Simplified overlay network solution for broad Enterprise security use-cases
- Granular fabric overlay solution based on network access control policy
- Flexible central policy enforcement with external fabric domain

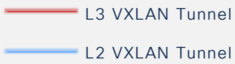

Efficient Layer 2 Broadcast domain

- 2 mechanics to handle **B**roadcast, **U**nknown Unicast and Link-Local **M**ulticast (BUM):
    - Ingress-Replication – Convert each BUM packet to multiple Unicast packets and transmit to each remote VTEP
    - Multicast-Replication – Convert each BUM packet to single Multicast packets and transmit in Underlay network
- Multicast replication offers significant system, network and end-user level performance benefits

Efficient Layer 2 Broadcast domain

Point-to-Point

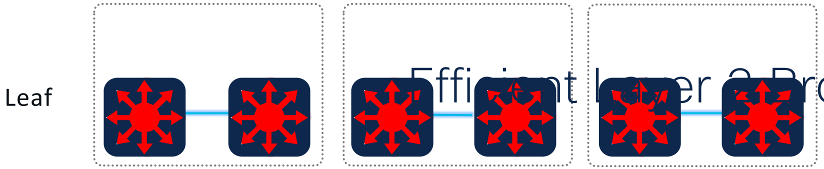

Efficient Layer 2 Broadcast domain

VLAN	VNI	BUM
401	40001	Ingress Replication
402	40002	
403	40003	

Hub-n-Spoke

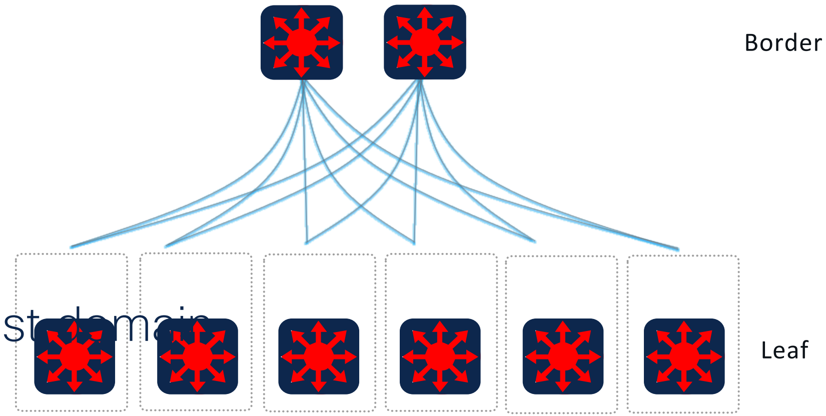

VLAN	VNI	BUM
301	30001	Ingress Replication
302	30002	
303	30003	

Scalable L2 BUM

- Per L2VNI BUM replication-type support. Deterministic BUM traffic management with BUM Rate-Limiter
- BUM replication-type selection based on Layer 2 overlay topology
- Controlled Multicast BUM based on broadcast domain boundary (*n* x L2VNI ID : 1 Multicast Group)
- Simplified Ingress-Replication for point-to-point Layer 2 overlay fabric

L3 VXLAN Tunnel
L2 VXLAN Tunnel

Efficient Layer 2 Broadcast domain

Full – Mesh

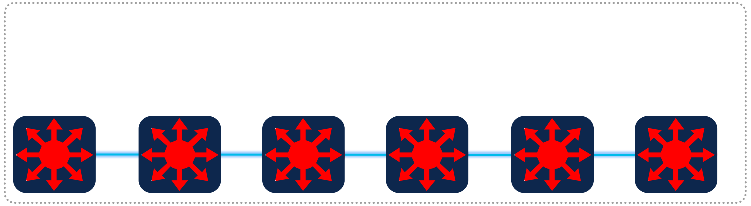

VLAN	VNI	BUM
101	10001	
102	10002	Multicast Group – 1
103	10003	

Partial – Mesh

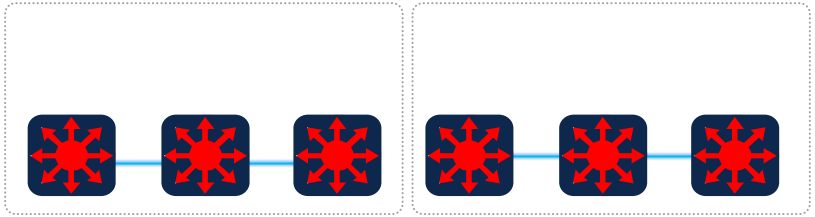

VLAN	VNI	Group
201	20001	
202	20002	Multicast Group – 1
203	20003	

VLAN	VNI	BUM
201	20001	
202	20002	Multicast Group – 2
203	20003	

Scalable L2 BUM

- Per L2VNI BUM replication-type support. Deterministic BUM traffic management with BUM Rate-Limiter
- BUM replication-type selection based on Layer 2 overlay topology
- Controlled Multicast BUM based on broadcast domain boundary (*n* x L2VNI ID : 1 Multicast Group)
- Simplified Ingress-Replication for point-to-point Layer 2 overlay fabric

L3 VXLAN Tunnel
L2 VXLAN Tunnel

Secure VN Extension

Design 1 – Stateful VN Extension

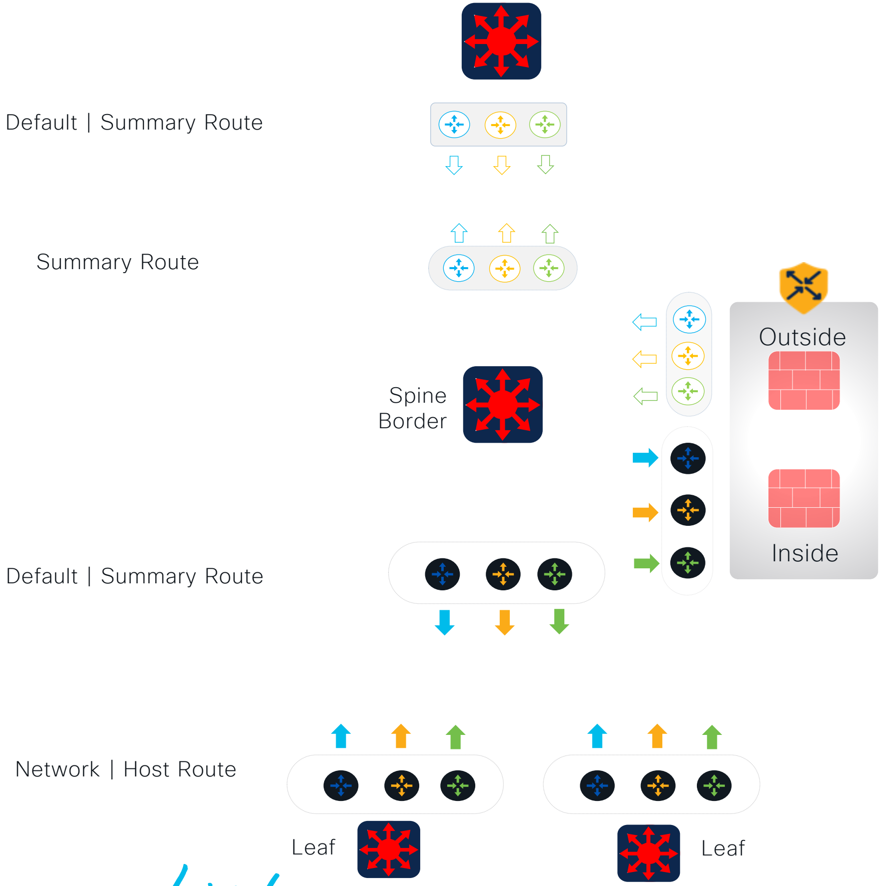

Secure Shared Extension

Design 2 – Stateful Global Extension

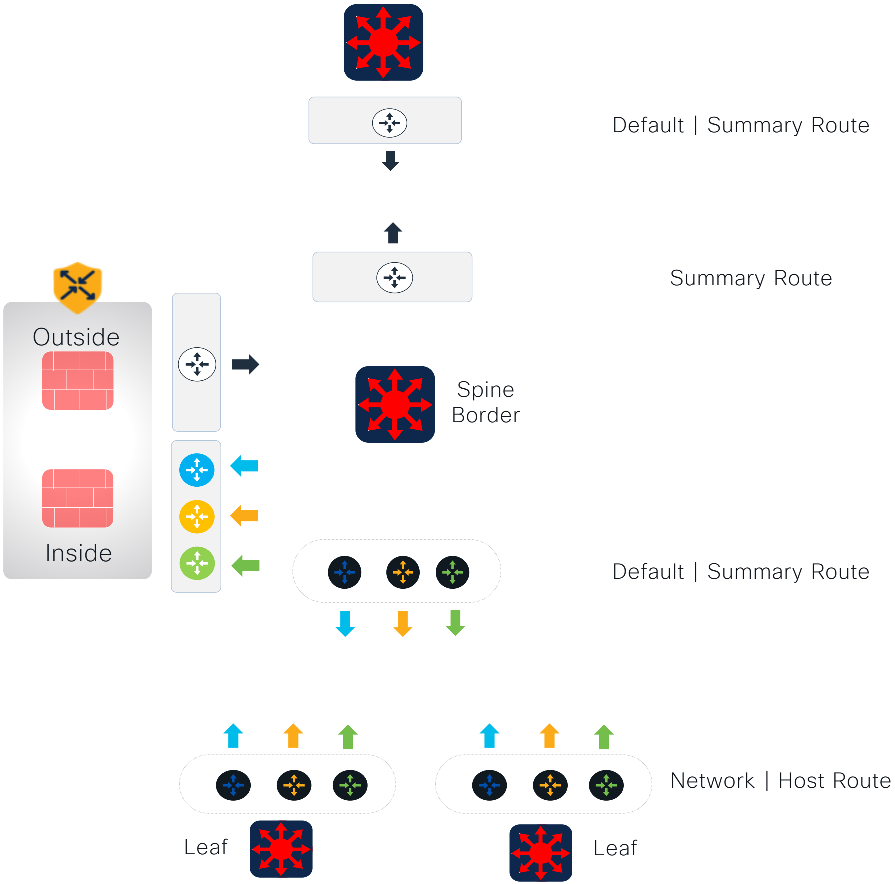

Multicast over VXLAN

Multicast Routing

Underlay Multicast

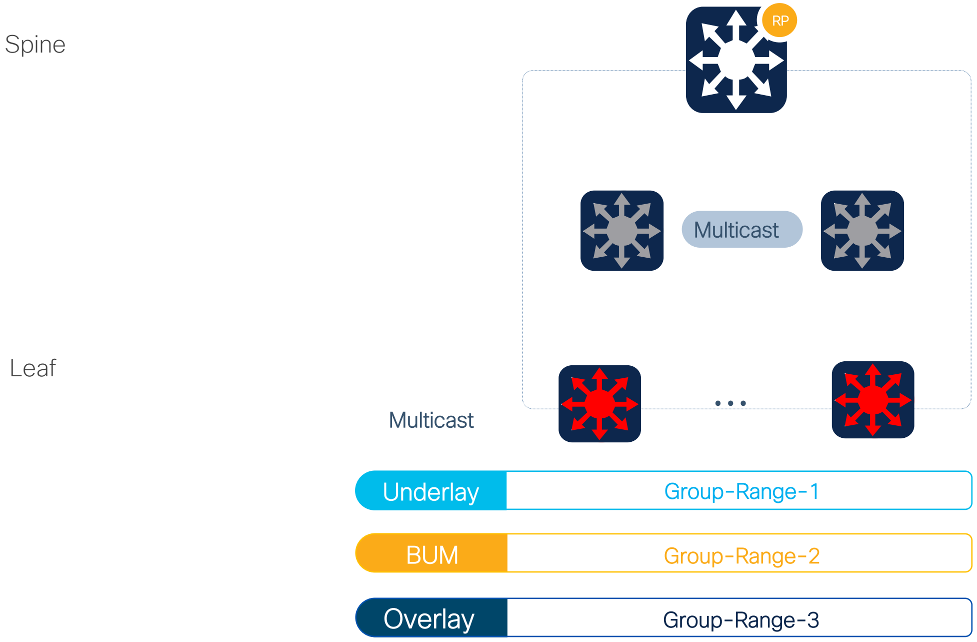

Multicast RP integrated on Spine or separate system

Non-overlapping Multicast Group for different purpose

Recommended to large scale EVPN deployments

Default MDT Group Range for Overlay TRM Multicast

Layer 3 Overlay

Distributed Anycast RP

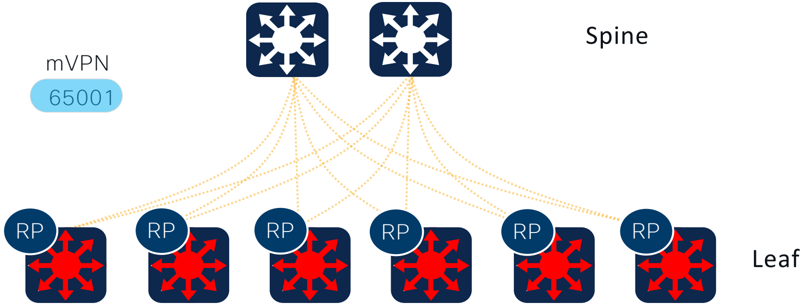

VRF	RP	IP	MDT
Blue	Anycast	Lo1: 10.1.1.101	239.1.1.101
Yellow	Anycast	Lo2: 10.2.1.101	239.2.1.101
Green	Anycast	Lo3: 10.3.1.101	239.3.1.101

Layer 3 Overlay

Fabric Border RP

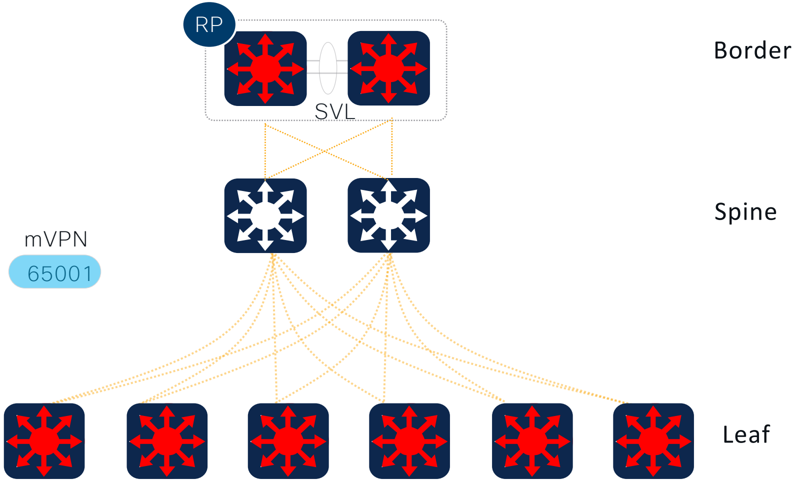

VRF	RP	IP	MDT
Blue	Anycast	Lo1: 10.1.1.101	239.1.1.101
Yellow	Anycast	Lo2: 10.2.1.101	239.2.1.101
Green	Anycast	Lo3: 10.3.1.101	239.3.1.101

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

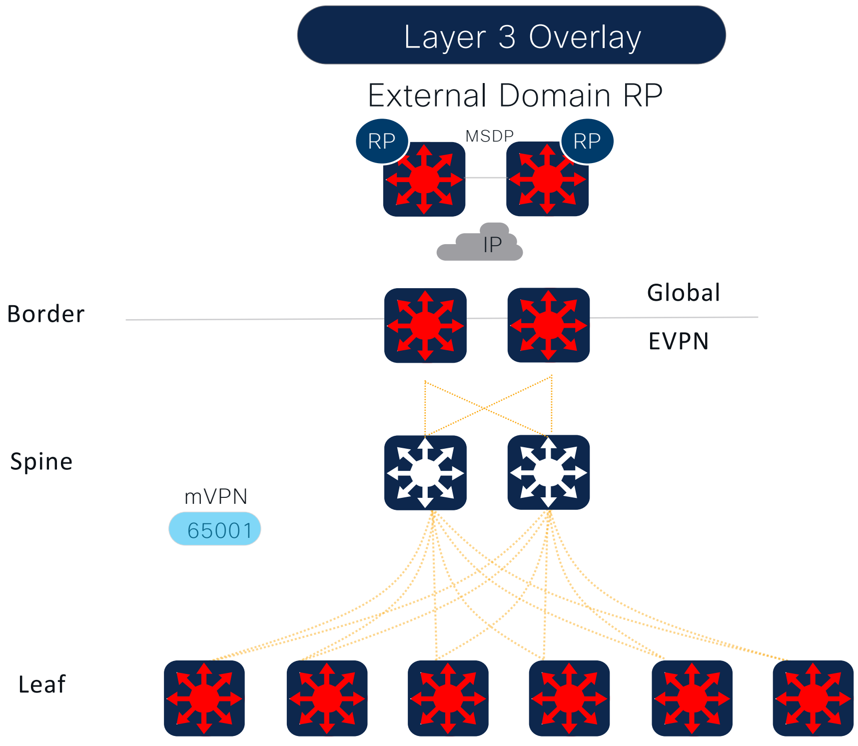

VRF	RP	IP	MDT
Blue	Anycast	Lo1: 10.1.1.101	239.1.1.101
Yellow	Anycast	Lo2: 10.2.1.101	239.2.1.101
Green	Anycast	Lo3: 10.3.1.101	239.3.1.101

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

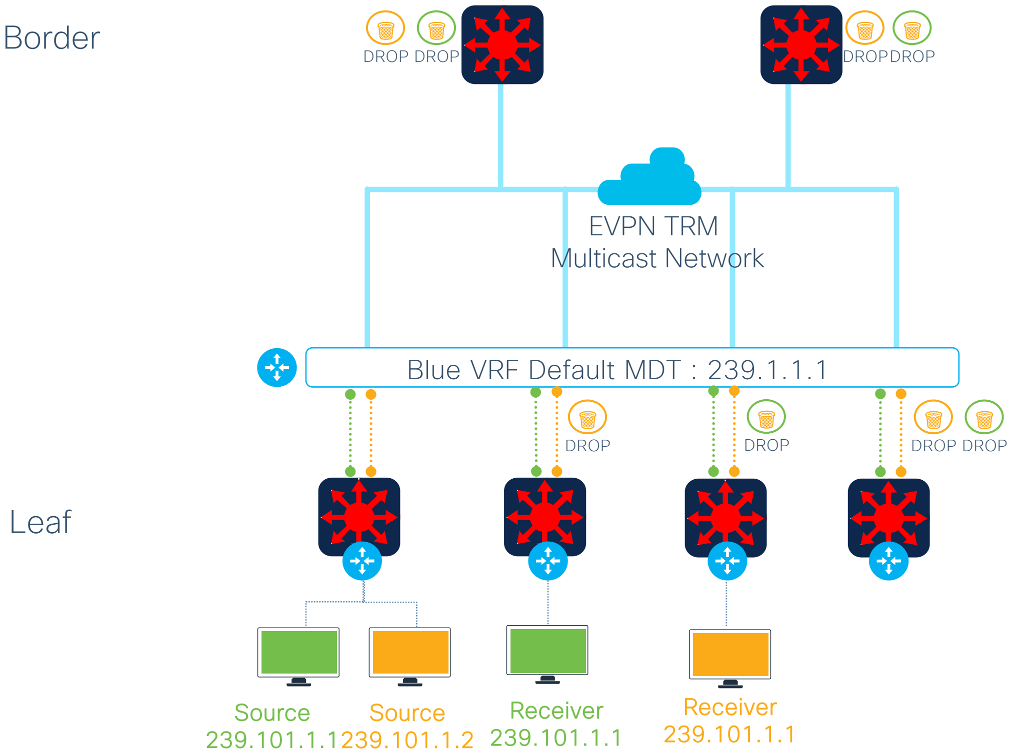

Challenges

Non-selective overlay Multicast replication

Inessential core network bandwidth utilization

Redundant system resources utilization

Limited scale for dense network environment

TRM Data MDT

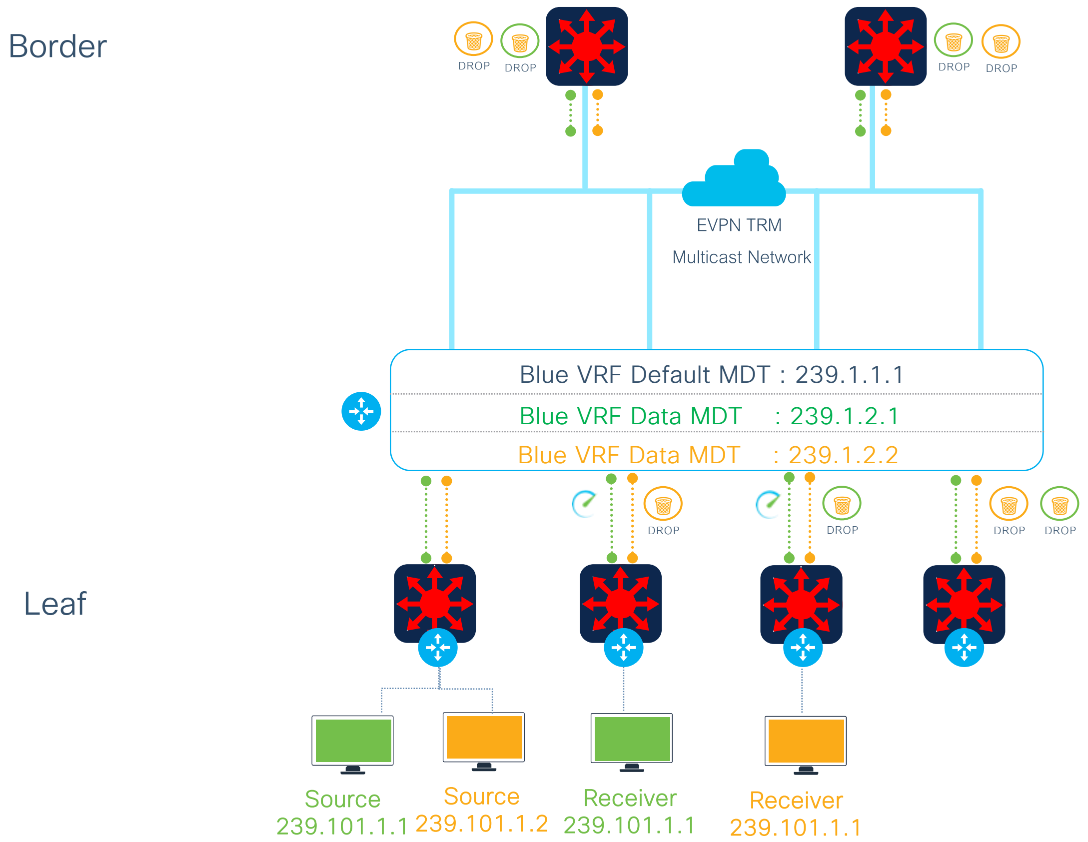

Key Benefits

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

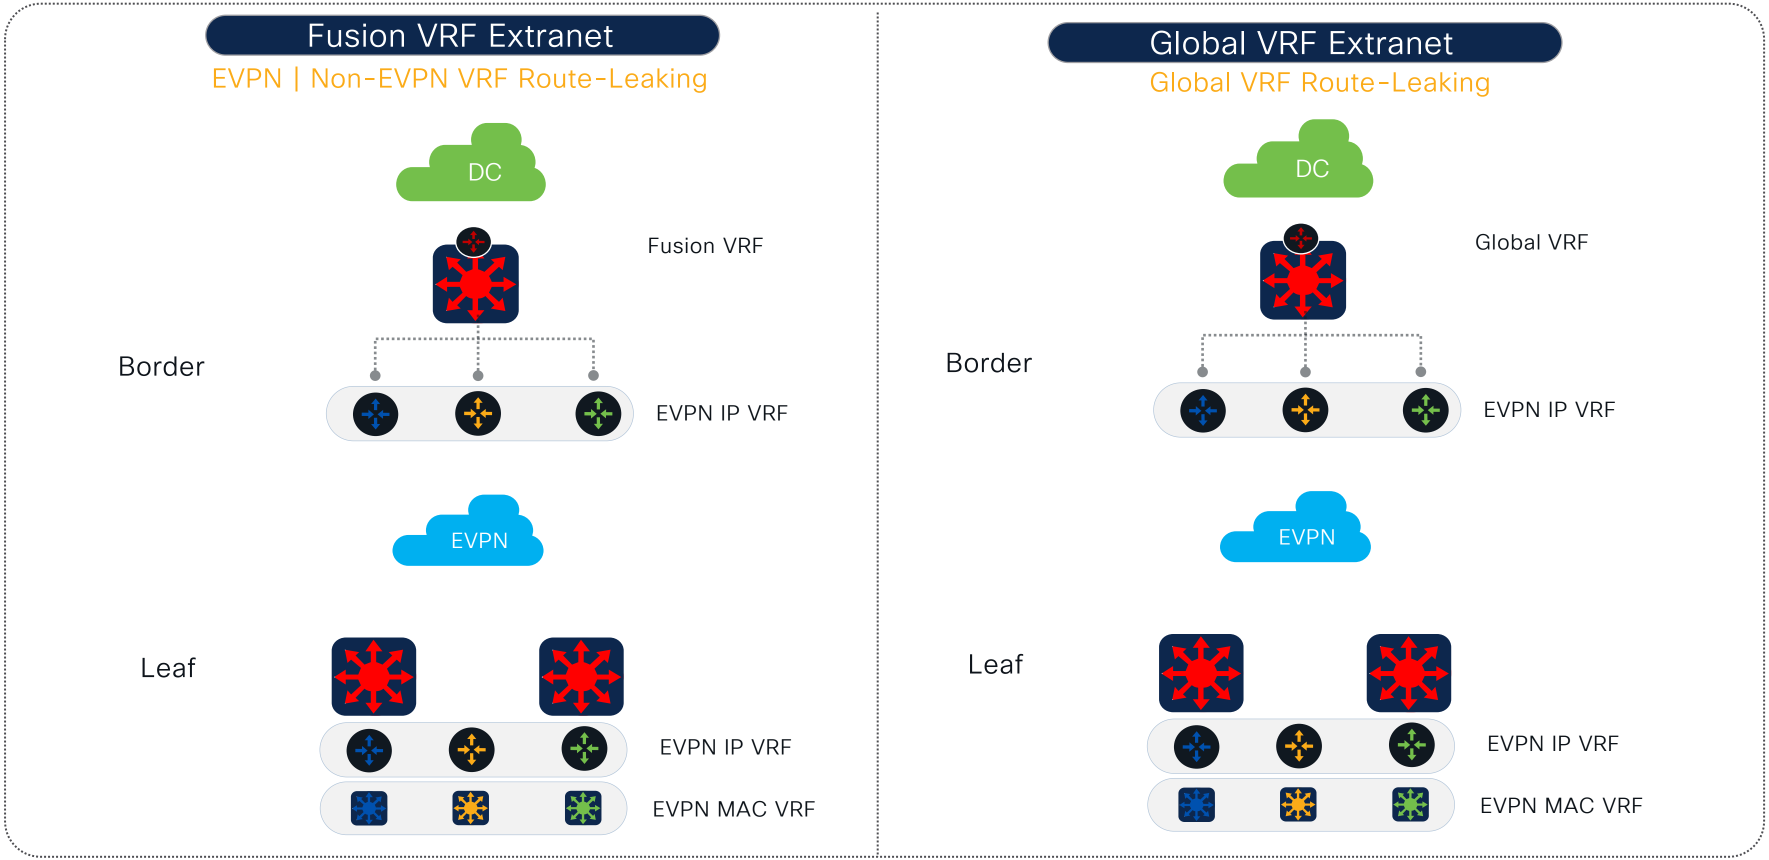

Integrated Extranet

Policy-based stateless extranet Unicast routing

Flexible route-leaking solution - EVPN-EVPN | EVPN-Non EVPN VRF | EVPN-Global

Various external Unicast routing protocol handoff

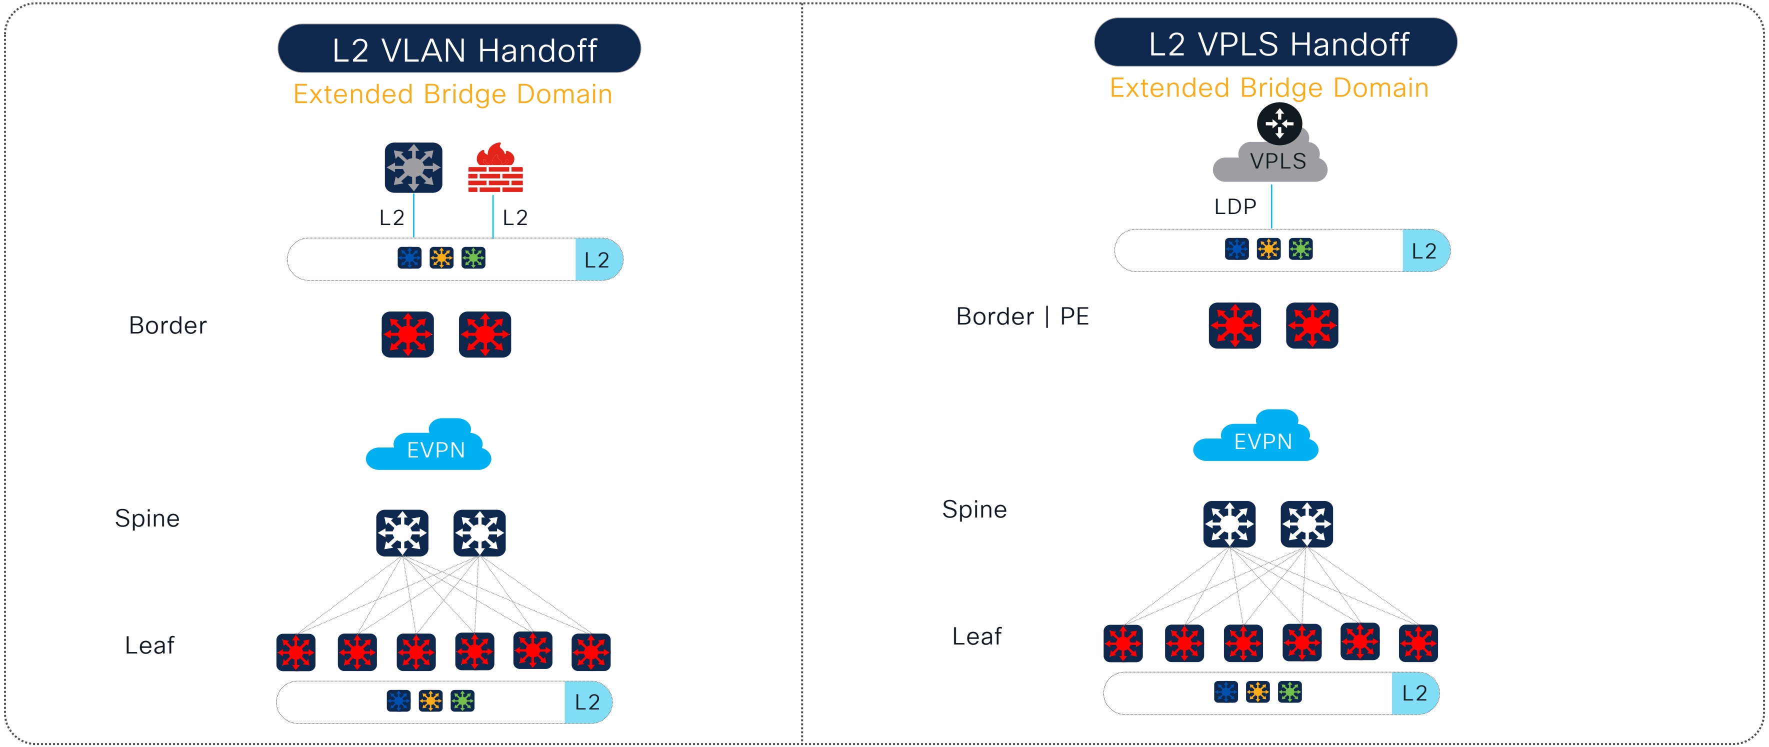

Integrated Extranet

- 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

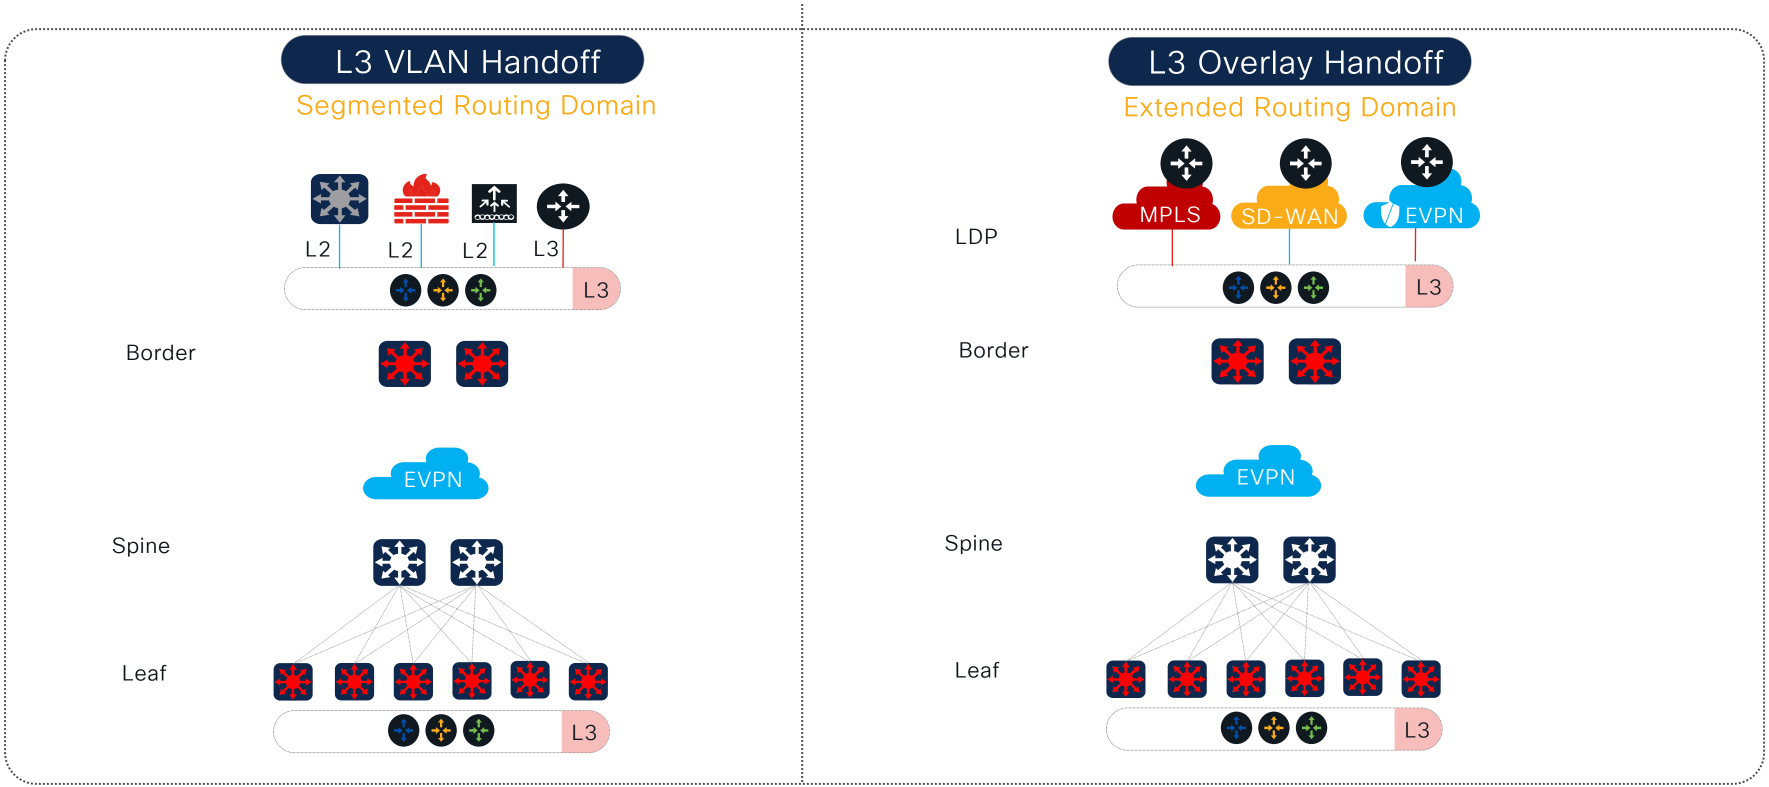

Integrated Extranet

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

Secure Encrypted Fabric

Single Cluster – Non-Hierarchical Fabric

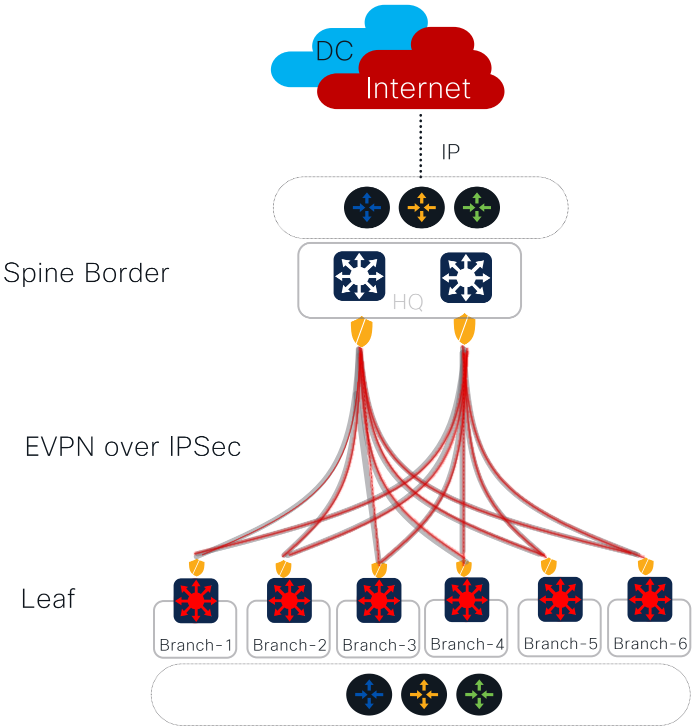

Secure Encrypted Fabric

Multi Cluster – Hierarchical Fabric

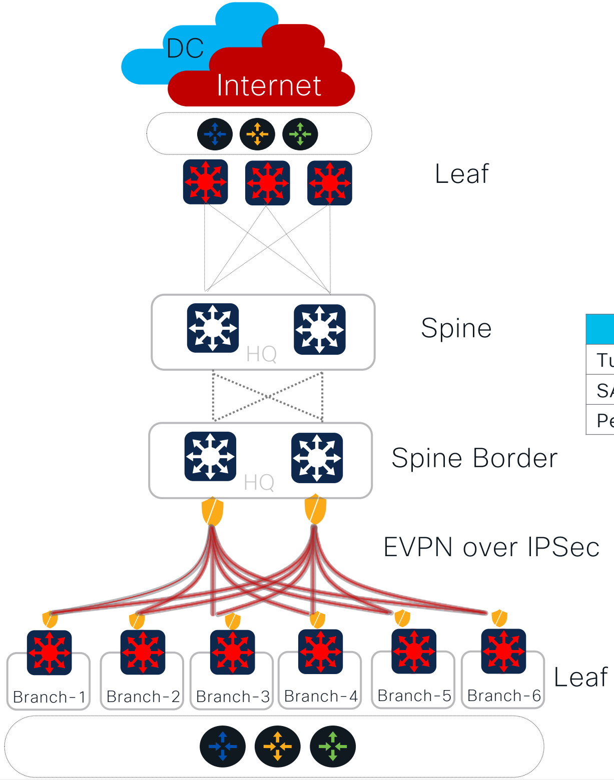

Scale and Performance Matrix

IPSec		EVPN	
Tunnel	128	Peers	128
SA	256	VRF | L3VNI	256
Performance	100Gbps	Unicast Prefix	39000

Encrypted EVPN Fabric

- High performance Catalyst 9300-X/9400X IPsec underlay network solution
- Simplified and scalable Layer 3 overlay fabric with integrated or co-located Spine/RR
- Single fabric cluster across WAN or “stitch” to EVPN fabric at central-office
- Unicast | Multicast support for IPv4 and IPv6 in overlay

BGP-EVPN Overlay's with IPv6 Underlay

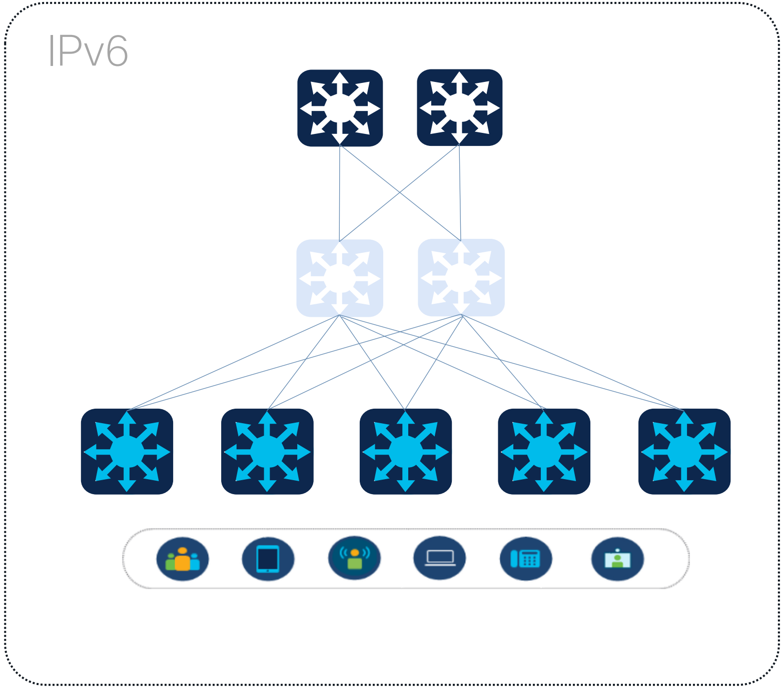

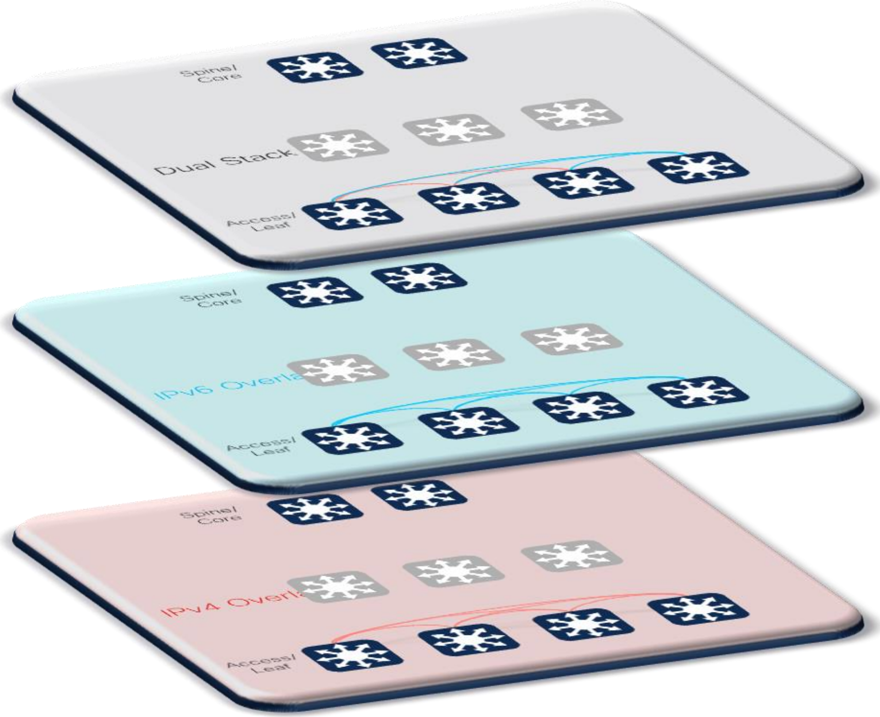

Seamless IPv6 Migration

Ingress Replication
Overlay-v4 Unicast
Overlay-v6 Unicast

Multicast Replication
Overlay-v4 Unicast
Overlay-v6 Unicast

Dual-Stack
V6 peers only
Mixed: V4/V6 and Dual Stack VTEP's
All Dual Stack VTEPS

BGP-EVPN Seamless Migration to IPv6 Underlay

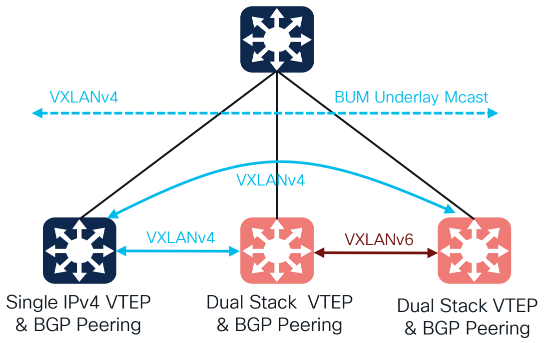

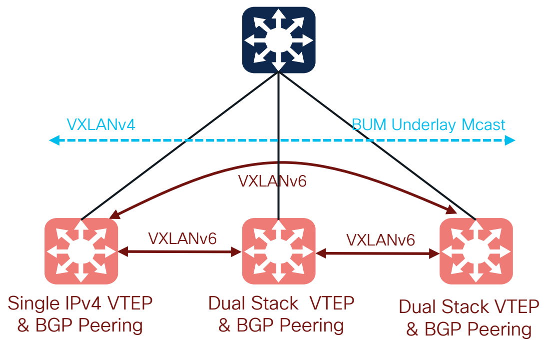

- Unicast traffic and BUM traffic with IR, dual stack VTEP will communicate with other IPv4 VTEP;s using VXLANv4 and with other Dual Stack VTEP using VXLANv6
- BUM traffic with Multicast Replication: Underlay Multicast IPv4 is used

```
interface Loopback0
 ip address 172.168.26.1 255.255.255.255
 ip ospf 100 area 0
 ipv6 address 2006:1::2/128
 ipv6 ospf 1 area 0
!
 interface nve1
 no ip address
 load-interval 30
 source-interface Loopback0
 host-reachability protocol bgp
 vxlan encapsulation dual-stack prefer-ipv4 underlay-mcast ipv4
 member vni 111110 ingress-replication
 member vni 222220 mcast-group 232.10.10.10 ff0e::12
 member vni 110110 vrf CORP
 member vni 220220 vrf GUEST
```

BGP-EVPN Seamless Migration to IPv6 Underlay

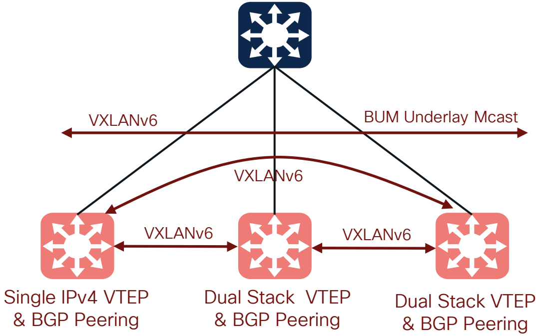


```
interface Loopback0
 ip address 172.168.26.1 255.255.255.255
 ip ospf 100 area 0
 ipv6 address 2006:1::2/128
 ipv6 ospf 1 area 0
!
 interface nve1
 no ip address
 load-interval 30
 source-interface Loopback0
 host-reachability protocol bgp
 vxlan encapsulation dual-stack prefer-ipv4 underlay-mcast ipv6
 member vni 111110 ingress-replication
 member vni 222220 mcast-group 232.10.10.10 ff0e::12
 member vni 110110 vrf CORP
 member vni 220220 vrf GUEST
```

- Once all VTPE’s are Dual-stack capable, the underlay Mcast can be switched to ipv6 Multicast

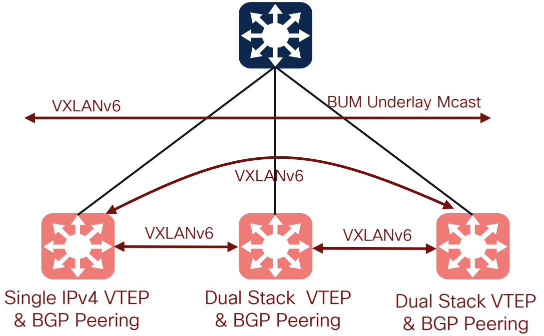


```
interface Loopback0
ipv6 address 2006:1::2/128
 ipv6 ospf 1 area 0
!
 interface nve1
 no ip address
 load-interval 30
 source-interface Loopback0
 host-reachability protocol bgp
 vxlan encapsulation ipv6
 member vni 111110 ingress-replication
 member vni 222220 mcast-group ff0e::12
 member vni 110110 vrf CORP
 member vni 220220 vrf GUEST
```

- Dual Stack VTEP migration to complete VXLANv6

Fabric Segmentation Options

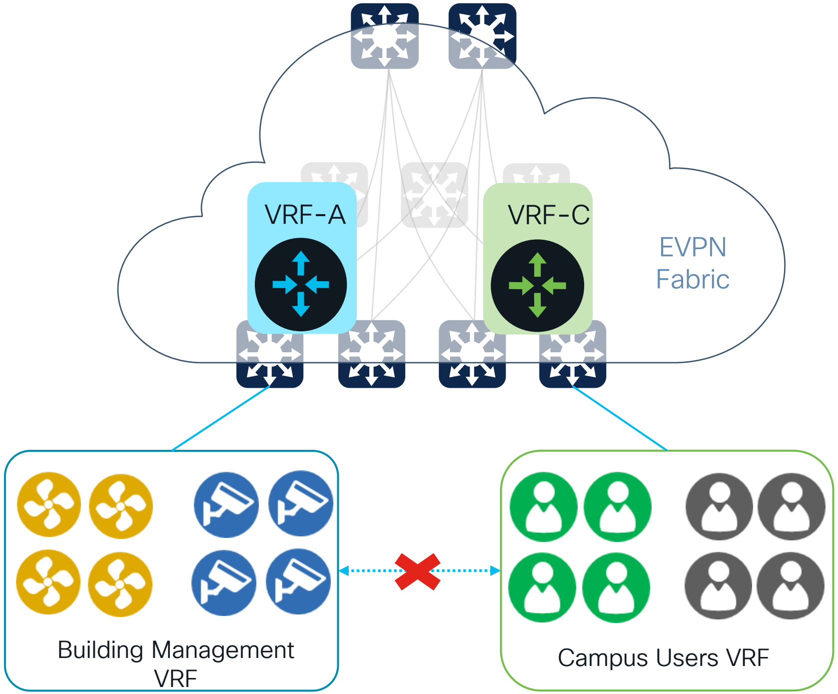

Macro Segmentation: No communication between VRF’s

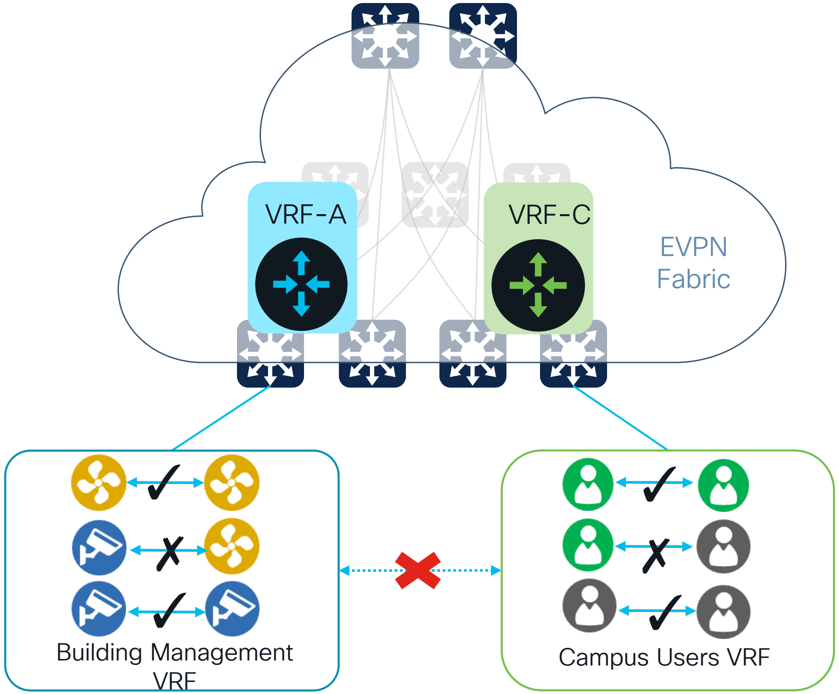

Micro Segmentation: Second level Segmentation between groups within a VRF

Software
Cisco IOS® XE 17.11.1

BGP EVPN – Role based Access Control

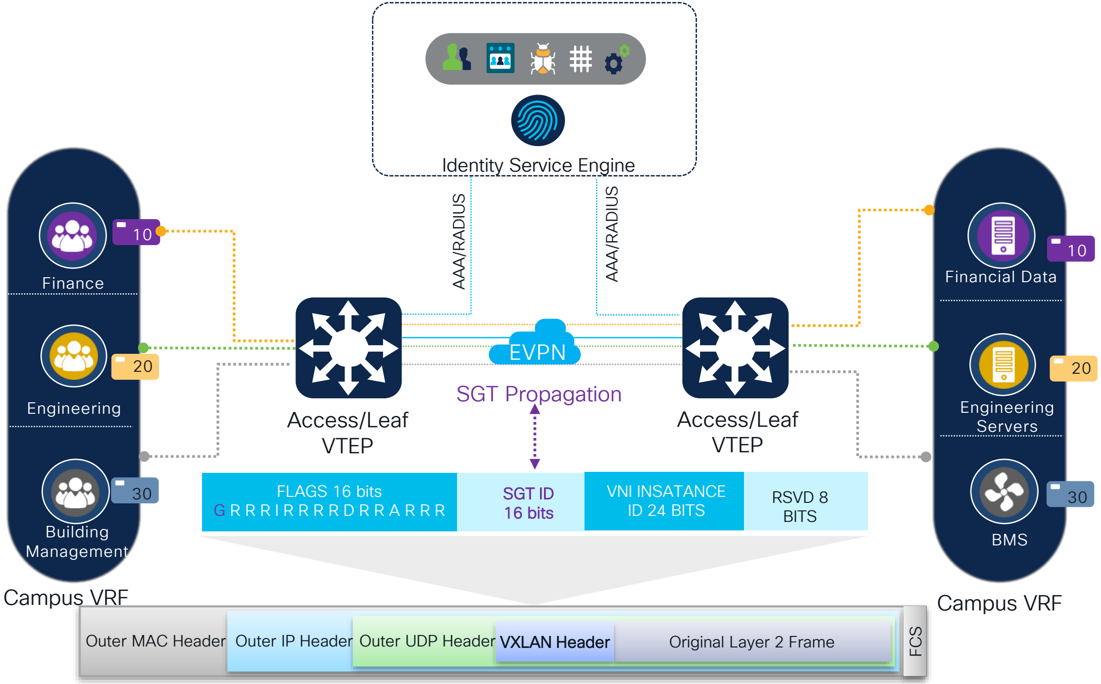

- Role Based Access Control
- Scalable policy based on User role
  Dynamic or Static Policy enforcement
- Centralized Policy Management for Dynamic policy provisioning

Automation

IOS XE Programmability

The NETCONF, RETCONF and gNMI are programmatic interfaces that provide additional methods for interfacing with the IOS XE device

YANG data models define the data that is available for configuration and streaming telemetry

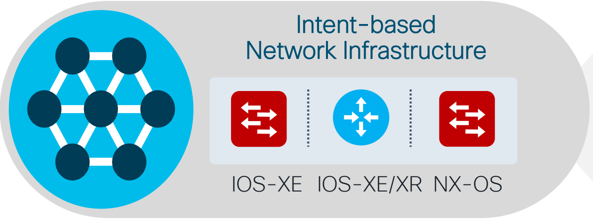

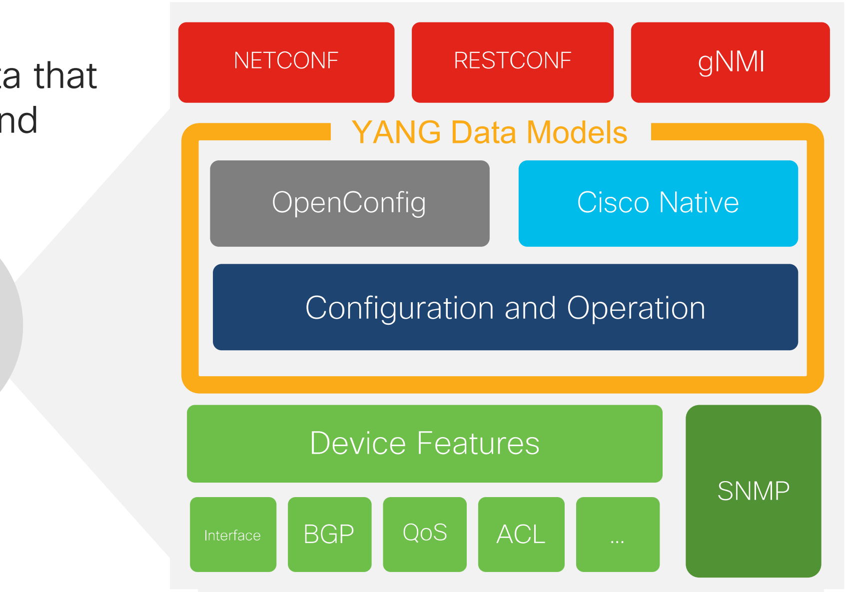

	Underlay IP	Underlay IP Routing	Underlay Multicast	Spine	PIM RP	BGP L2VPN ASN	MAC VRF ID	IP VRF RD	BUM	MAC VRF RT Import	MAC VRF RT Export	IP VRF RT Import	IP VRF RT Export	Layer 3 Multicast	Border Handoff
Layer 3 Routed Overlay Network															
Design-1	IPv4	OSPFv2	PIM-SM	RR	RR	iBGP	N/A	<RID:VLAN>	N/A	N/A	N/A	<ASN:VLAN>	<ASN:VLAN>	N/A	Layer 3 VRF
Design-2	IPv4 | IPv6	OSPFv2 | v3	PIM-SM	RR	RR	iBGP	N/A	<RID:VLAN>	N/A	N/A	N/A	<ASN:VLAN>	<ASN:VLAN>	N/A	Layer 3 VRF
Design-3	IPv4	OSPFv2	PIM-SM	RR	RR	iBGP	N/A	<RID:VLAN>	N/A	N/A	N/A	<ASN:VLAN>	<ASN:VLAN>	TRMv4	Layer 3 VRF
Design-4	IPv4 | IPv6	OSPFv2 | v3	PIM-SM	RR	RR	iBGP	N/A	<RID:VLAN>	N/A	N/A	N/A	<ASN:VLAN>	<ASN:VLAN>	TRMv4 | v6	Layer 3 VRF
Design-5	IPv4	OSPFv2	PIM-SM	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	<ASN:VLAN>	N/A	Layer 3 VRF
Design-6	IPv4 | IPv6	OSPFv2 | v3	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	<ASN:VLAN>	N/A	Layer 3 VRF
Design-7	IPv4	OSPFv2	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	<ASN:VLAN>	TRMv4	Layer 3 VRF
Design-8	IPv4 | IPv6	OSPFv2 | v3	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	<ASN:VLAN>	TRMv4 | v6	Layer 3 VRF
Design-9	IPv4	OSPFv2	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	<ASN:VLAN>	N/A	Layer 3 VRF
Design-10	IPv4 | IPv6	OSPFv2 | v3	PIM-SM	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	<ASN:VLAN>	N/A	Layer 3 VRF
Design-11	IPv4	OSPFv2	PIM-SM	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	<ASN:VLAN>	<ASN:VLAN>	TRMv4	Layer 3 VRF
Design-12	IPv4 | IPv6	OSPFv2 | v3	PIM-SM	RR	RR	iBGP	N/A	<RID:VLAN>	Multicast	<RID:VLAN>	<RID:VLAN>	<ASN:VLAN>	<ASN:VLAN>	TRMv4 | v6	Layer 3 VRF
Layer 2 Cross-Connect Overlay Network															
Design-13	IPv4	OSPFv2	PIM-SM	RR	RR	iBGP	N/A	N/A	Multicast	<RID:VLAN>	<RID:VLAN>	N/A	N/A	N/A	Layer 3 VRF
Design-14	IPv4 | IPv6	OSPFv2 | v3	PIM-SM	RR	RR	iBGP	N/A	N/A	Multicast	<RID:VLAN>	<RID:VLAN>	N/A	N/A	N/A	Layer 3 VRF
Design-15	IPv4	OSPFv2	PIM-SM	RR	RR	iBGP	N/A	N/A	Multicast	<RID:VLAN>	<RID:VLAN>	N/A	N/A	N/A	Layer 3 VRF
Design-16	IPv4 | IPv6	OSPFv2 | v3	PIM-SM	RR	RR	iBGP	N/A	N/A	Multicast	<RID:VLAN>	<RID:VLAN>	N/A	N/A	N/A	Layer 3 VRF

Design-Based Playbooks

https://github.com/cisco

EVPN Ansible – Solution Playbook

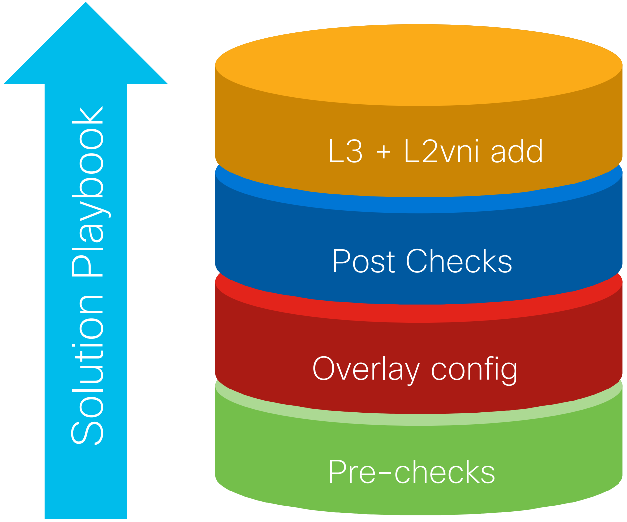

Same playbook to add L3/L2 VNI's
Eg: Add one or multiple L3/l2vni using same playbook

Framework for post-check
Eg: BGP status up/down, overlay ping checks

Solution level deployment
Eg: Ipv4 + Ipv6 + TRM in a single playbook

Framework for pre-checks
Eg: License check, underlay reachability check

Simple to Use

- Single playbook for complete solution
- Single inventory file to add Leaf/Spine variables

DEMO

BGP EVPN Ansible Automation

```yaml
- name: SHOW IP BGP L2VPN EVPN SUMMARY
  any_errors_fatal: true
  cisco.ios.ios_command:
    commands:
      - "show ip bgp l2vpn evpn summary | begi  Up/Dow "
  register: after_show_bgp
#
- name: DEBUG
  ansible.builtin.debug:
    msg:
      - "{{ after_show_bgp }}"
#
- name: "ASSERT SUCCESS OR FAILURE BASED ( :00:00 OR NOT NEVER AND NEIGHBOORS LEAFS LISTED CORRECTLY"
  any_errors_fatal: true
  assert:
    that:
      - after_show_bgp.stdout
      - after_show_bgp.stdout
      - after_show_bgp
    fail_msg: "--->
    success_msg: "
    quiet: no
  with_items:
    - "{{ groups['leafs'] }}"
  register: after_assert
#
- name: DEBUG
  ansible.builtin.debug:
    msg:
      - "{{ after_assert }}"
```

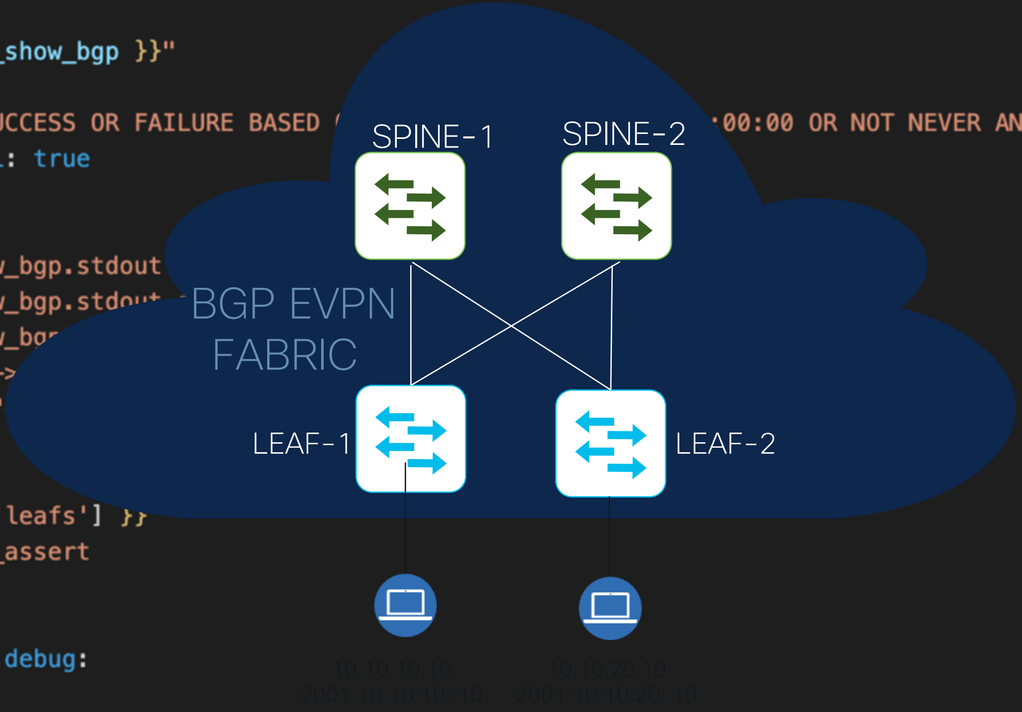

ssh netadmin@172.26.193.242

netadmin@172.26.193.242 (ssh)

```
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
SPINE-1#
```

netadmin@172.26.193.243 (ssh)

```
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
SPINE-2#
```

netadmin@172.26.193.240 (ssh)

```
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
LEAF-1#
```

netadmin@172.26.193.241 (ssh)

```
[OK]
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
LEAF-2#
```

raj@raj-virtual-machine: ~ (ssh)

```
From 10.10.10.10 icmp_seq=745 Destination Host Unreachable
From 10.10.10.10 icmp_seq=746 Destination Host Unreachable
From 10.10.10.10 icmp_seq=747 Destination Host Unreachable
From 10.10.10.10 icmp_seq=748 Destination Host Unreachable
From 10.10.10.10 icmp_seq=749 Destination Host Unreachable
From 10.10.10.10 icmp_seq=750 Destination Host Unreachable
From 10.10.10.10 icmp_seq=751 Destination Host Unreachable
From 10.10.10.10 icmp_seq=752 Destination Host Unreachable
From 10.10.10.10 icmp_seq=753 Destination Host Unreachable
From 10.10.10.10 icmp_seq=754 Destination Host Unreachable
From 10.10.10.10 icmp_seq=755 Destination Host Unreachable
From 10.10.10.10 icmp_seq=756 Destination Host Unreachable
```

raj@raj-ubuntu-2: ~ (ssh)

```
From 10.10.20.10 icmp_seq=744 Destination Host Unreachable
From 10.10.20.10 icmp_seq=745 Destination Host Unreachable
From 10.10.20.10 icmp_seq=746 Destination Host Unreachable
From 10.10.20.10 icmp_seq=747 Destination Host Unreachable
From 10.10.20.10 icmp_seq=748 Destination Host Unreachable
From 10.10.20.10 icmp_seq=749 Destination Host Unreachable
From 10.10.20.10 icmp_seq=750 Destination Host Unreachable
From 10.10.20.10 icmp_seq=751 Destination Host Unreachable
From 10.10.20.10 icmp_seq=752 Destination Host Unreachable
From 10.10.20.10 icmp_seq=753 Destination Host Unreachable
From 10.10.20.10 icmp_seq=754 Destination Host Unreachable
From 10.10.20.10 icmp_seq=755 Destination Host Unreachable
```

raj@raj-virtual-machine: ~ (ssh)

```
From 2001:10:10:10::10 icmp_seq=740 Destination unreachable: Address unreachable
From 2001:10:10:10::10 icmp_seq=741 Destination unreachable: Address unreachable
From 2001:10:10:10::10 icmp_seq=742 Destination unreachable: Address unreachable
From 2001:10:10:10::10 icmp_seq=743 Destination unreachable: Address unreachable
From 2001:10:10:10::10 icmp_seq=744 Destination unreachable: Address unreachable
From 2001:10:10:10::10 icmp_seq=745 Destination unreachable: Address unreachable
From 2001:10:10:10::10 icmp_seq=746 Destination unreachable: Address unreachable
From 2001:10:10:10::10 icmp_seq=747 Destination unreachable: Address unreachable
From 2001:10:10:10::10 icmp_seq=748 Destination unreachable: Address unreachable
```

raj@raj-ubuntu-2: ~ (ssh)

```
From 2001:10:10:20::10 icmp_seq=738 Destination unreachable: Address unreachable
From 2001:10:10:20::10 icmp_seq=739 Destination unreachable: Address unreachable
From 2001:10:10:20::10 icmp_seq=740 Destination unreachable: Address unreachable
From 2001:10:10:20::10 icmp_seq=741 Destination unreachable: Address unreachable
From 2001:10:10:20::10 icmp_seq=742 Destination unreachable: Address unreachable
From 2001:10:10:20::10 icmp_seq=743 Destination unreachable: Address unreachable
From 2001:10:10:20::10 icmp_seq=744 Destination unreachable: Address unreachable
From 2001:10:10:20::10 icmp_seq=745 Destination unreachable: Address unreachable
From 2001:10:10:20::10 icmp_seq=746 Destination unreachable: Address unreachable
```

inventory_raj.yml — AUTOMATION

inventory_raj.yml .../2 | l3_edge_vlan_ipv4_ipv6_inner_loop.yml | invent

ANSIBLE > Ansible-05-02 > bgp-evpn-ansible-main > example_use_cases > l3_loop > 2 > inve

```yaml
---
#
leafs:
#
  hosts:
#
    sw-access-leaf-1:
      ansible_host: 172.26.193.240
      role: leaf
      rid: 172.168.1.1
      vrfs:
        - vrf_name: VRF_OVERLAY_2010
          vrf_core_vlan_id: 2010
          vrf_l3_vni_id: 20010
          vrf_edge_vlan:
            - id: 10
              ip: 10.10.10.1
              mask: 255.255.255.0
              ipv6: 2001:10:10:10::1
              prefix: 64
            - id: 20
              ip: 10.20.10.1
              mask: 255.255.255.0
              ipv6: 2001:20:10:10::1
```

PROBLEMS OUTPUT DEBUG CONSOLE TERMINAL zsh - 2

```
→ 2
```

EVPN Ansible – Feature level Playbook

playbook_access_add_preview.yml	adding L2VNI and L3VNI
playbook_access_incremental_commit.yml	initial commit for release/2.x.x
playbook_access_incremental_preview.yml	initial commit for release/2.x.x
playbook_cleanup.yml	initial commit for release/2.x.x
playbook_dhcp_add_commit.yml	adding L2VNI and L3VNI
playbook_dhcp_add_preview.yml	adding L2VNI and L3VNI
playbook_dhcp_delete_commit.yml	dhcp incremental commit
playbook_dhcp_delete_preview.yml	dhcp incremental commit
playbook_output.yml	fix playbook_output
playbook_overlay_commit.yml	adding L2VNI and L3VNI
playbook_overlay_delete_commit.yml	ipv6_incremental
playbook_overlay_delete_generate.yml	initial commit for release/2.x.x
playbook_overlay_delete_ipv6_commit.yml	adding L2VNI and L3VNI
playbook_overlay_delete_ipv6_generate.yml	adding L2VNI and L3VNI
playbook_overlay_delete_ipv6_preview.yml	adding L2VNI and L3VNI
playbook_overlay_delete_preview.yml	initial commit for release/2.x.x
playbook_overlay_incremental_commit.yml	adding L2VNI and L3VNI
playbook_overlay_incremental_generate.yml	adding L2VNI and L3VNI
playbook_overlay_incremental_ipv6_commit.yml	adding L2VNI and L3VNI
playbook_overlay_incremental_ipv6_generate.yml	adding L2VNI and L3VNI
playbook_overlay_incremental_ipv6_preview.yml	ipv6_incremental
playbook_overlay_incremental_preview.yml	adding L2VNI and L3VNI
playbook_overlay_precheck.yml	initial commit for release/2.x.x
playbook_overlay_preview.yml	adding L2VNI and L3VNI

Feature specific Playbooks

↓

Add/remove a feature

↓ Get playbooks below

Ansible Playbooks

EVPN Automation with Terraform

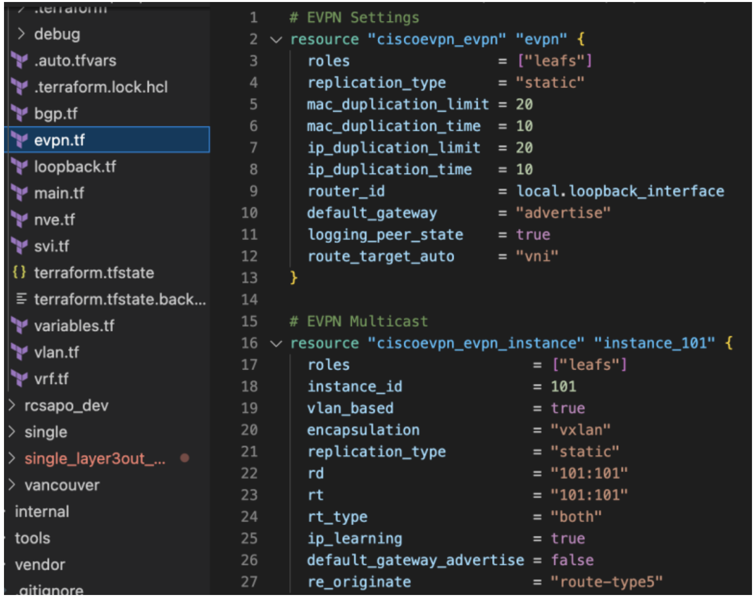

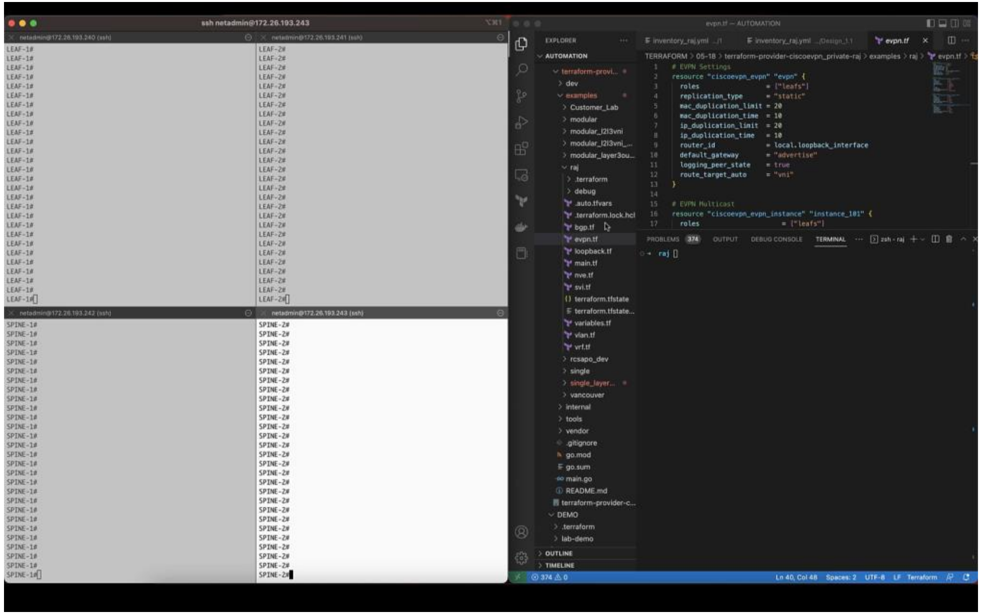

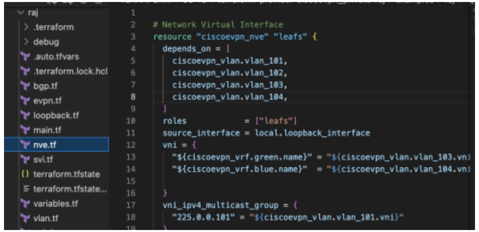

[Terraform Provider](#)
[Terraform Examples](#)

Catalyst 9000 EVPN Reference

Configuration Guide

Completed Chapters

BGP EVPN VXLAN Overview
Configuring EVPN VXLAN Layer 2 Overlay Network
Configuring EVPN VXLAN Layer 3 Overlay Network
Configuring EVPN VXLAN Integrated Routing and Bridging
Configuring Spine Switches in a BGP EVPN VXLAN Fabric
Configuring DHCP Relay in a BGP EVPN VXLAN Fabric
Configuring VXLAN-Aware Flexible NetFlow
Configuring Tenant Routed Multicast
Configuring EVPN VXLAN External Connectivity
Cisco DNA Service for Bonjour Overview
Configuring Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks
Troubleshooting BGP EVPN VXLAN
Feature History and Information for BGP EVPN VXLAN

More Coming Soon ...

Reference

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9600/software/release/17-5/configuration_guide/vxlan/b_175_bgp_evpn_vxlan_9600_cg.html

Would You Like to Know More?

Catalyst 9000 Series Enterprise Switches

- **cisco.com/go/cat9K**
- Cisco Catalyst 9000 at-a-Glance
- Cisco Catalyst 9000 Family FAQ
- Catalyst 9000 Series - Cisco Community
- Catalyst 9000 Series - CiscoLive Library

FREE

cs.co/cat9kbook

Cisco Catalyst TV @ YouTube

This channel is all about Cisco Catalyst Platforms and its services and software solutions. Subscribe and Explore Playlists [Catalyst Switching]() and [Catalyst Programmability & Automation]() for videos and demos by the Technical Marketing Engineers on latest, relevant and exciting topics.

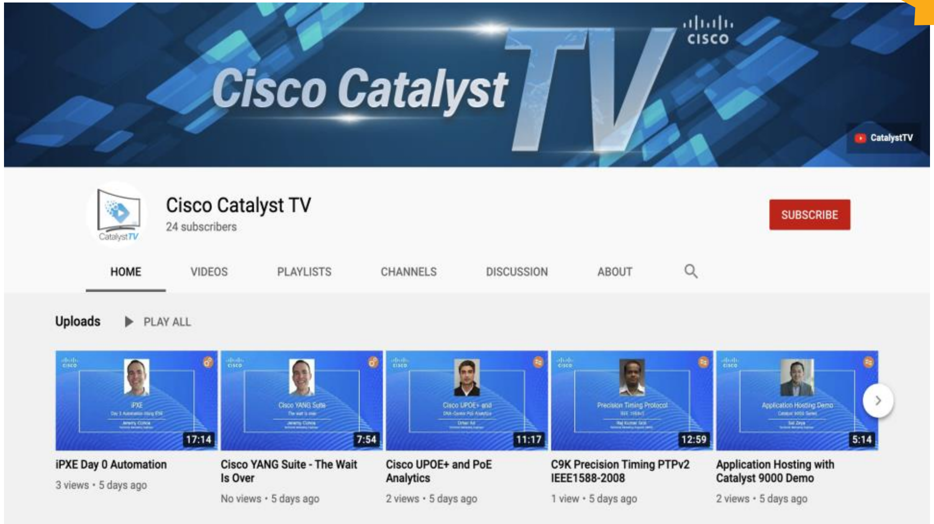

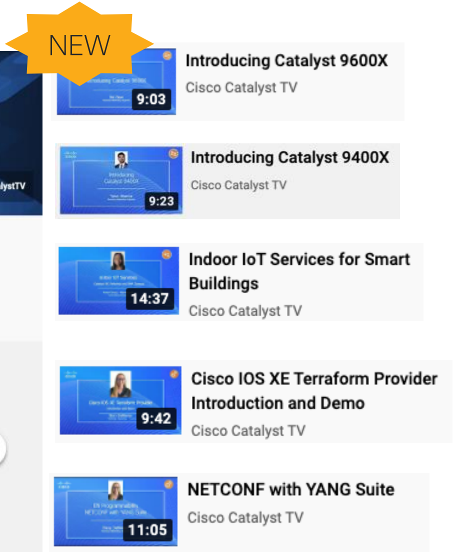

SUBSCRIBE

Cs.co/CatalystTV

Complete your Session Survey

- Please complete your session survey after each session. Your feedback is important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (open from Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Session Catalog and clicking the "Attendee Dashboard" at https://www.ciscolive.com/emea/learn/sessions/session-catalog.html

Continue Your Education

Visit the Cisco Showcase for related demos.

Book your one-on-one Meet the Engineer meeting.

Attend any of the related sessions at the DevNet, Capture the Flag, and Walk-in Labs zones.

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

The bridge to possible

Thank you

cisco Live!

ALL IN