The bridge to possible

BGP EVPN in Enterprise Campus

Building Scalable Fabrics with Catalyst 9000 Switches

Manish Sharma, Technical Marketing Engineer

BRKENS-2092

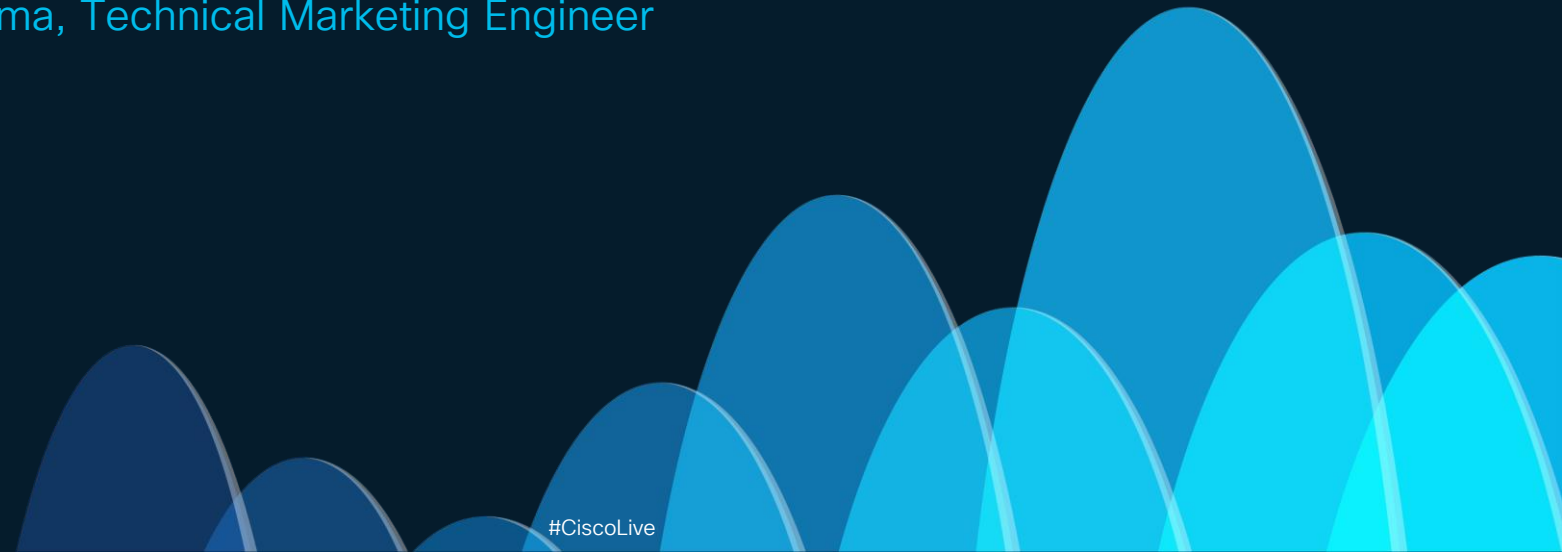

cisco Live!

#CiscoLive

Cisco Webex App

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKENS-2092

Questions?

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

How

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

Webex spaces will be moderated by the speaker until June 7, 2024.

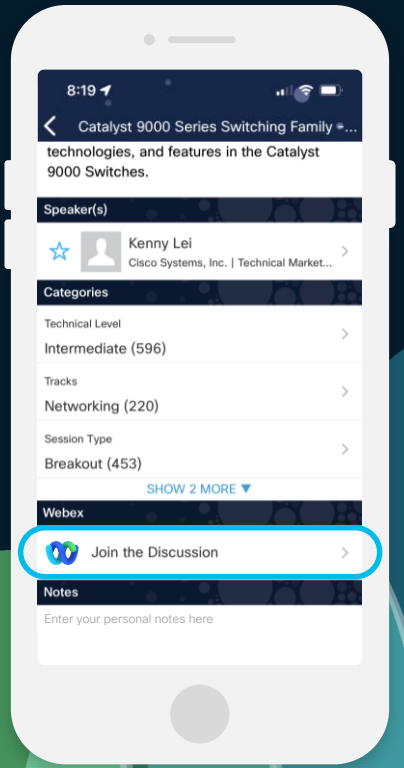

Agenda

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

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

2019 16.9.1

Catalyst 9000 Spine | Leaf | Border Support
Overlay Type : Layer 2 Cross-Connect
Overlay Type : Centralized Gateway (Asymmetric IRB)
IPv4 ARP / IPv6 ND Suppression for Distributed AnyCast Gateway
IPv4 ARP / IPv6 ND Suppression for Layer 2 Leaf
EVPN L2 Multi-homing with Cisco StackWise Virtual
Central Switching Wireless support
EVPN MAC/IP mobility improvement for Distributed Wireless Mode
Firewall integration for Service-Insertion
IPv6 host in overlay
DHCPv4/v6 Relay in Default VRF
VXLAN 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

2019 16.12.1

Catalyst 9000 Hybrid Role - Spine + Leaf + Border Support
BGP EVPN Layer 2 Multi-home ESI support
Distributed AnyCast RP for Tenant Routed Multicast (TRM)
PIM SM and SSM Underlay support for TRM
IPv4/IPv6 host overlay Multicast with Default MDT
EVPN Fabric External Domain RP
L2/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
NVO SNMP MIB for BGP EVPN
Border : Multi-VRF IPv4 and IPv6 Multicast Handoff
Border : EVPN to VPLS Access-VFI Multi-home support
Border : EVPN to VPLS IRB Interworking domain
Border : EVPN to MPLS Multicast VPNv4 Integration
Border : EVPN to External Multicast Domain Integration

2020 17.3.1

Catalyst 9k EVPN Journey

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 Multicast VPN Handoff
Border : EVPN to Global VN Extranet support
Border : EVPN to Non-EVPN VN Extranet support

2021 17.6.1

Per-VLAN Peer-to-Peer Protected Mode
Optimized Layer 2 Overlay Multicast
VXLANv6 : Layer 3 Overlay over VXLANv6
VXLANv6 Underlay – IPv6 BGP EVPN AF Peering support
VXLANv6 Underlay – RFC 5549 / RFC 8950 Support
VXLANv6 Overlay – Layer 2 Overlay over VXLANv6
VXLANv6 Overlay – EVPN RT-3 (Ingress-Replication) over VXLANv6
VXLANv6 Overlay – Border Layer 2 802.1Q IPv6 Handoff
VXLANv6 Overlay – Distributed AnyCast GW Overlay over VXLANv6
VXLANv6 Overlay – EVPN RT-2 & RT-5 (IPv4 | IPv6 VRF) over VXLANv6
VXLANv6 Overlay – Border Layer 3 802.1Q IPv6 Handoff
VXLANv6 Overlay – Border Layer 3 MPLS VPNv6 Handoff
VXLANv6 Overlay – VXLAN – v4 + v6 Neighbor Support
VXLANv6 Overlay – VXLAN – v4 + v6 Support for DAG
VXLANv6 Overlay – EVPN RT- 2 | 3 | 5 over VXLAN v4 + v6
VXLANv6 Overlay – VXLANv4 to VXLANv6 Migration w/ IR
VXLANv6 Overlay – VXLAN – v4 + v6 Manual Switching
VXLANv6 Overlay – VXLANv4 to VXLANv6 Migration w/ Multicast Rep
VXLANv6 Overlay – Border Layer

2022 17.9.1

Advanced RT-2 and RT-5 BGP EVPN Route-Map support
Per-VLAN BGP EVPN ESI Layer 2 Multi-home support
EVPN Micro-Segmentation with Cisco TrustSec integration
CLI Simplicity – Dynamic BGP EVPN Address-Family Peering Support
CLI Simplicity – Dynamic Multicast VPN 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
C9500X | C9600X – Optimized Layer 2 Overlay Multicast
VXLANv6 Overlay – TRM Multicast – v4 Support for DAG
VXLANv6 Overlay – AnyCast RP over VXLANv6 support
VXLANv6 Underlay – IPv6 BGP mVPN AF Peering support
VXLANv6 Overlay – TRM Multicast – v6 Support for DAG
C9500X | C9600X – Per-VLAN BGP EVPN ESI Layer 2 Multi-home support
Scalable Fabric : Multi-Tenant Stacked VLAN (QinQ) Layer 2 Overlay Networks

2023 17.12.1/13.1

Cisco Catalyst 9000

BGP EVPN VXLAN Fabric

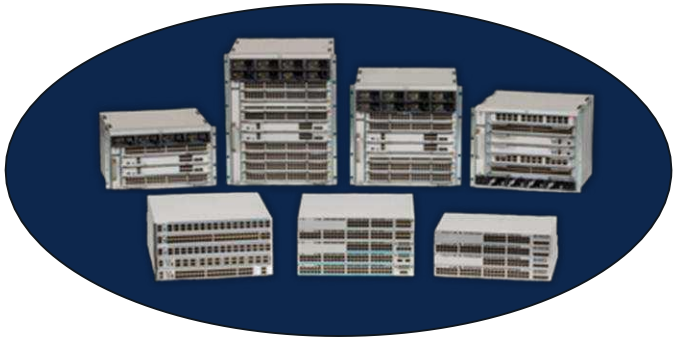

Enterprise | Healthcare | Education | Financial | Public Sector | Manufacturing | Hospitality | Media | Transportation | Retail

VXLAN Overview

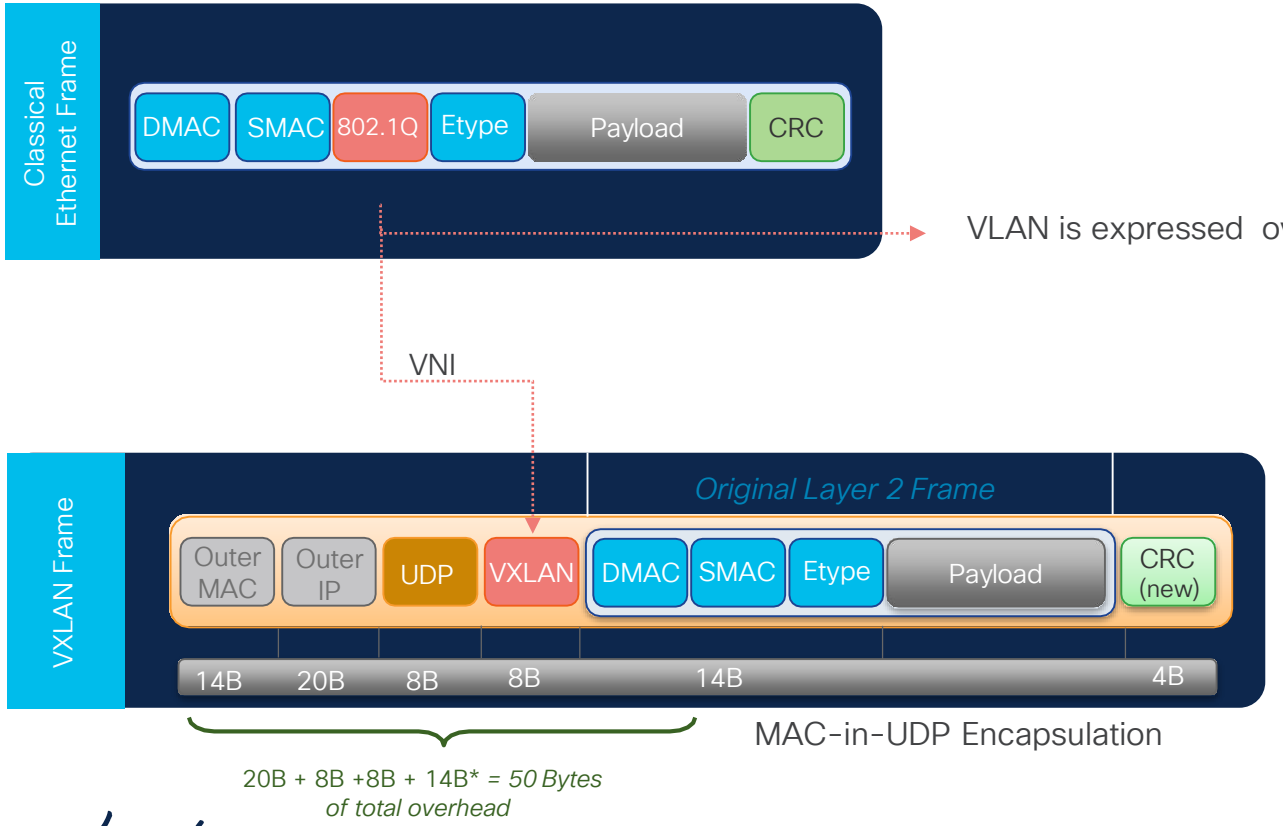

VXLAN leverages the VNI field with a total address space of 24 bits

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

VXLAN with BGP EVPN

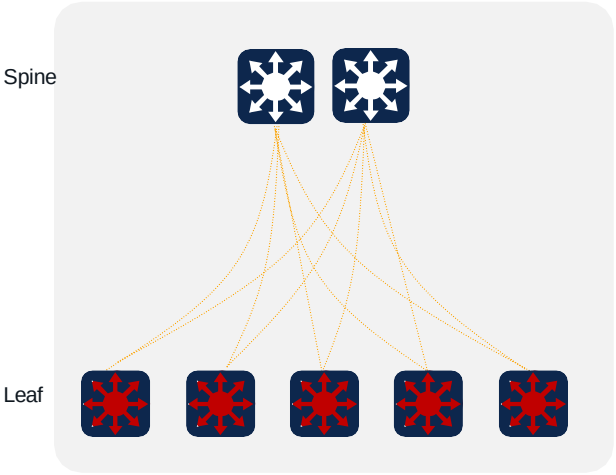

- Standards based Overlay (VXLAN) with Standards based Control-Plane (BGP)
- Layer-2 MAC and Layer-3 IP information distribution by Control-Plane (BGP)
- Forwarding decision based on Control-Plane (minimizes flooding)
- Integrated Routing/Bridging (IRB) for Optimized Forwarding in the Overlay
- Multi-Tenancy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

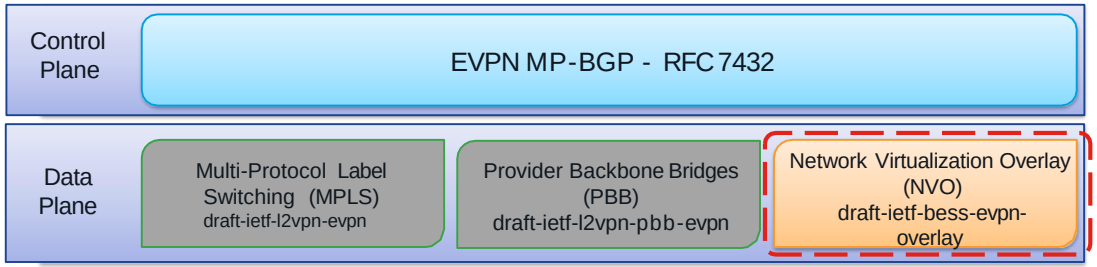

EVPN over NVO Tunnels (VXLAN)

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

BGP EVPN System Role

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.

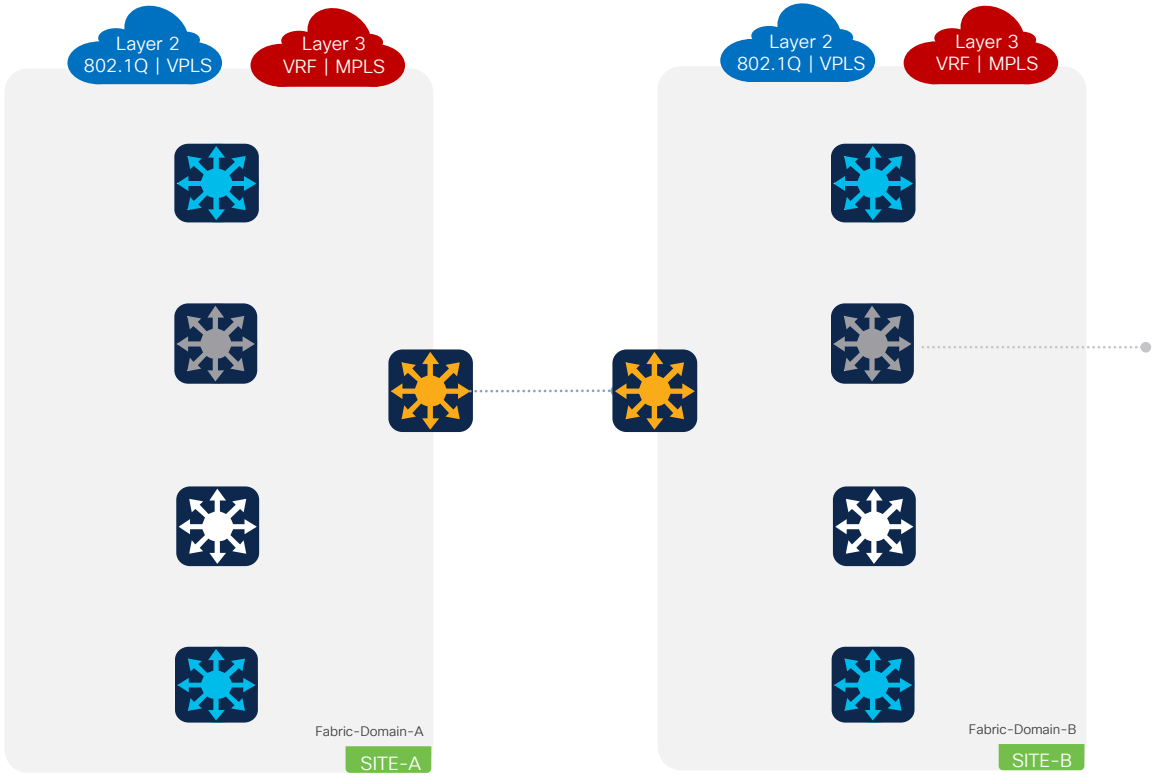

BGP EVPN System Role

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.

★ - Recommended

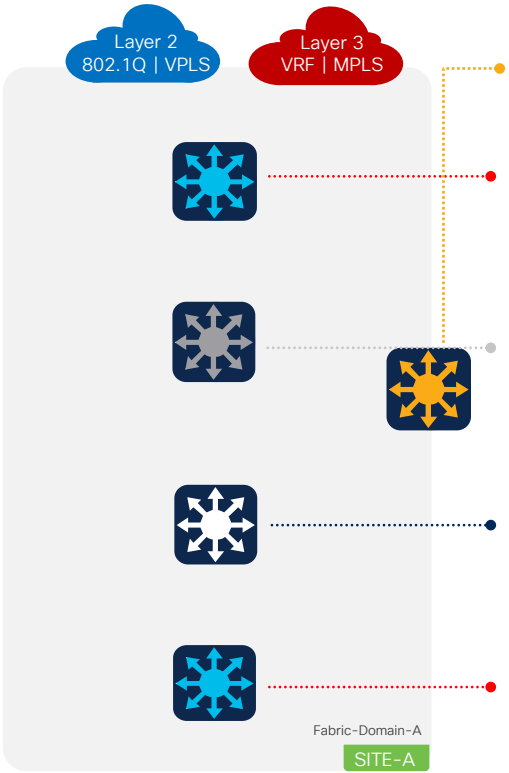

Catalyst EVPN Scale and Performance Matrix

Cisco Catalyst BGP EVPN Configuration Guide
Scale and Performance Chapter

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

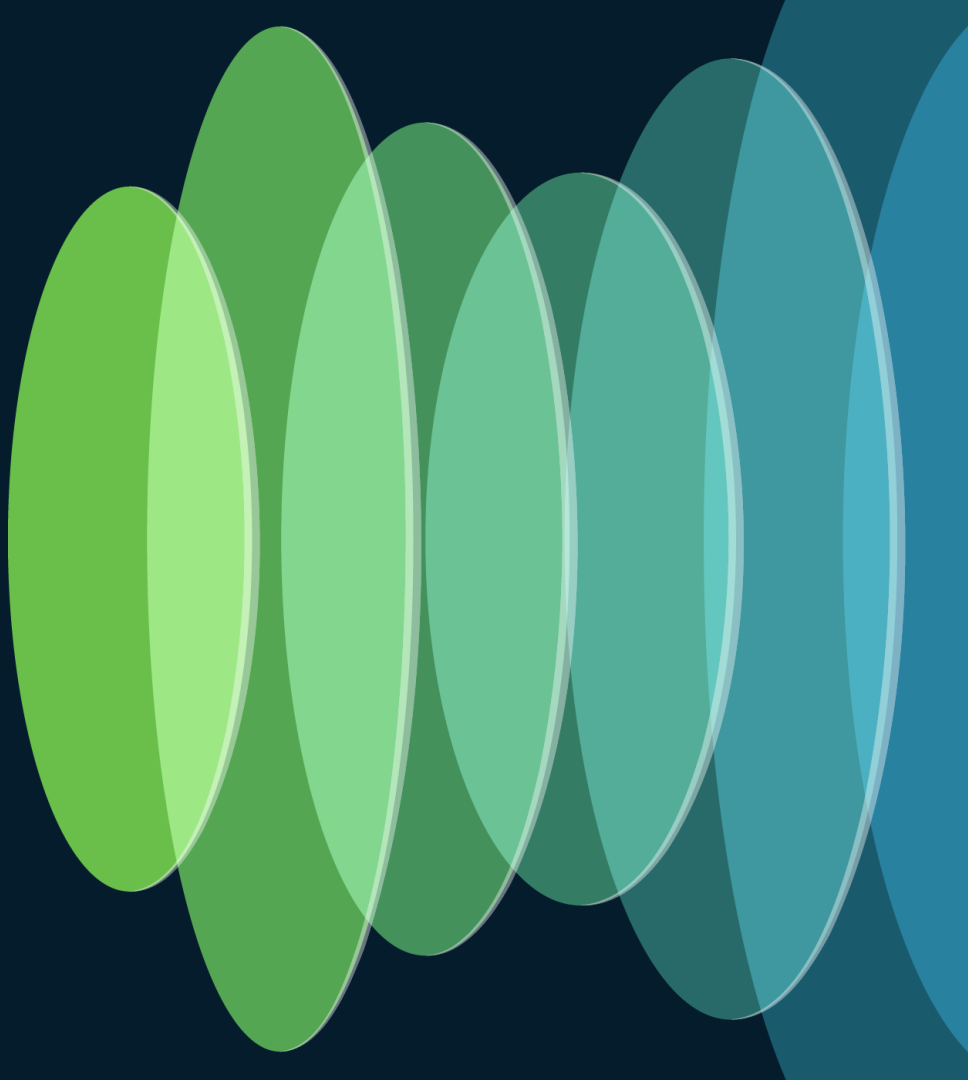

cisco Live!

Traditional Network Transition

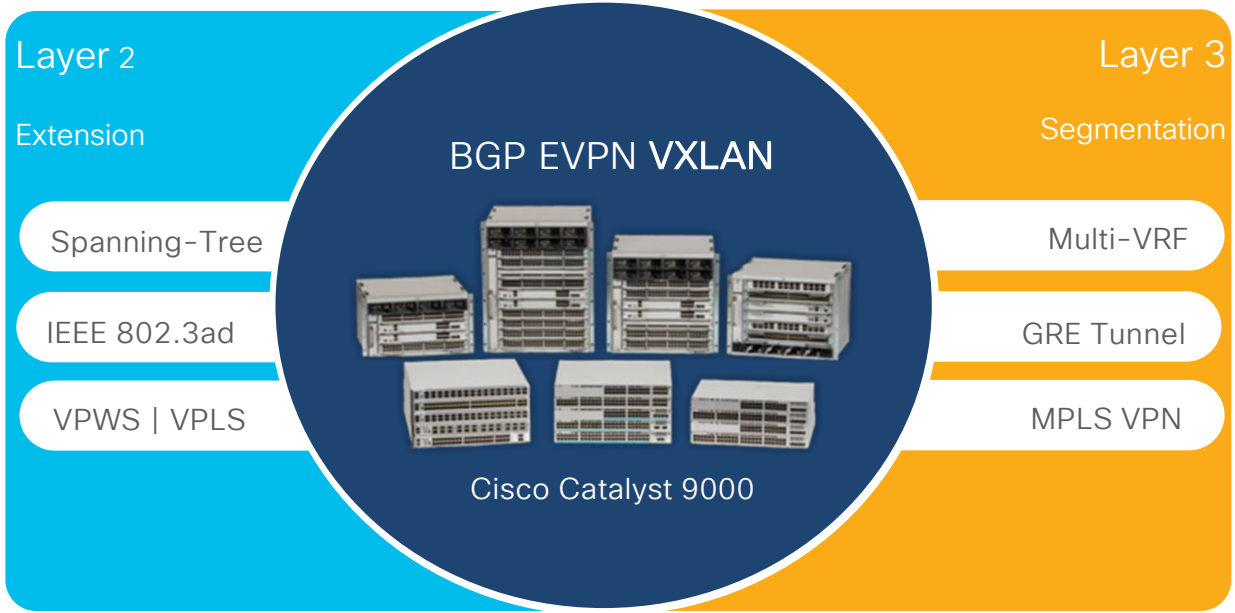

EVPN Evolution

- Product transition drives architecture transitions
- Convergence of traditional L2 overlay to simplified and scalable fabric
- Transition classic L3 overlays to enterprise-grade scalable fabric
- Unified end-to-end common fabric architecture reducing cost and complexity

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

Enterprise BGP EVPN Reference Architecture

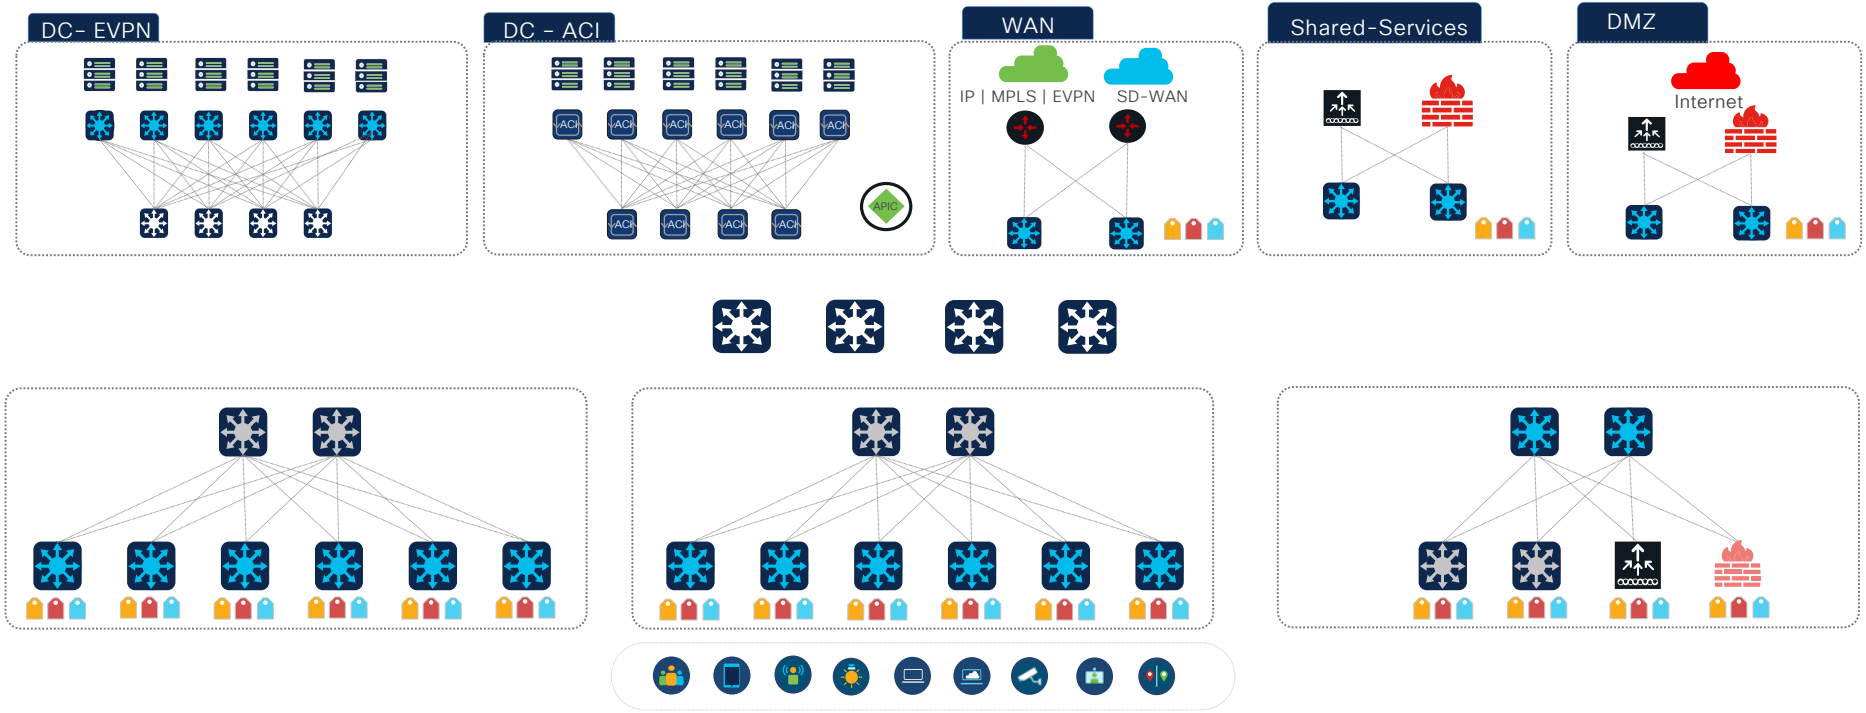

Industry Standard

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

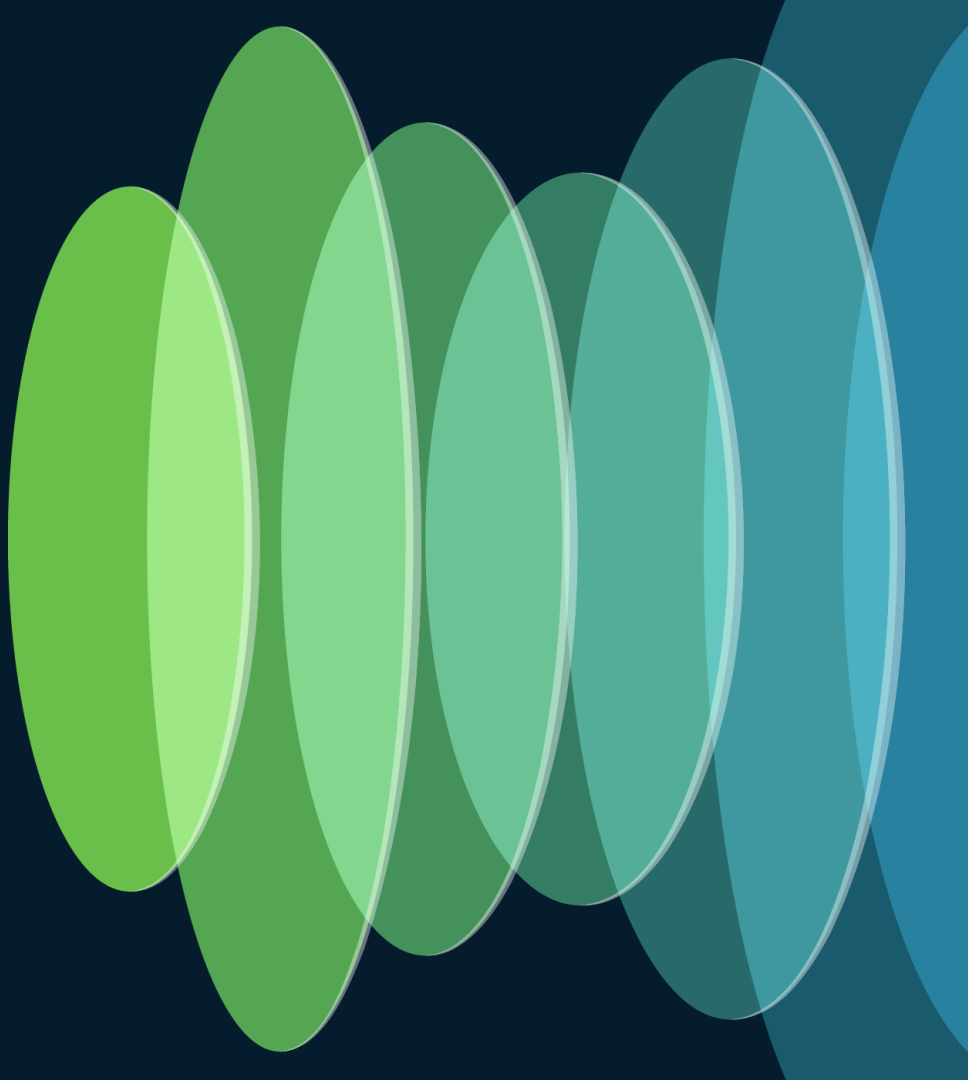

VXLAN Constructs

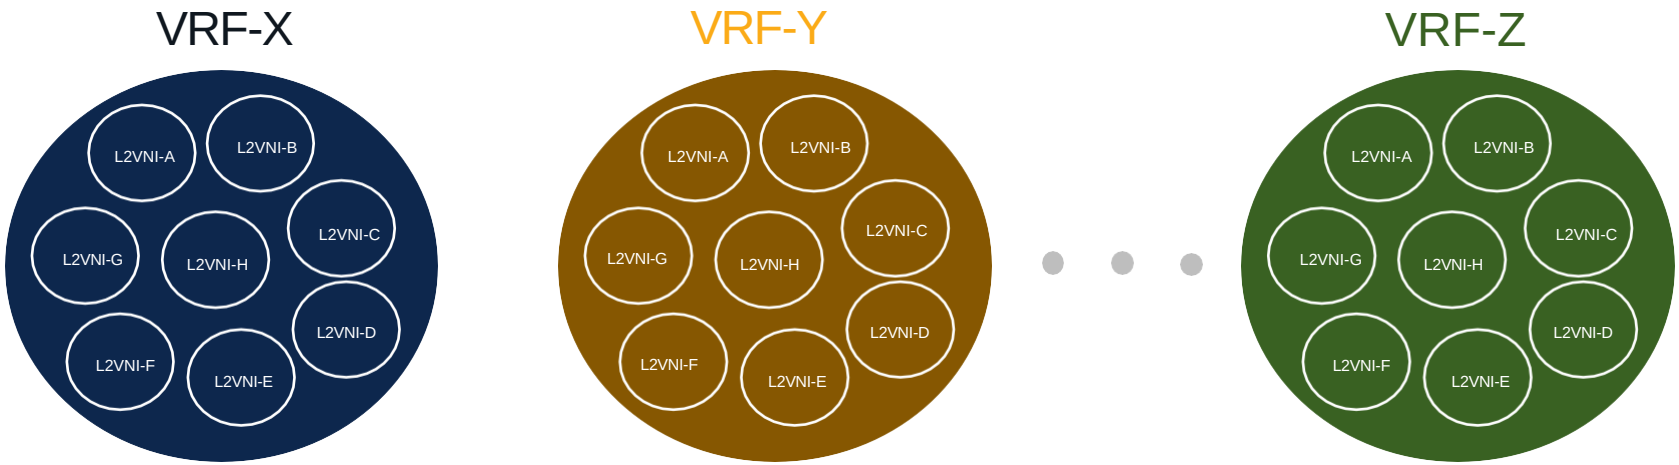

VNI – Virtual Network Identifier

MP-BGP EVPN Route Type(s)

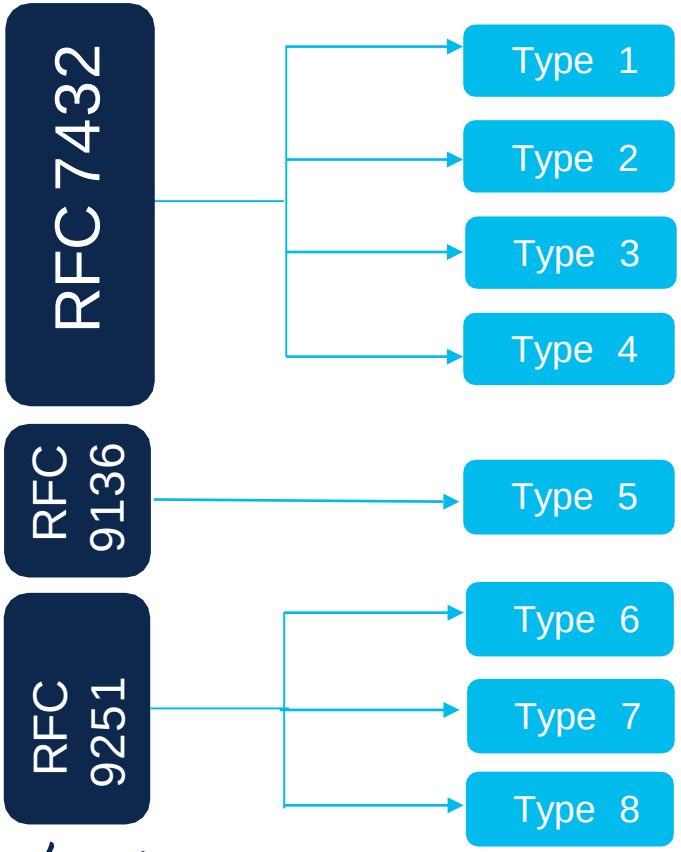

Ethernet Auto-Discovery (A-D) route

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)

Bridging and Routing

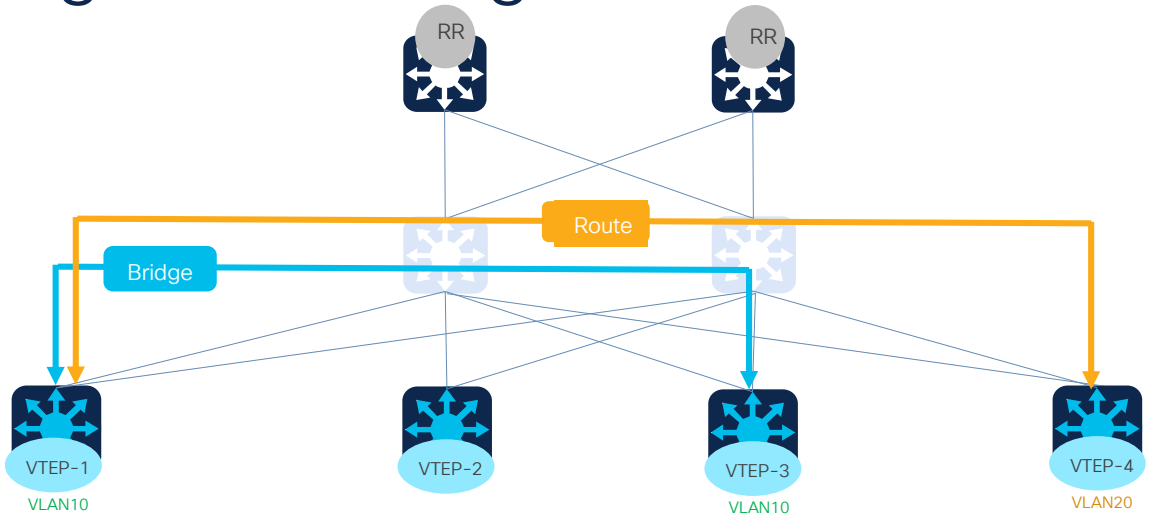

Host A
MAC: 0000.1111.1111
IP: 192.168.10.10

Host B
MAC: 0000.2222.2222
IP: 192.168.10.11

Host C
MAC: 0000.4444.4444
IP: 192.168.20.11

L3 VXLAN Tunnel
L2 VXLAN Tunnel

Overlay 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
 exit-address-family
```

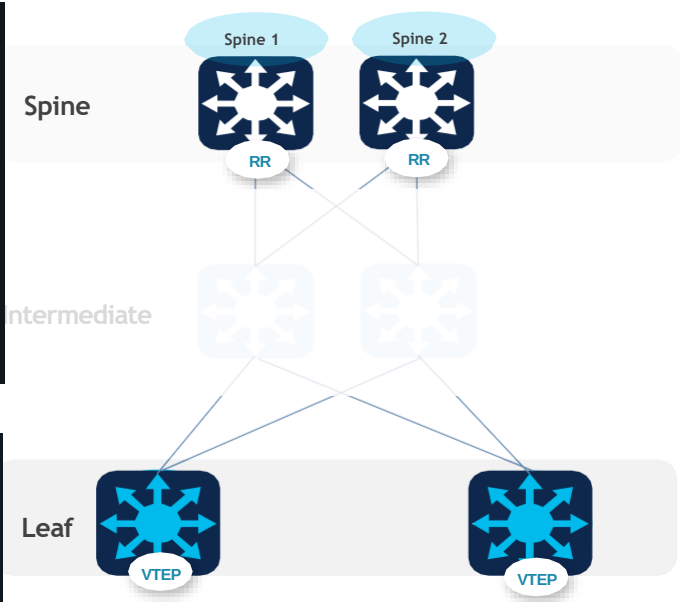

Spine-2

```
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
 exit-address-family
```

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
```

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

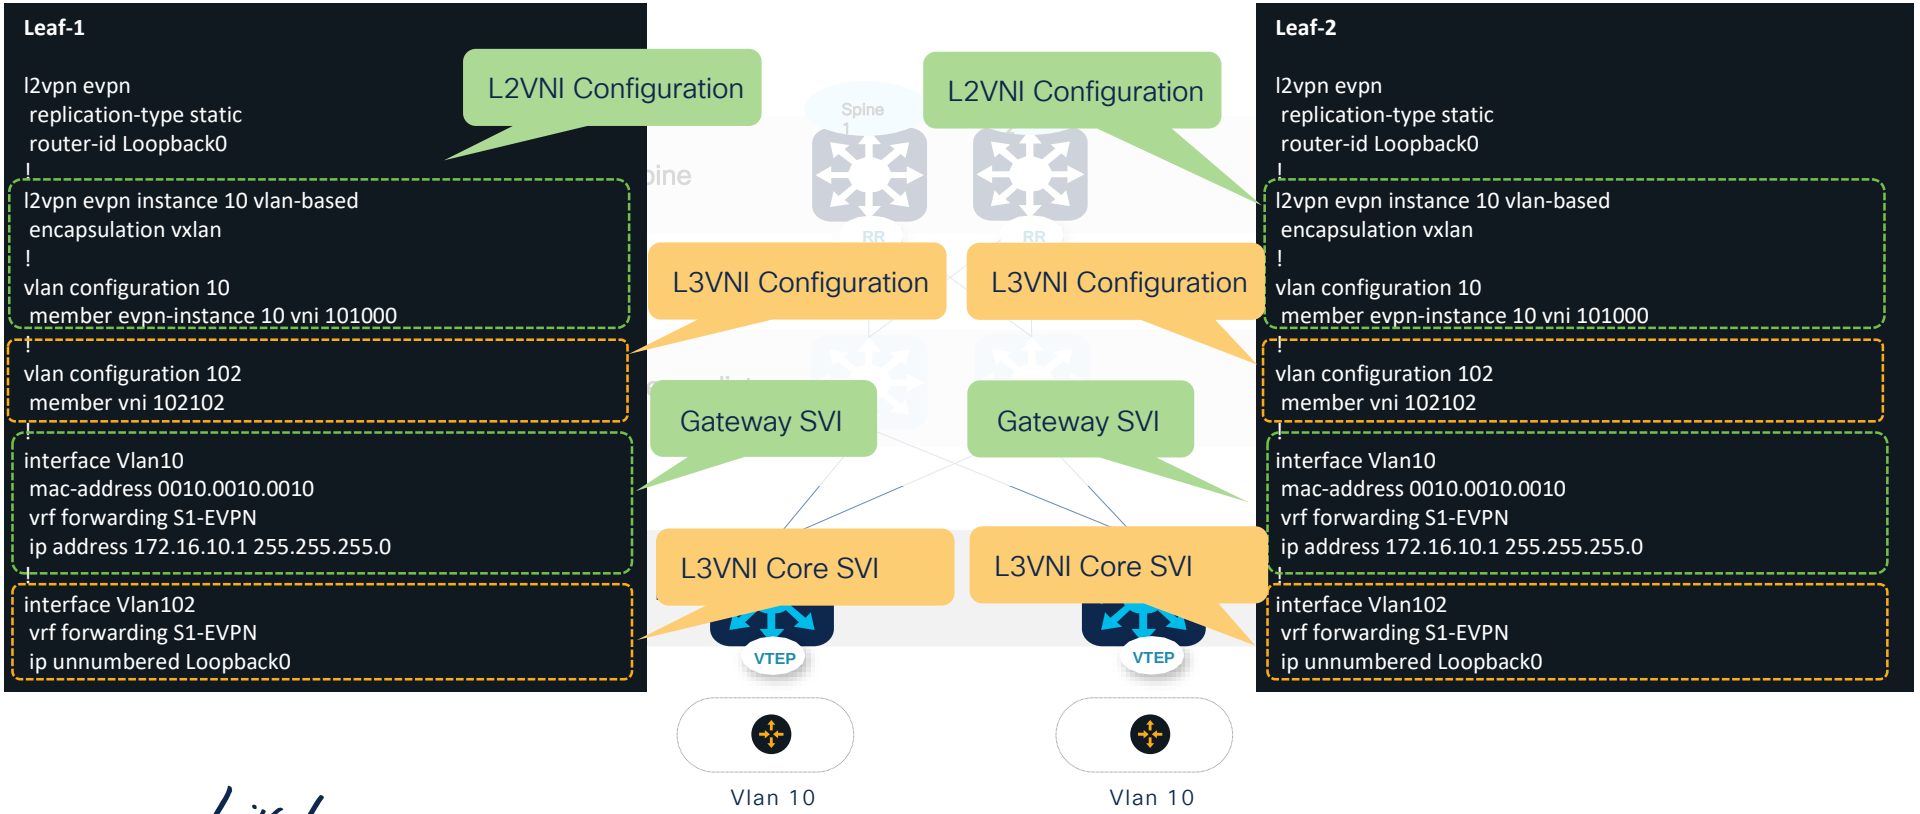

NVE Configuration

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
```

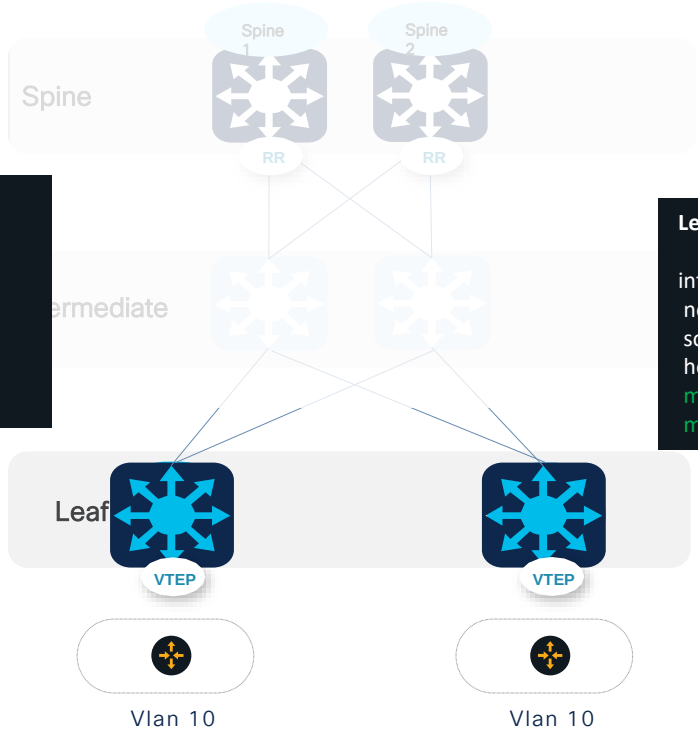

Leaf-1

```
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

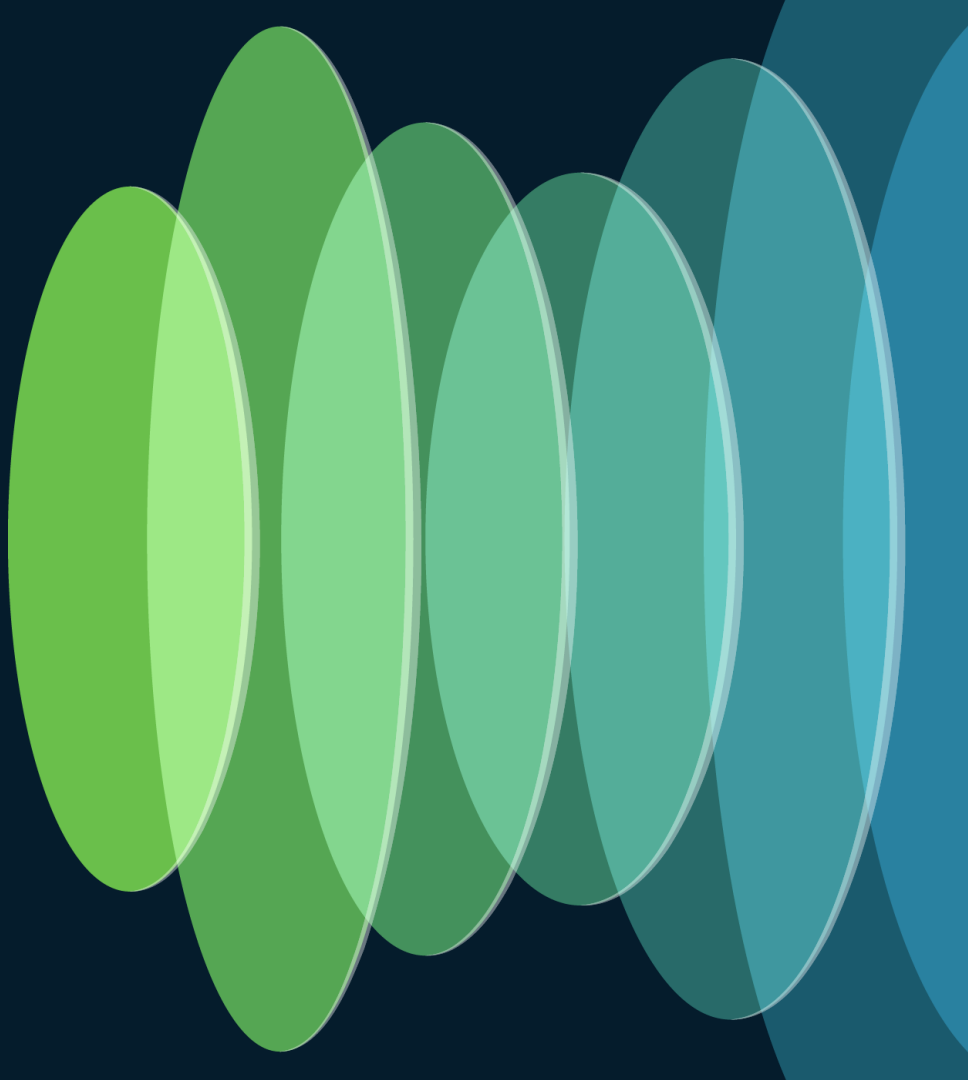

Underlay Design Considerations

Routed Access

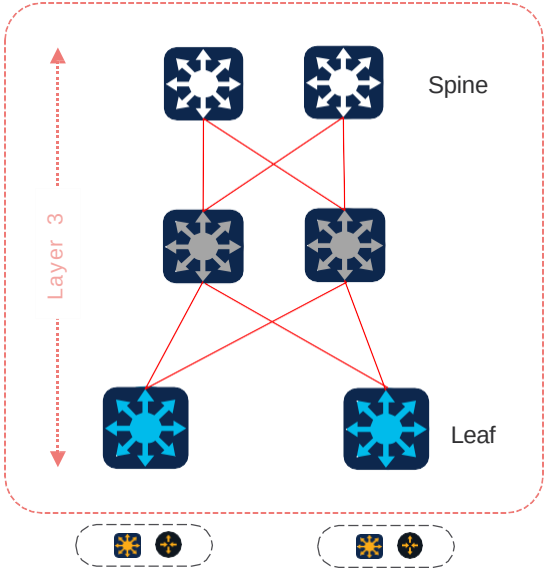

- Routed Access environment
- Leaf Layer – Routed Access
- Spine/RR – Direct | Multi-hop
- ECMP | Non-Blocking ports
- Underlay | Overlay IP Gateway
- L2 | L3 Overlay Support
- Multicast Support

StackWise-Virtual

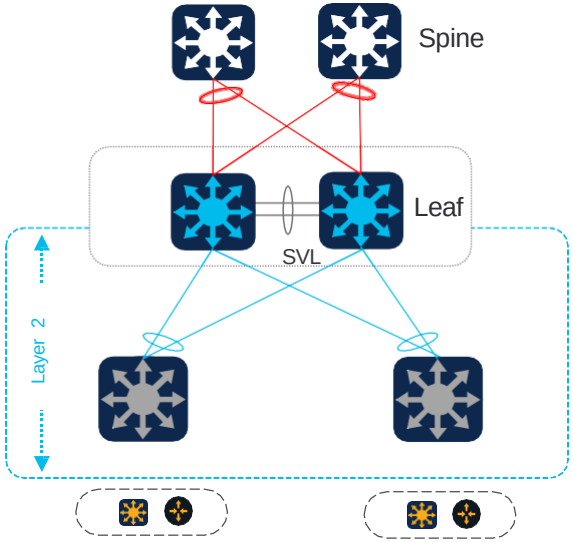

- Access – Traditional Layer 2
- Leaf Layer – Distribution
- Spine/RR – Direct | Multi-hop
- MEC | ECMP | Active/Active forwarding
- Underlay | Overlay IP Gateway
- L2 | L3 Overlay Support
- Multicast Support

Layer 2 | Layer 3 Distribution

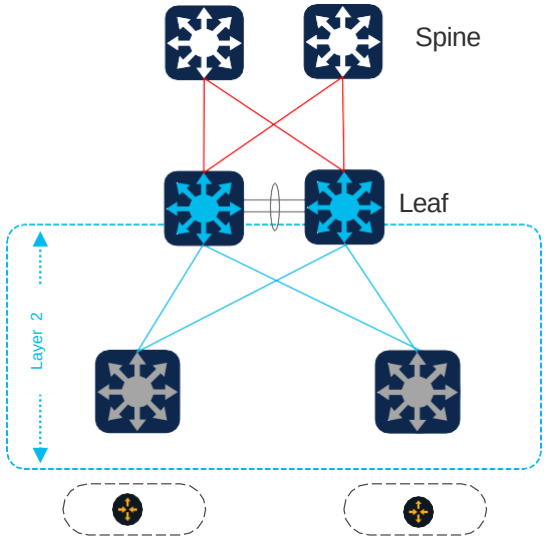

- Access – Traditional Layer 2
- Leaf Layer – Distribution
- Spine/RR – Direct | Multi-hop
- FHRP | ECMP | Multicast
- Underlay | Overlay IP Gateway
- L3 Overlay Support. No L2 Extension
- Multicast Support

Layer 2 Layer 3

EVPN ESI Dual-Home

(Ethernet Segment Identifier)

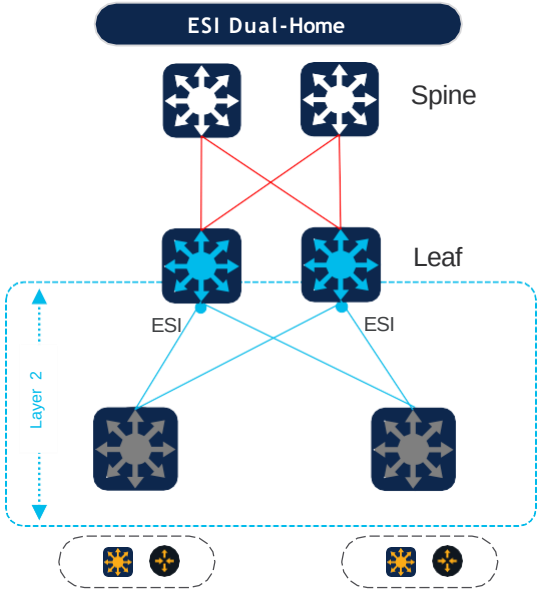

- Access – Traditional Layer 2
- Leaf Layer – Distribution
- Spine/RR – Direct | Multi-hop
- Per-ESI Anycast Gateway
- Per-VLAN | FHRP | ECMP | Multicast
- Active / Standby load-balancing
- L2 | L3 Overlay support
- Multicast Support

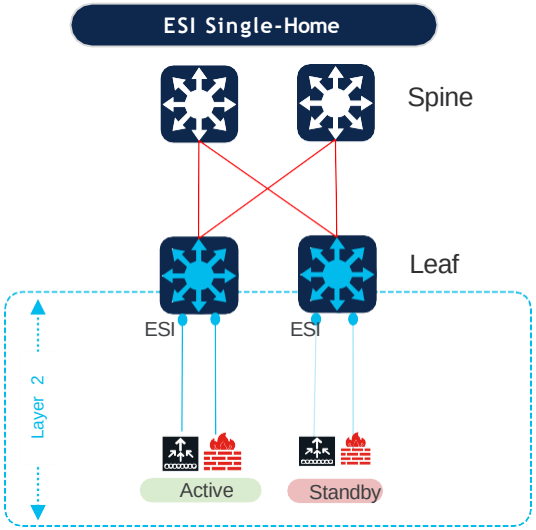

- Access – Traditional Layer 2
- Leaf Layer – Distribution
- Spine/RR – Direct | Multi-hop
- Per-ESI Anycast Gateway
- Per-VLAN | FHRP | ECMP | Multicast
- Active / Standby load-balancing
- L2 | L3 Overlay support
- Multicast Support

Layer 2 — Layer 3

Underlay Unicast Routing Design Alternatives

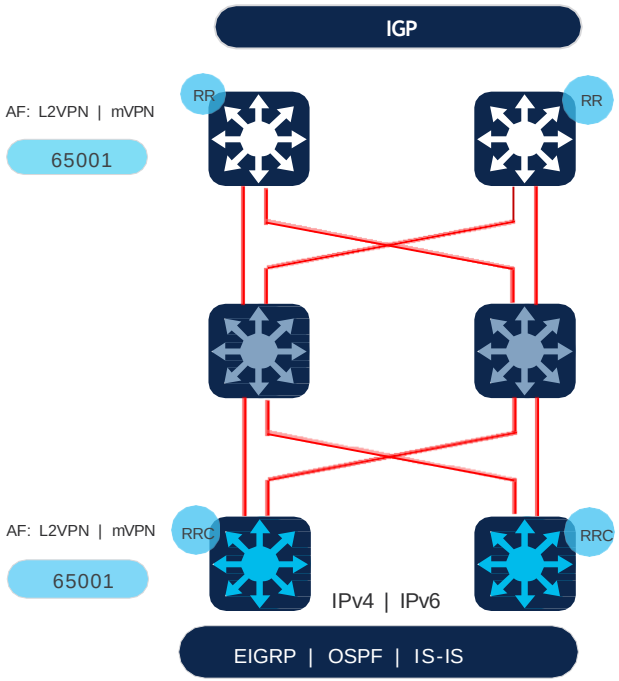

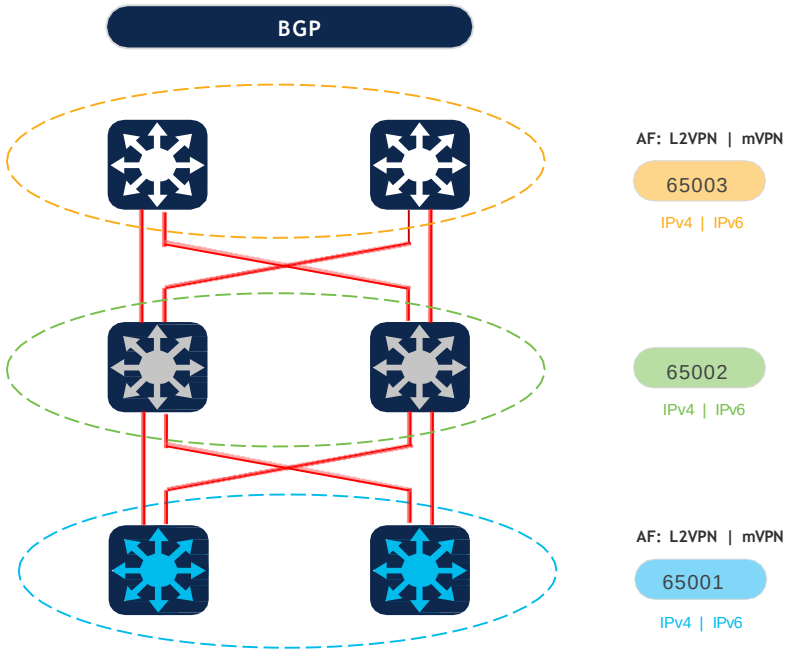

- Flexible Underlay Unicast alternatives – IGP (EIGRP/OSPF/IS-IS) or BGP
- Physical/Virtual Spine RR support – IOS-XE | NXOS | XR
- Secure link-layer underlay network encryption using MACSEC
- Underlay MTU size consideration. TCP MSS adjust supported.

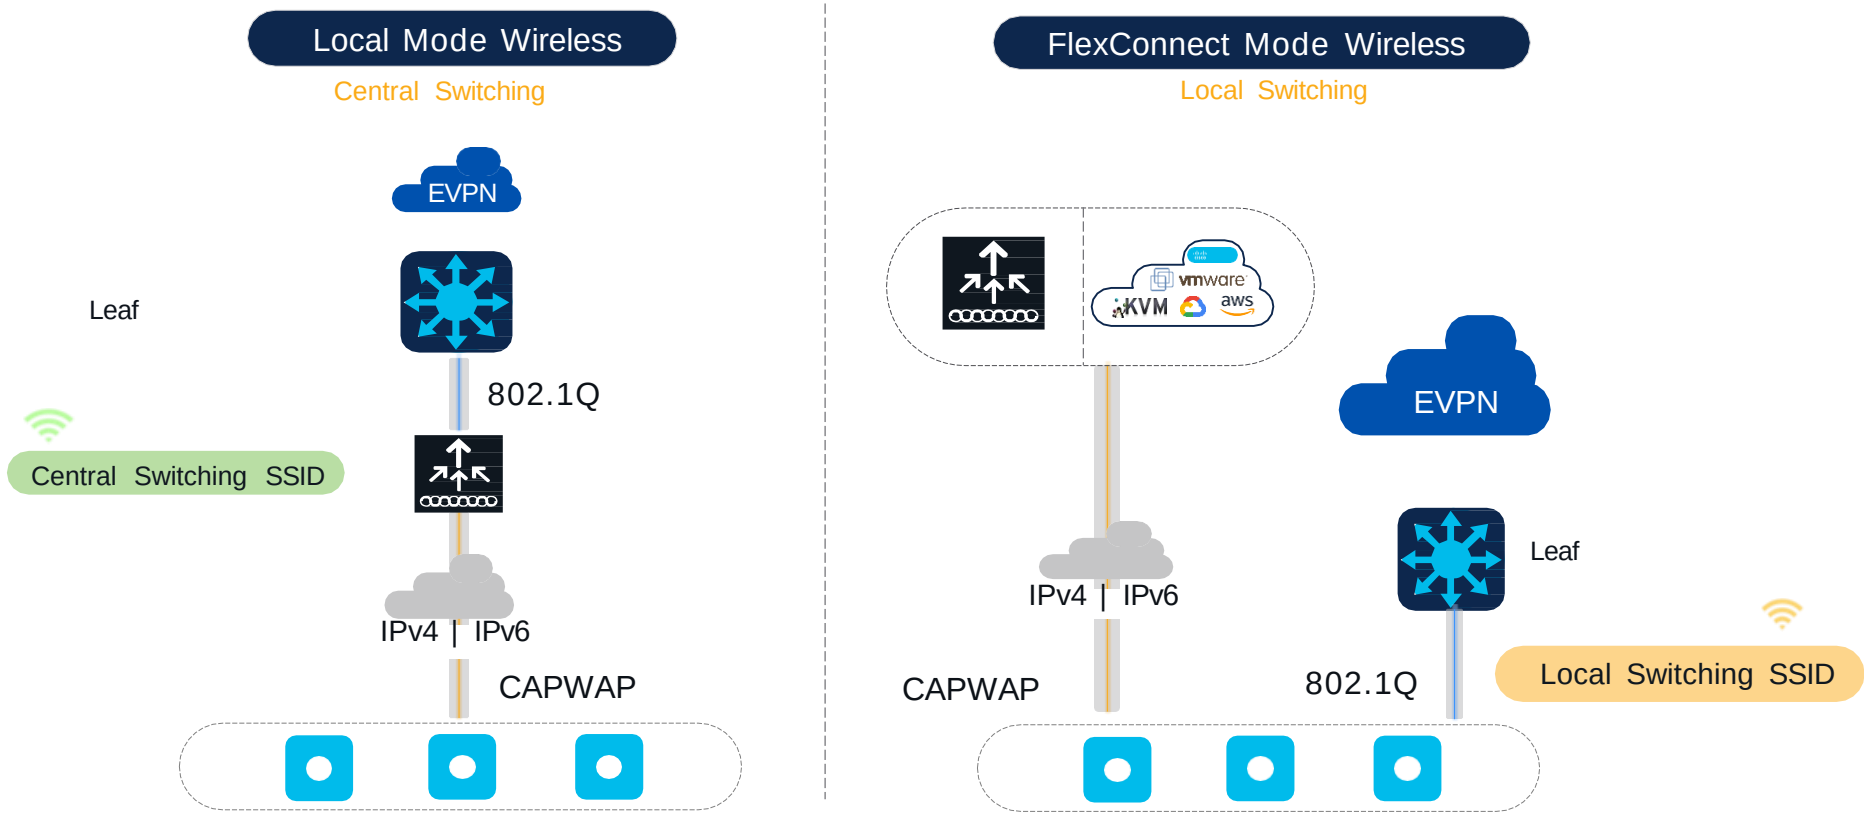

Wireless

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

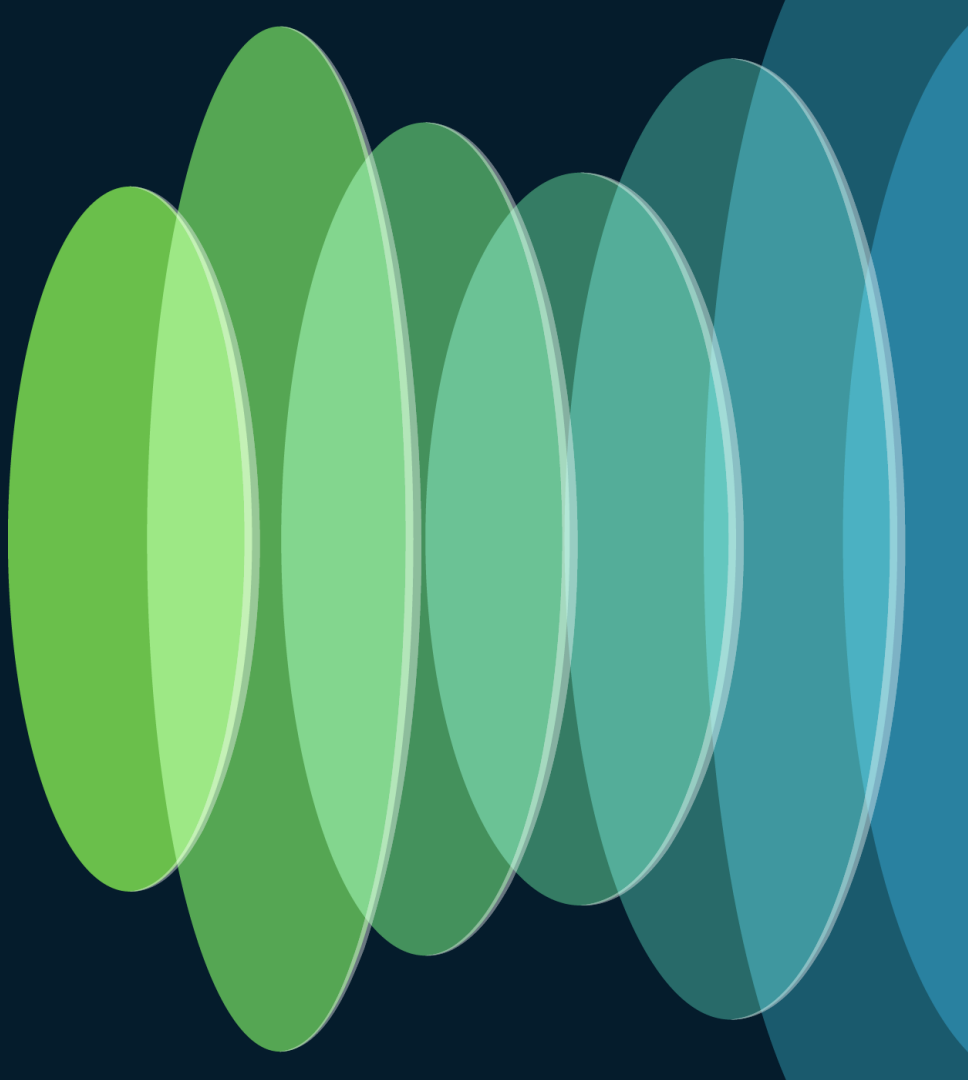

Flexible Routing and Bridging Overlay Types

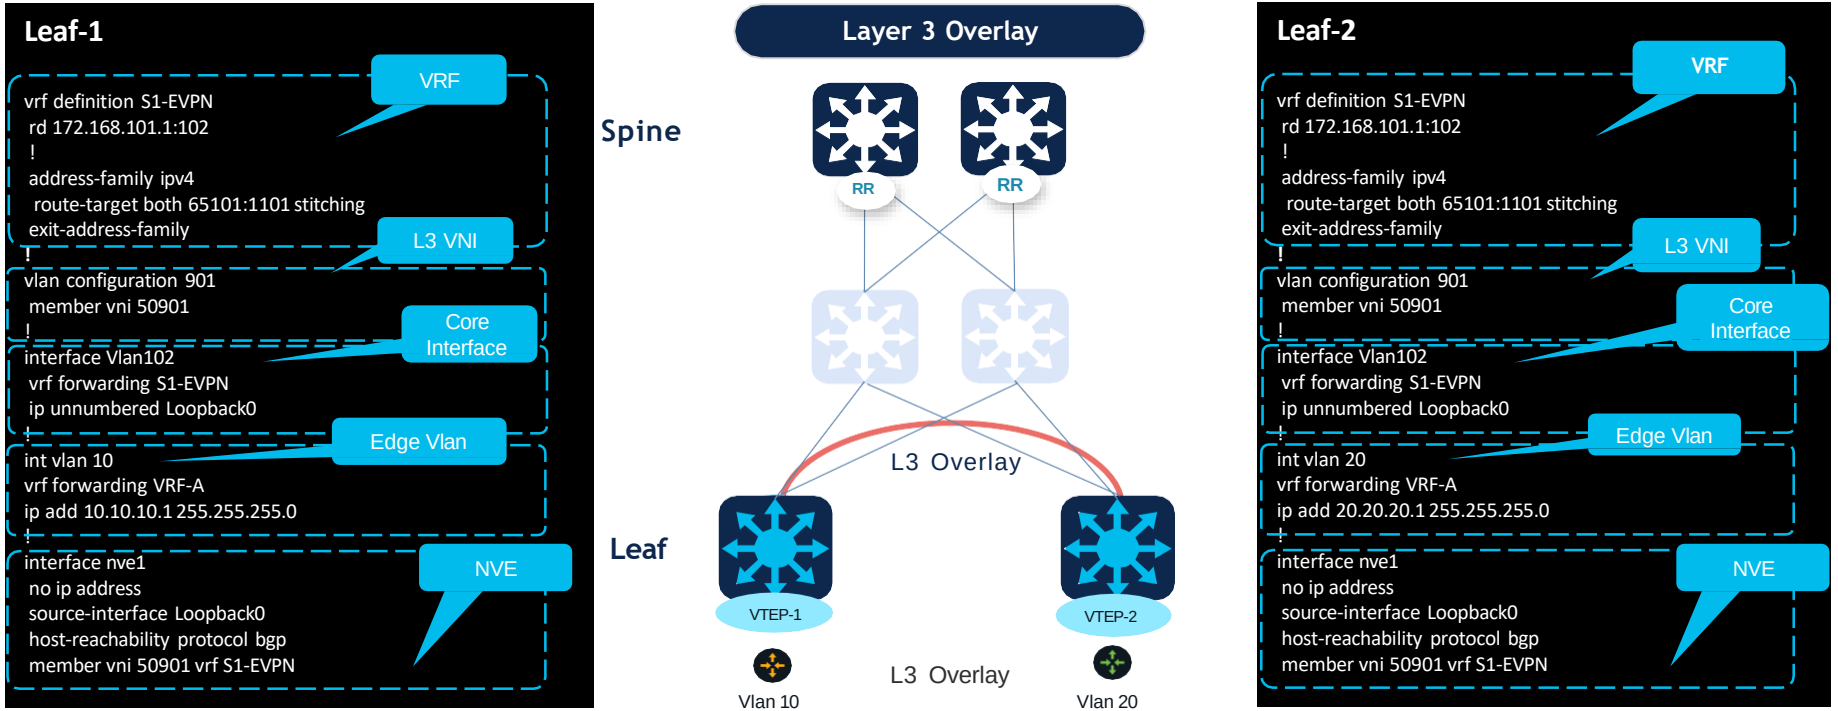

- Layer 3 overlay network allows host devices in different Layer 2 networks to send Layer 3 or routed traffic to each other
- The network forwards the routed traffic using a Layer 3 virtual network instance (VNI) and an IP VRF.

Flexible Routing and Bridging Overlay Types

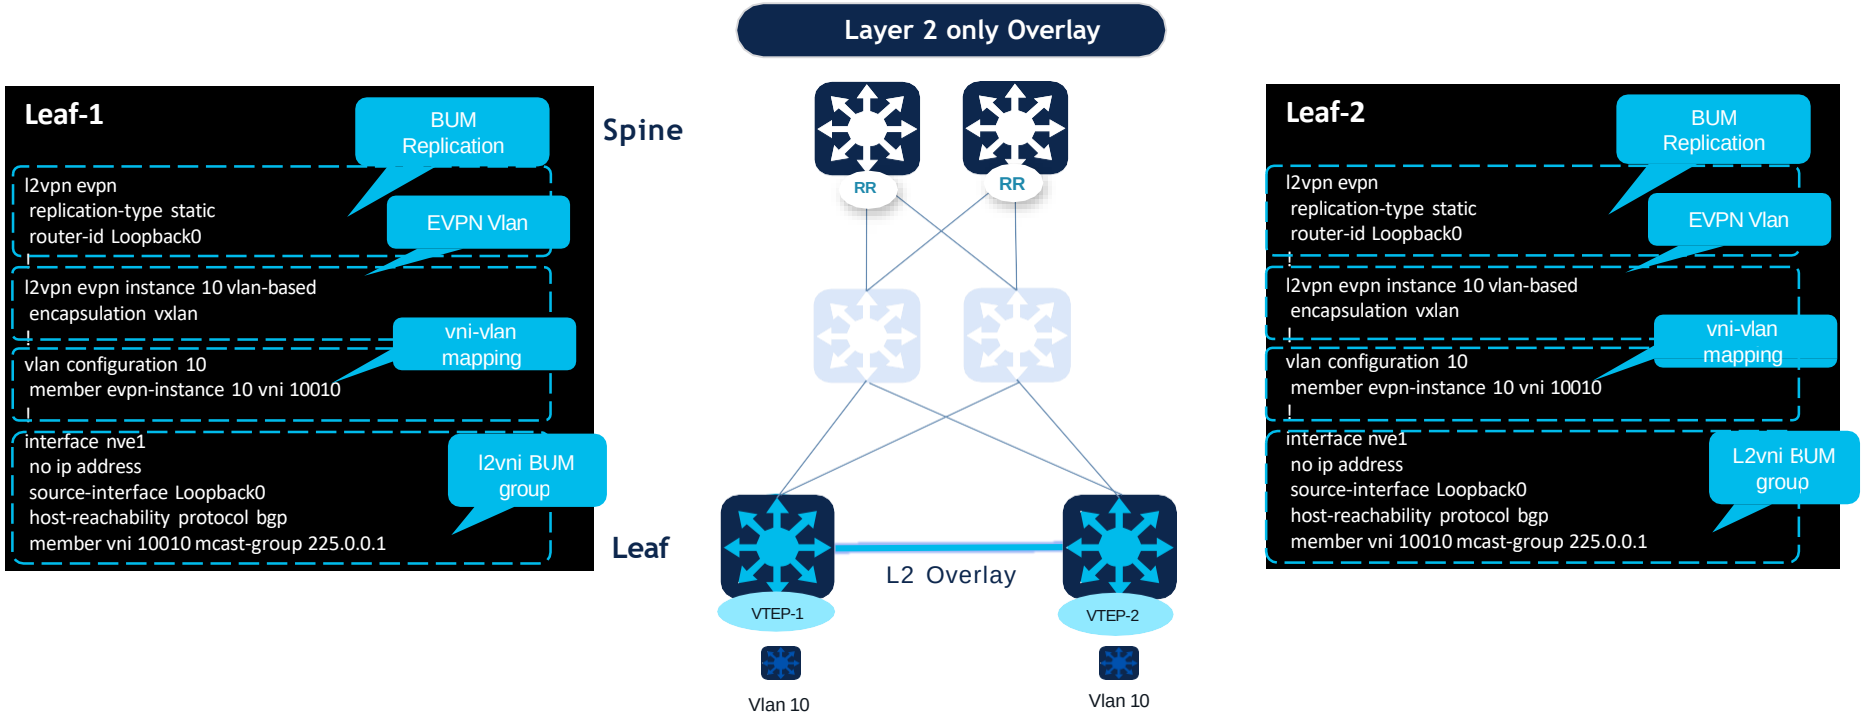

- L2 only stretch across the EVPN domain
- Flexible workload placement, host mobility, and optimal traffic forwarding across the BGP EVPN VXLAN fabric.

Flexible Routing and Bridging Overlay Types

Leaf-1

```
l2vpn evpn
 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 Vlan102
 vrf forwarding S1-EVPN
 ip unnumbered Loopback0
!
interface nve1
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
 member vni 102102 vrf S1-EVPN
 member vni 101000 ingress-replication
!
interface Vlan10
 mac-address 0001.0001.0001
 vrf forwarding S1-EVPN
 ip address 10.10.10.1 255.255.255.0
```

L2VNI

L3VNI

NVE

DAG

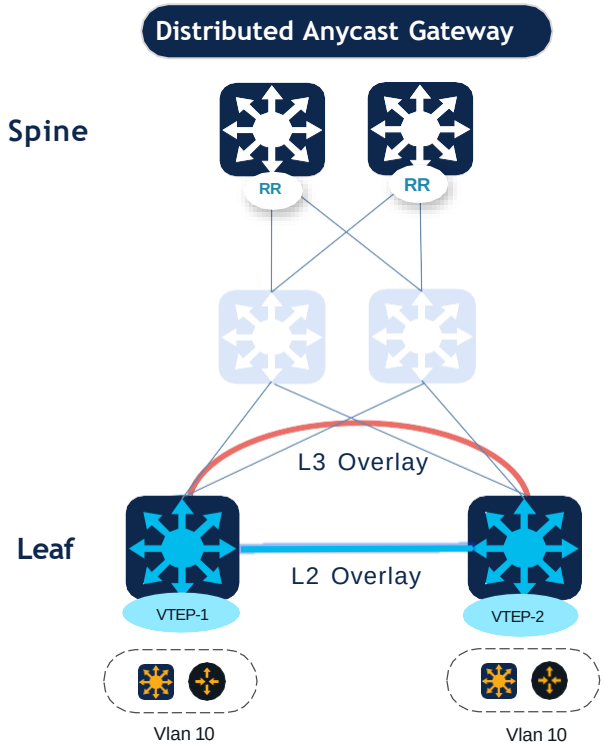

Leaf-2

```
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 Vlan102
 vrf forwarding S1-EVPN
 ip unnumbered Loopback0
!
interface nve1
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
 member vni 102102 vrf S1-EVPN
 member vni 101000 ingress-replication
!
interface Vlan10
 mac-address 0001.0001.0001
 vrf forwarding S1-EVPN
 ip address 10.10.10.1 255.255.255.0
```

L2VNI

L3VNI

NVE

DAG

- 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

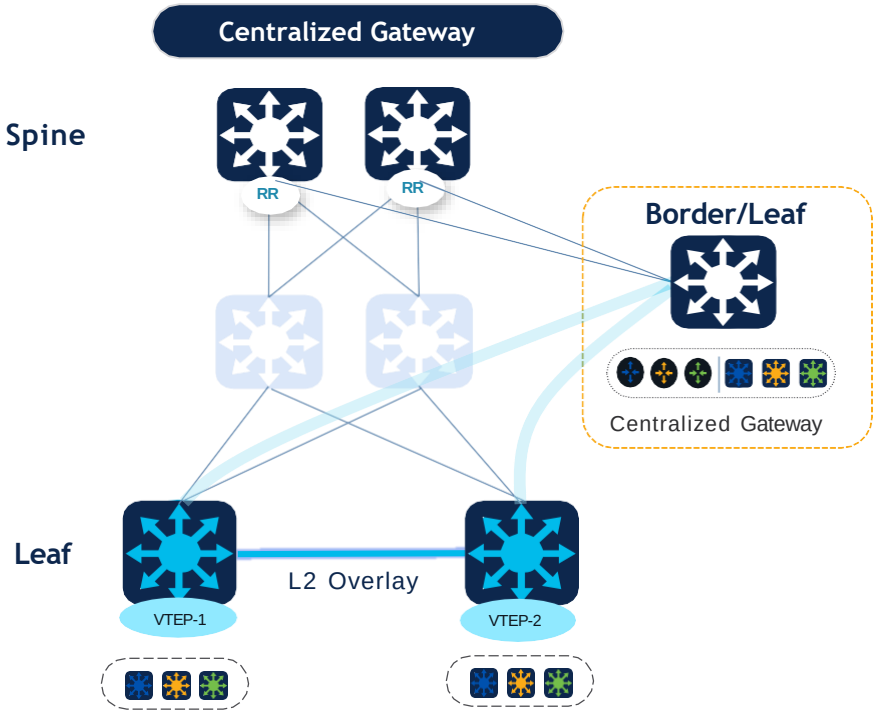

Flexible Routing and Bridging Overlay Types

Border/Leaf (CGW)

```
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
 !
```

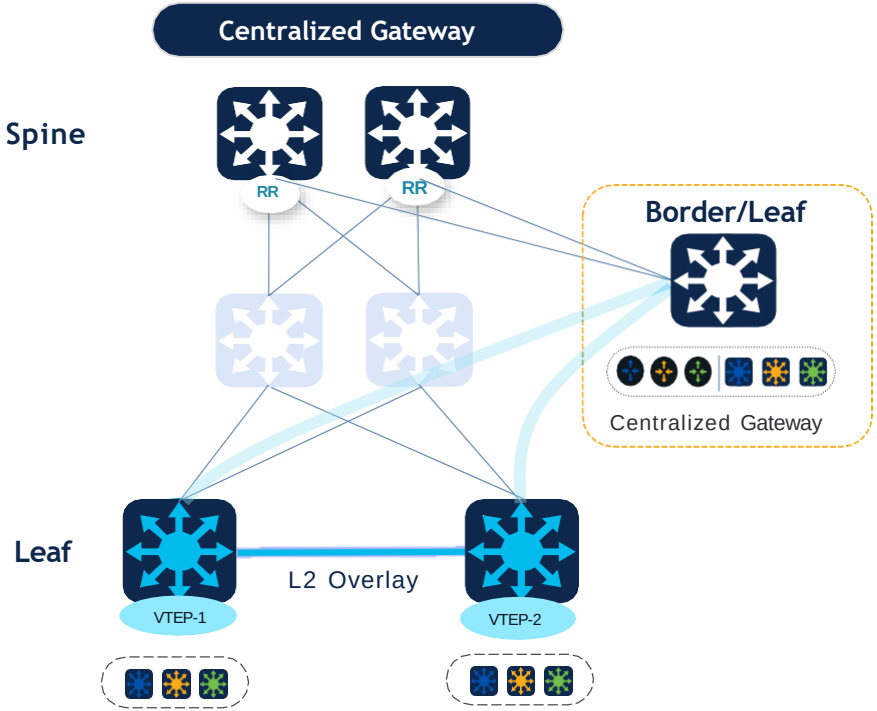

Flexible Routing and Bridging Overlay Topologies

Point to Point

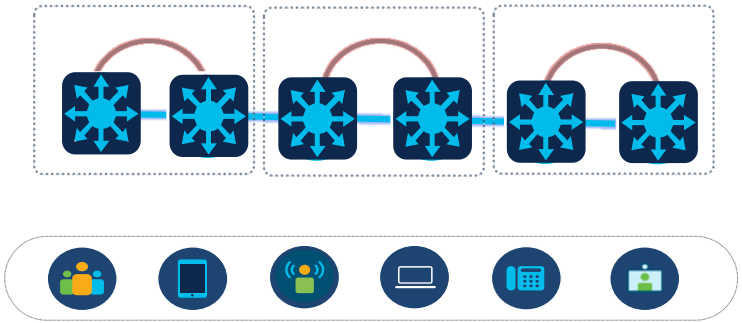

Hub & Spoke

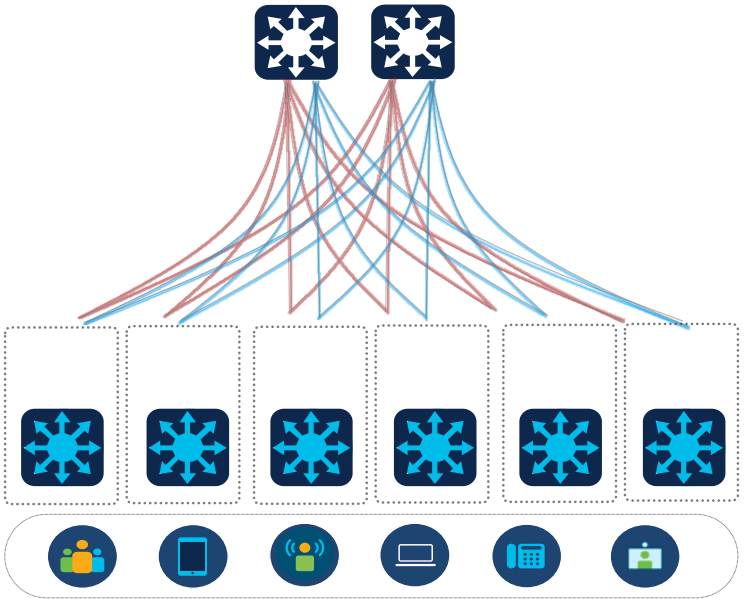

Flexible Routing and Bridging Overlay Topologies

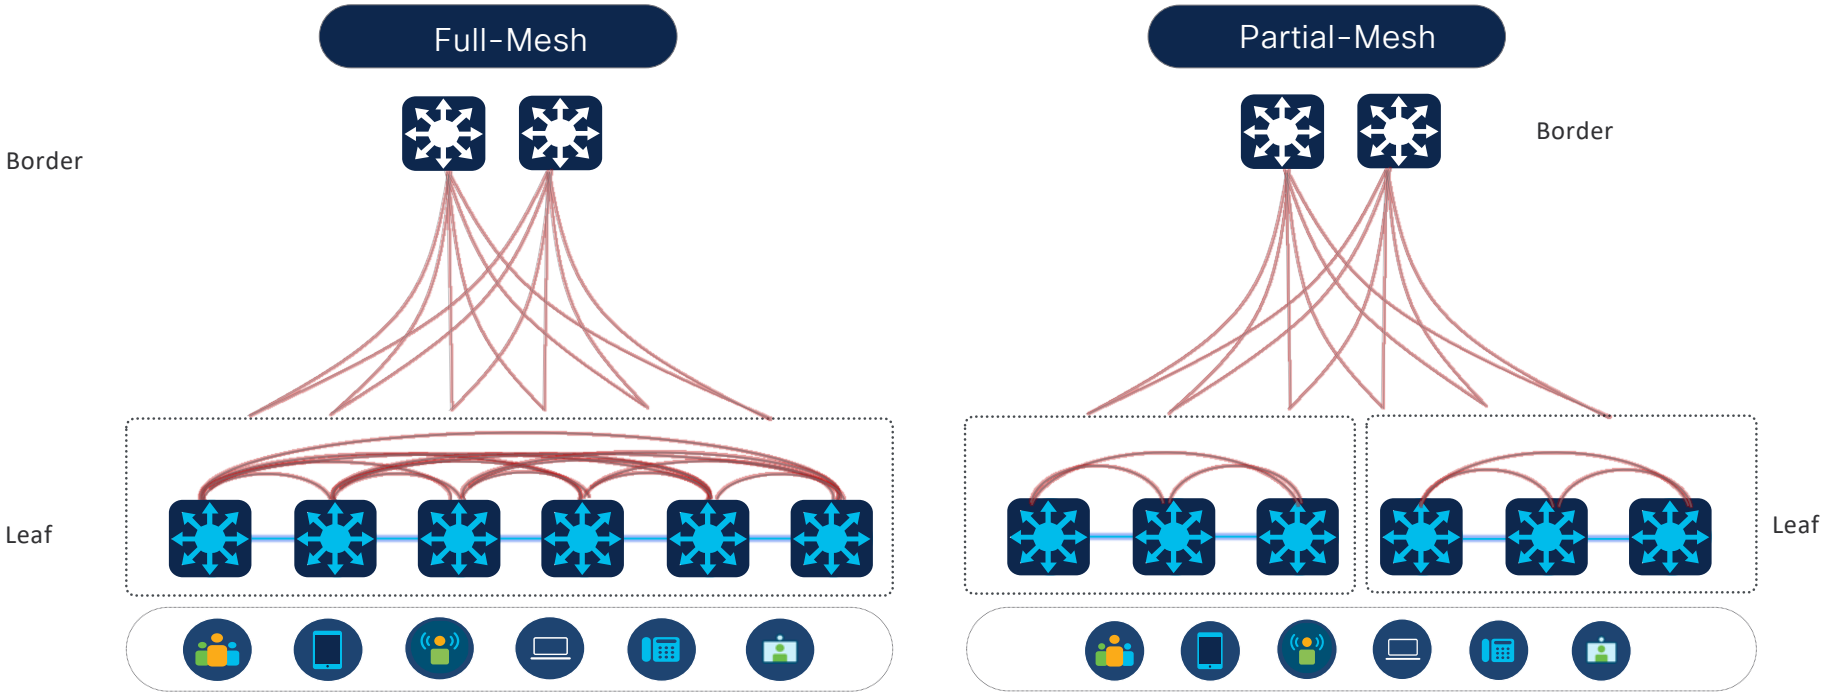

Segmentation Switch Group | Single Switch | VN | Subnet | VLAN | Port | App

Overlay Network Design- Hub-n-Spoke

Border:

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

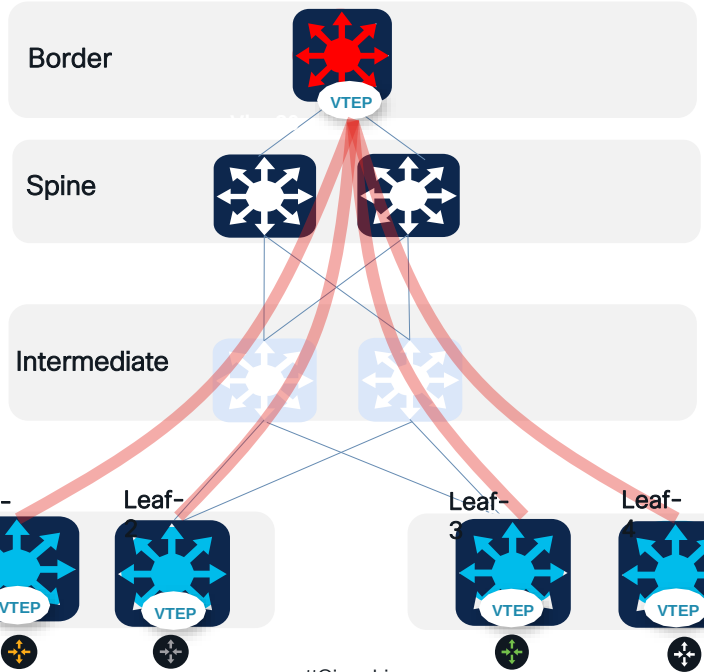

Leaf-1-Spoke:

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

Leaf-2-Spoke:

```
vrf definition S1-EVPN
 rd 3:3
 !
 address-family ipv4
  route-target import 1:1 stitching
  route-target export 2:2 stitching
  exit-address-family
```

Leaf-3-Spoke:

```
vrf definition S1-EVPN
 rd 4:4
 !
 address-family ipv4
  route-target import 1:1 stitching
  route-target export 2:2 stitching
exit-address-family
```

Leaf-4-Spoke:

```
vrf definition S1-EVPN
 rd 5:5
 !
 address-family ipv4
 route-target import 1:1 stitching
 route-target export 2:2 stitching
 exit-address-family
```

Overlay Network Design- Hub-n-Spoke

Border:

l2vpn evpn instance 11 vlan-based
encapsulation vxlan
route-target import 2:2
route-target export 1:1

Leaf-1-Spoke:

l2vpn evpn instance 11 vlan-based
encapsulation vxlan
route-target import 1:1
route-target export 2:2

Leaf-2-Spoke:

l2vpn evpn instance 11 vlan-based
encapsulation vxlan
route-target import 1:1
route-target export 2:2

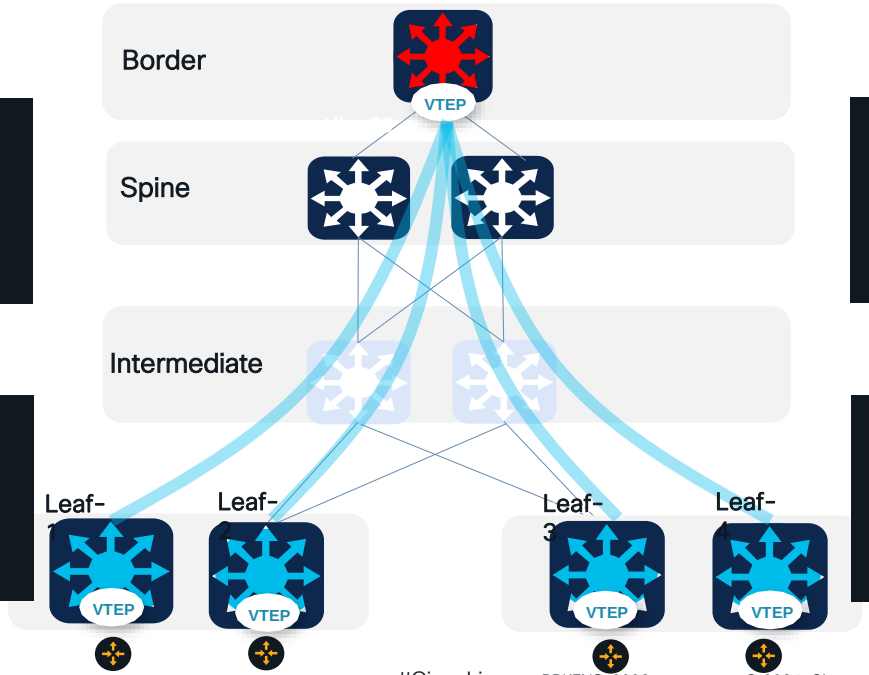

Leaf-3-Spoke:

l2vpn evpn instance 11 vlan-based
encapsulation vxlan
route-target import 1:1
route-target export 2:2

Leaf-4-Spoke:

l2vpn evpn instance 11 vlan-based
encapsulation vxlan
route-target import 1:1
route-target export 2:2

Efficient Layer 2 Broadcast domain

Point-to-Point

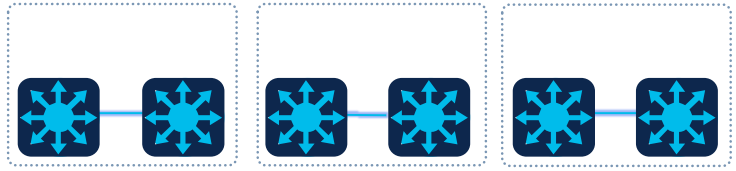

VLAN	VNI	BUM
401	40001	Ingress Replication
402	40002	
403	40003	

Hub-n-Spoke

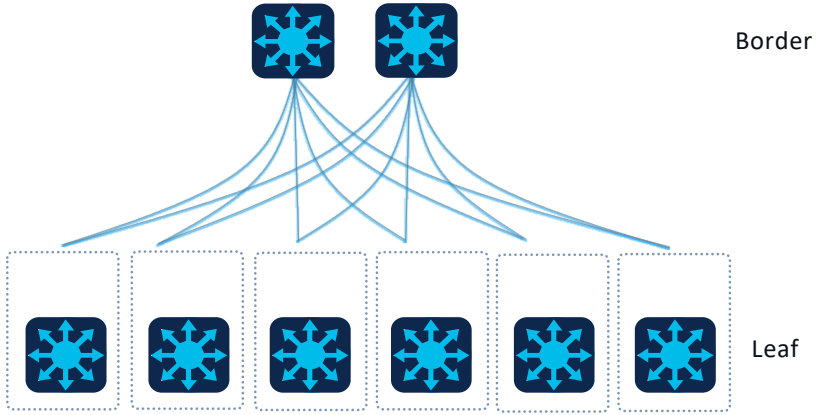

VLAN	VNI	BUM
301	30001	Ingress Replication
302	30002	
303	30003	

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

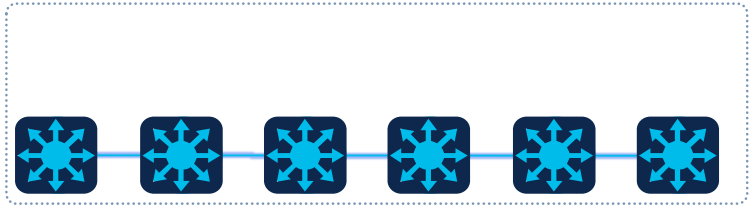

VLAN	VNI	BUM
101	10001	Multicast Group – 1
102	10002	
103	10003	

Partial – Mesh

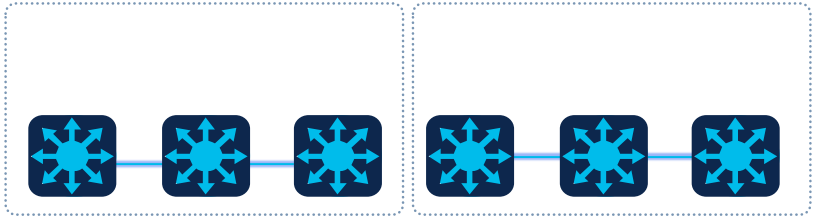

VLAN	VNI	Group
201	20001	Multicast Group – 1
202	20002	
203	20003	

VLAN	VNI	BUM
201	20001	Multicast Group – 2
202	20002	
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

Single Cluster Fabric Architecture

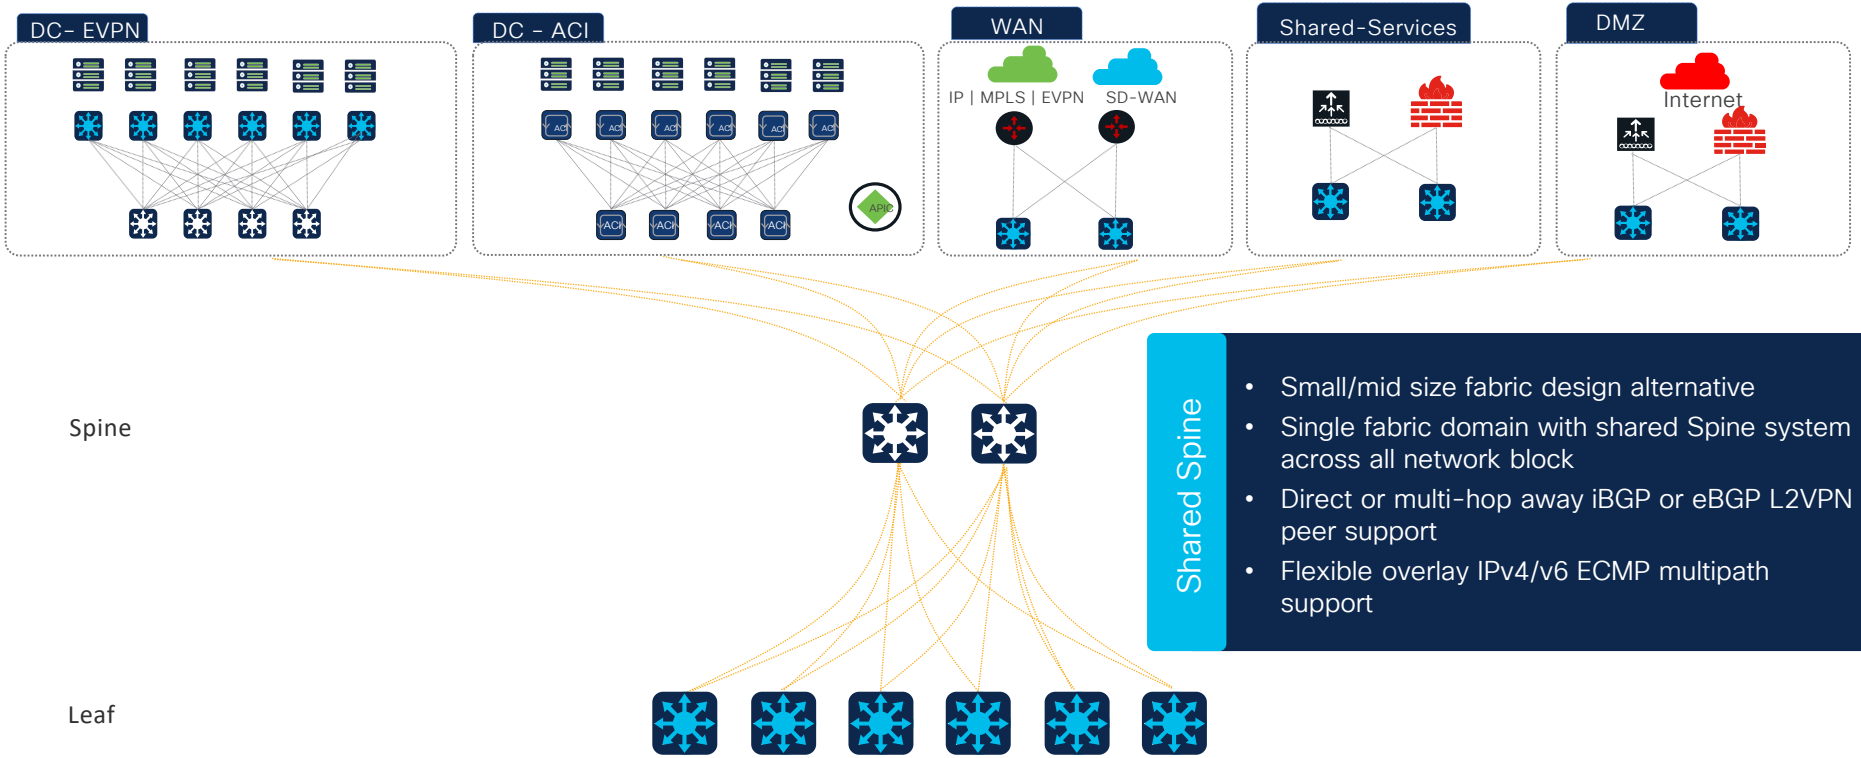

Shared Spine

- 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

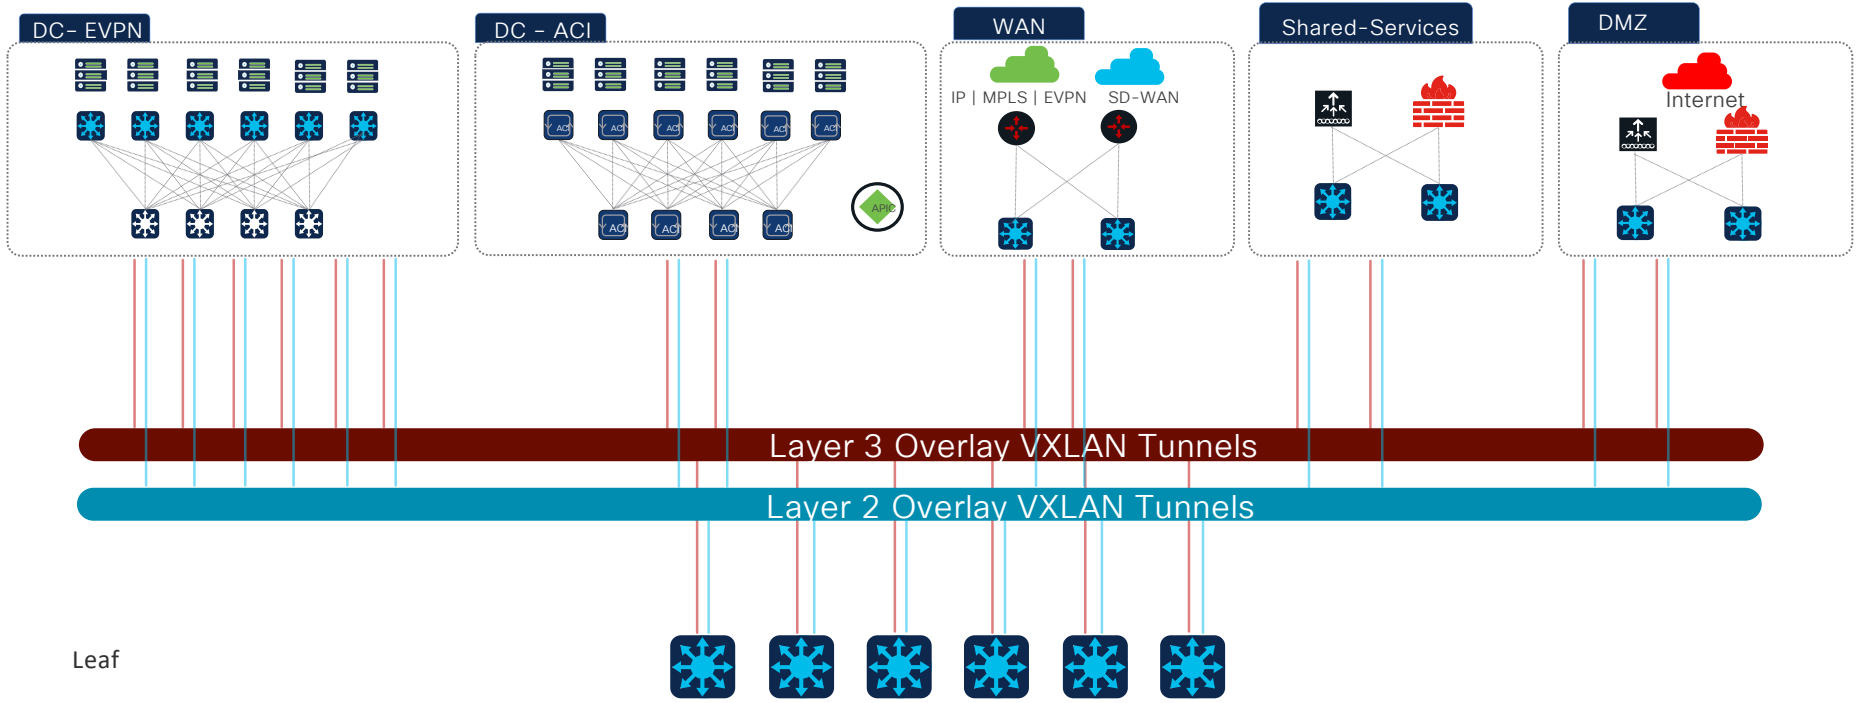

Non-Hierarchical Fabric

- 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

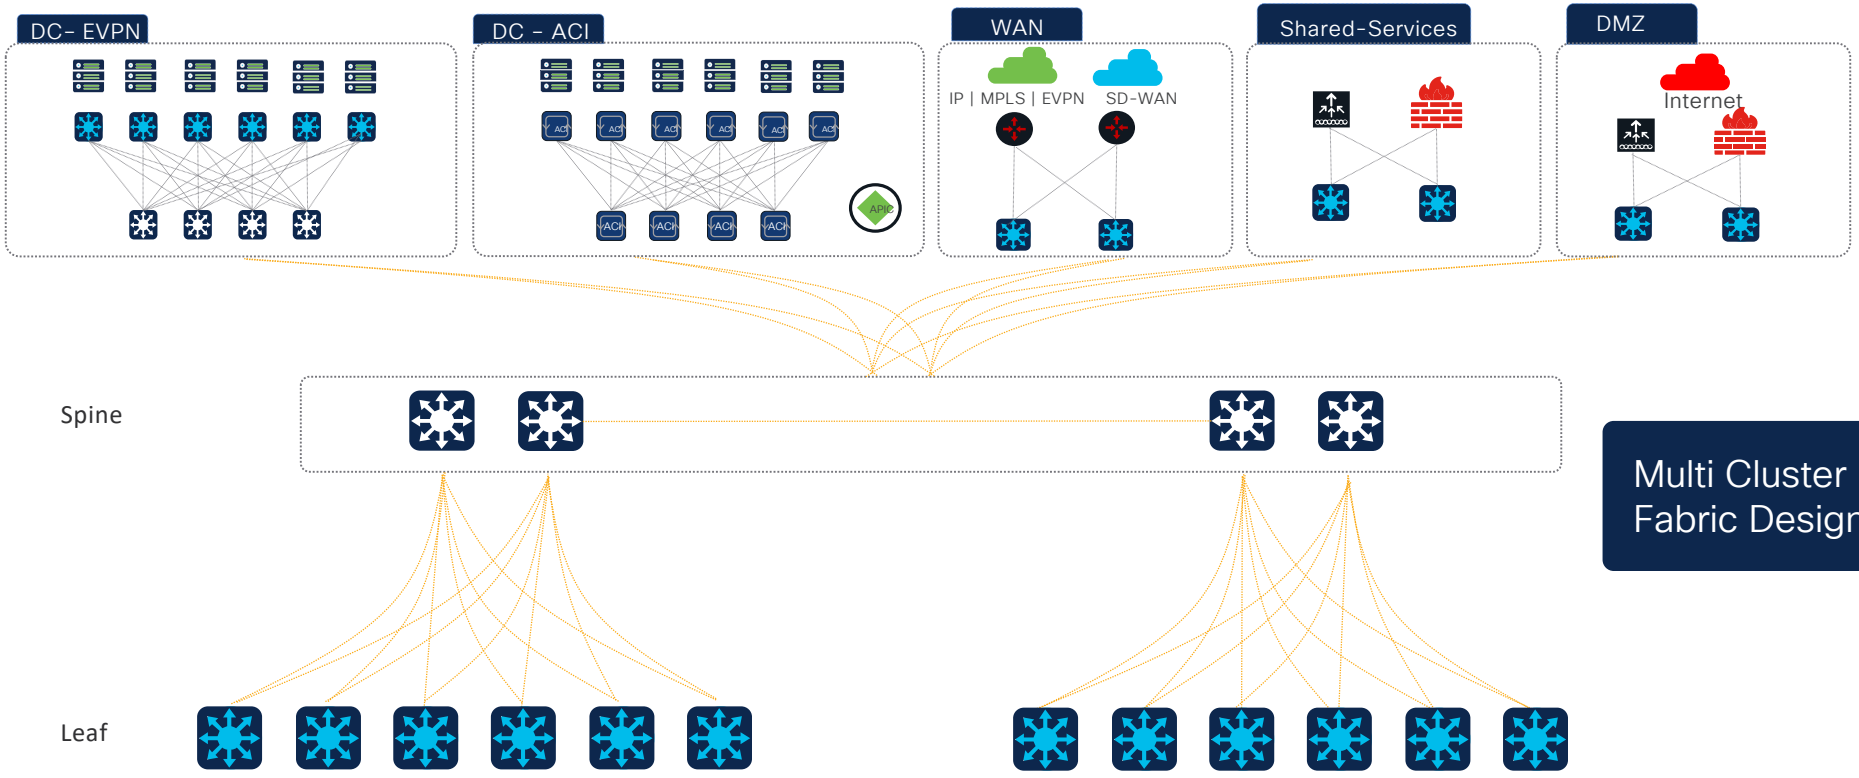

Multi Cluster Fabric Design

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.

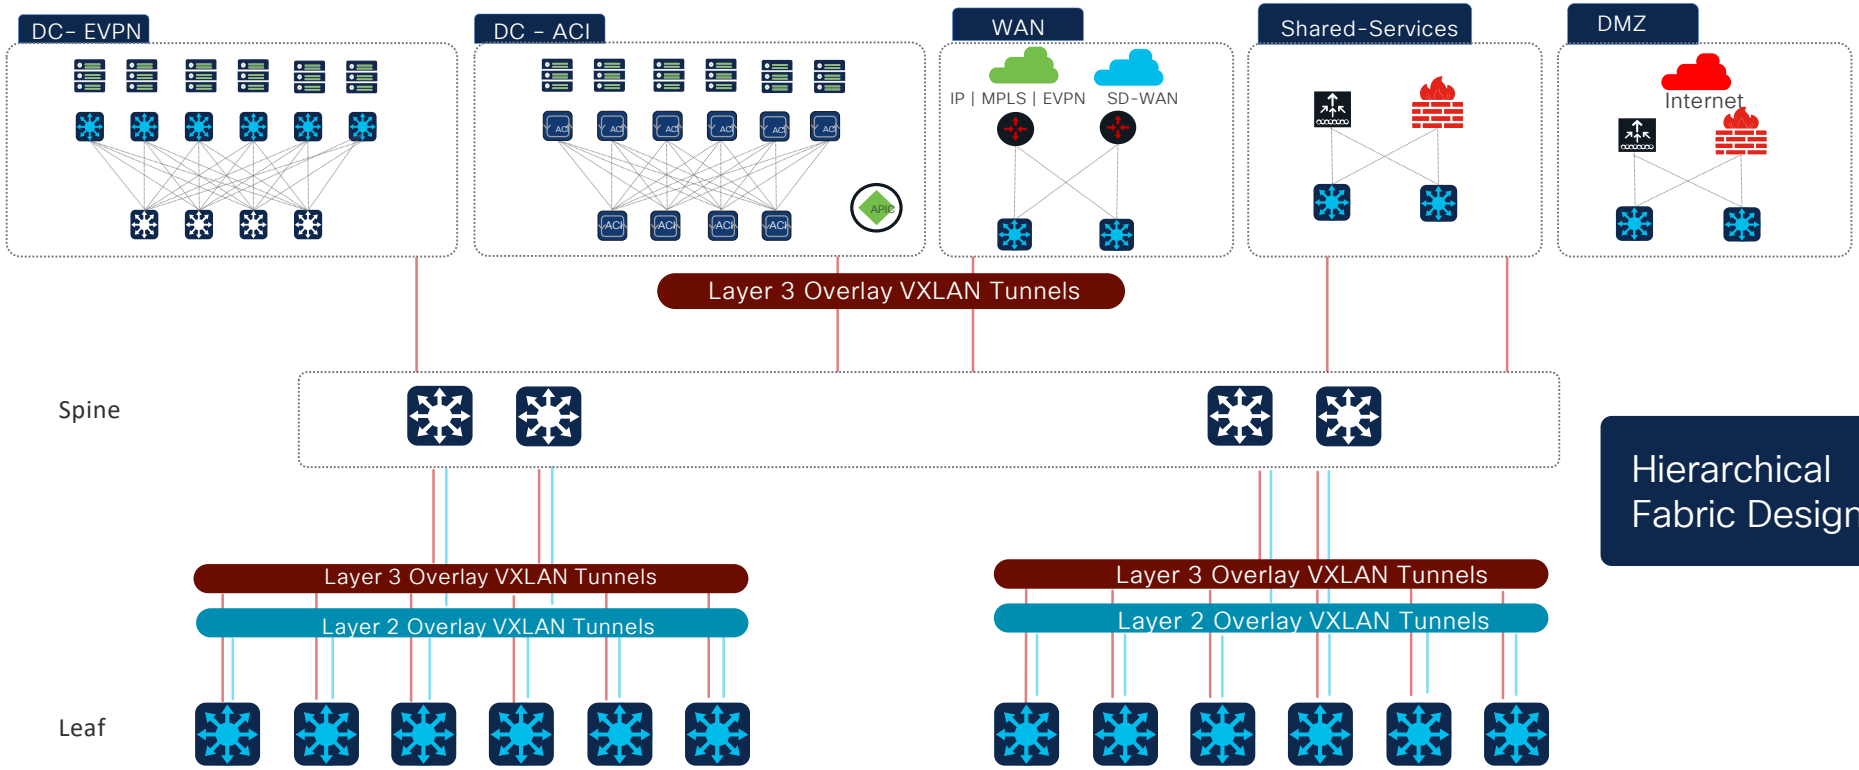

Hierarchical Fabric Design

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

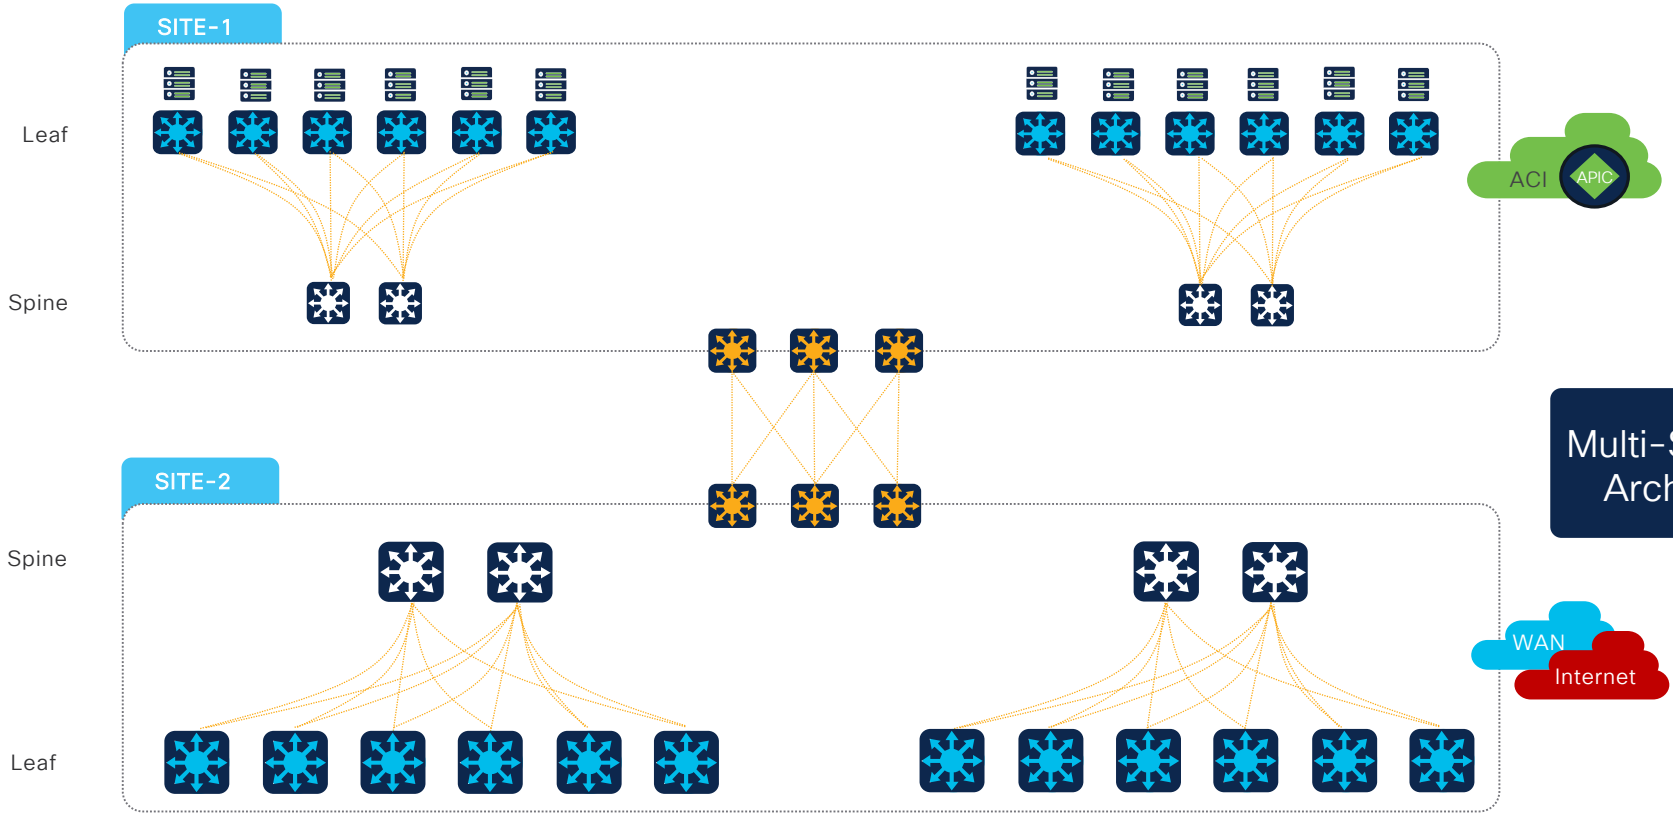

Multi-Site Fabric Architecture

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)

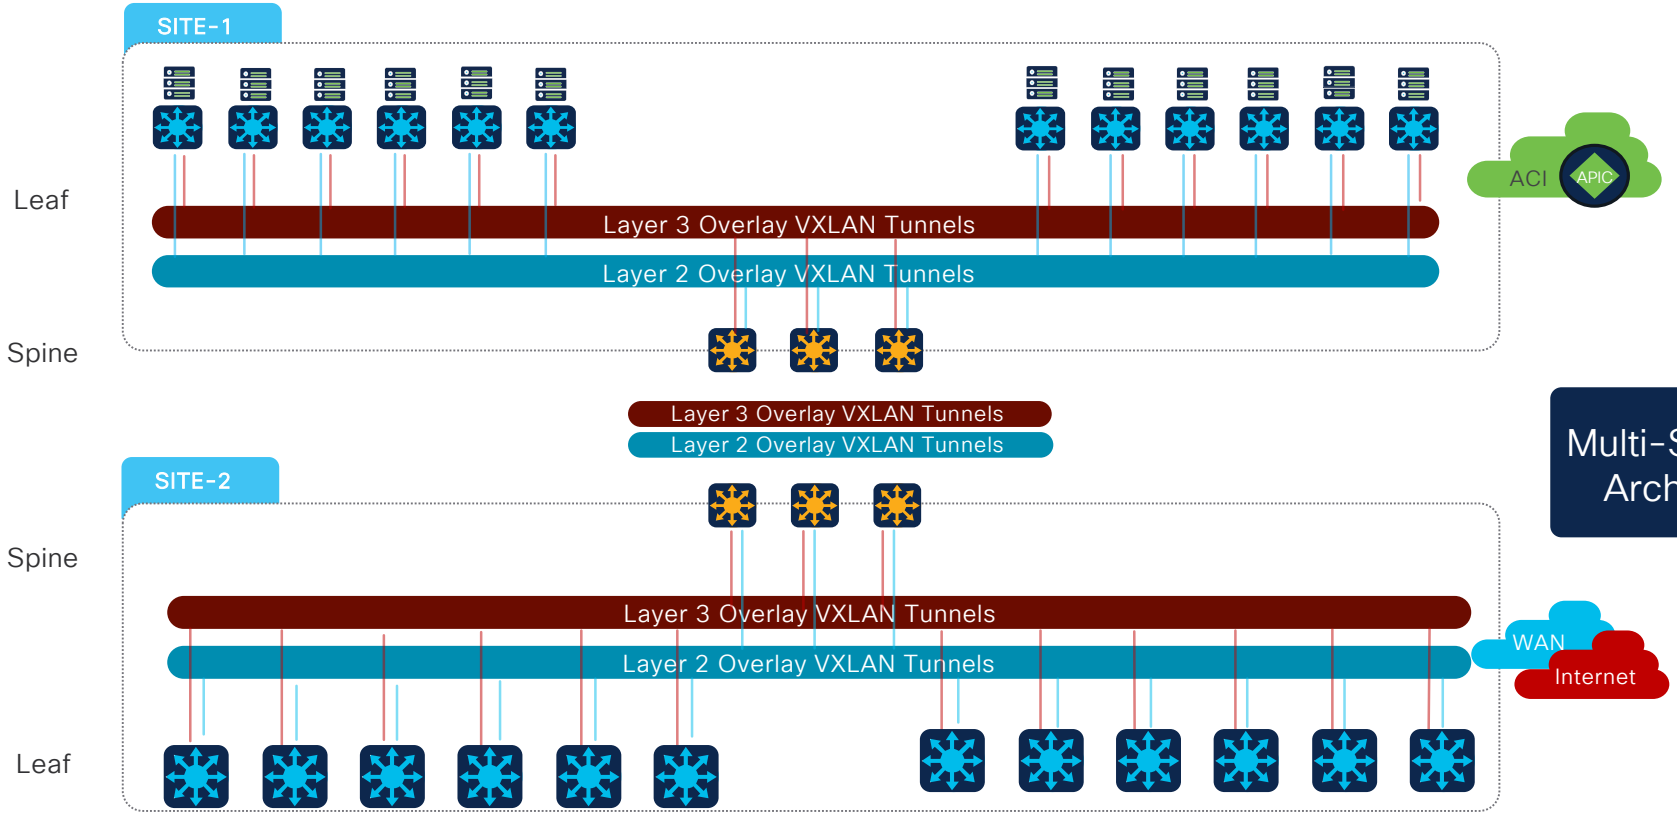

Multi-Site Fabric Architecture

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

Segmentation

Logical Local Grouping

Macro-Segmentation

Extended Group with EVPN

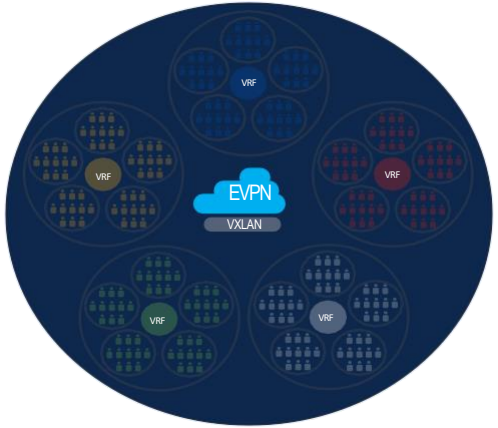

Micro-Segmentation

Policy-Plane enforced fabric

Software
Cisco IOS® XE 17.11.1

BGP EVPN – Micro Segmentation

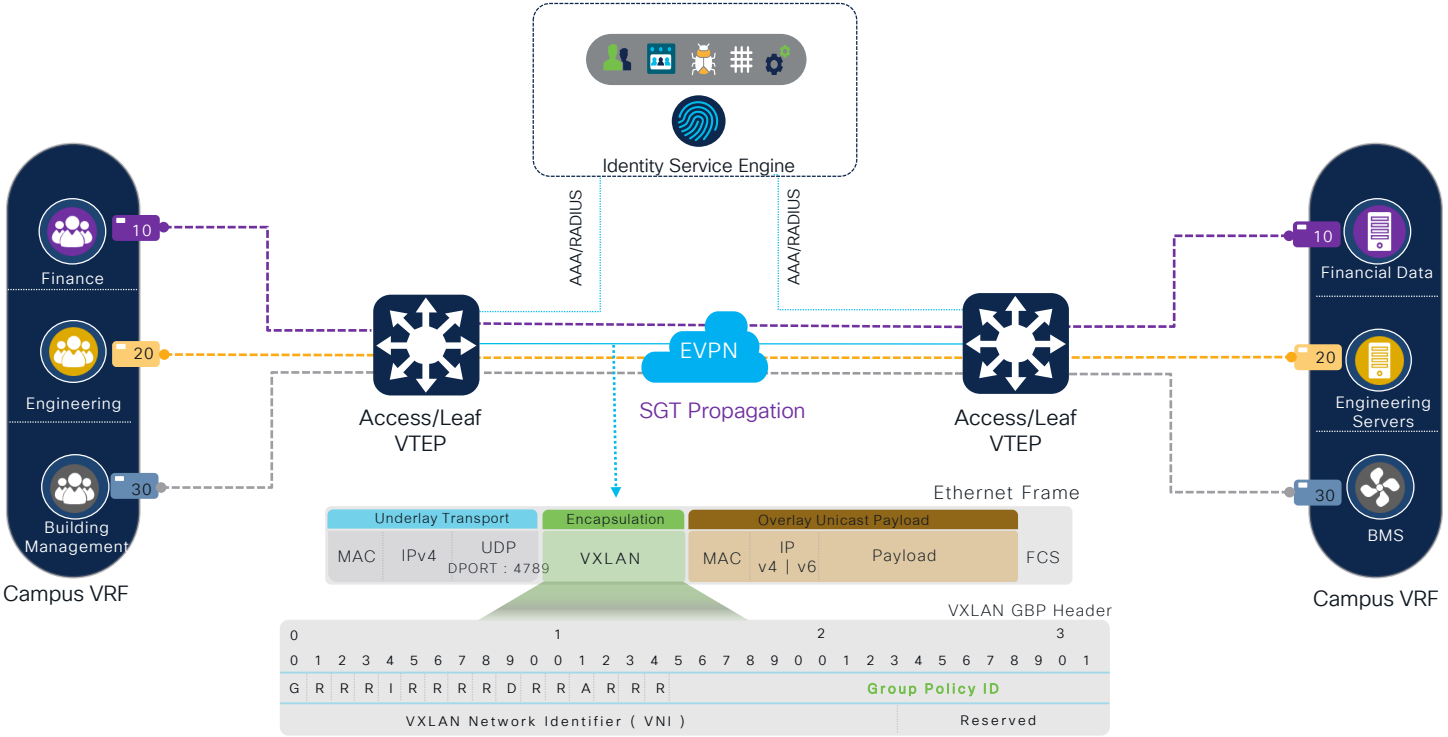

Unified Micro-Segmentation

- Seamless Secure Group Tag (SGT) integration from Wired / Wireless to BGP EVPN VXLAN fabric networks
- Cisco ISE combined with pxGrid provides Cisco and multi-vendor network and security product integration
- Dynamic or Static SGT mapping on TrustSec to VXLAN GBP support context-aware BGP EVPN Fabric with VXLAN Group-Based (GBP)
- Unified Policy Plane – Extend within and beyond EVPN fabric, Terminate to enforce and seamless Data-Center (ACI) integrations

Multicast over VXLAN

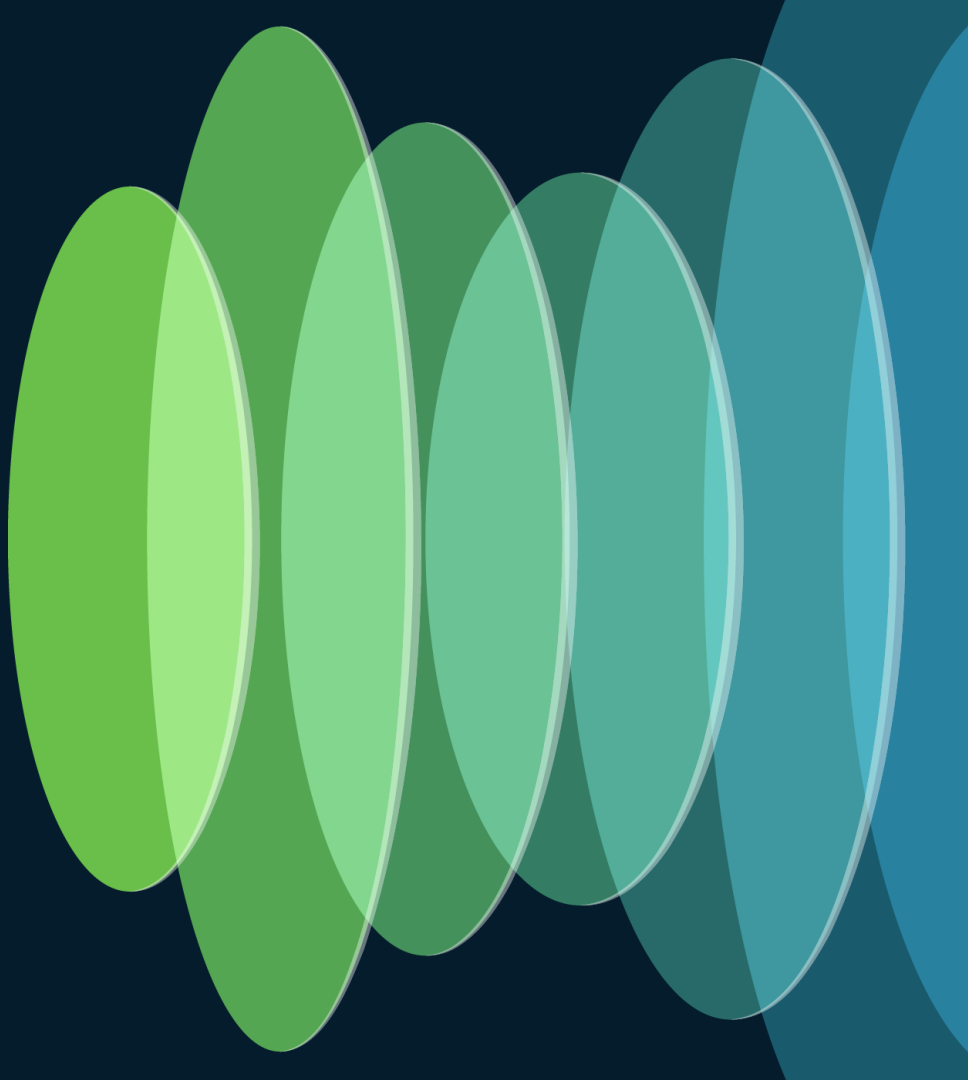

Multicast Routing

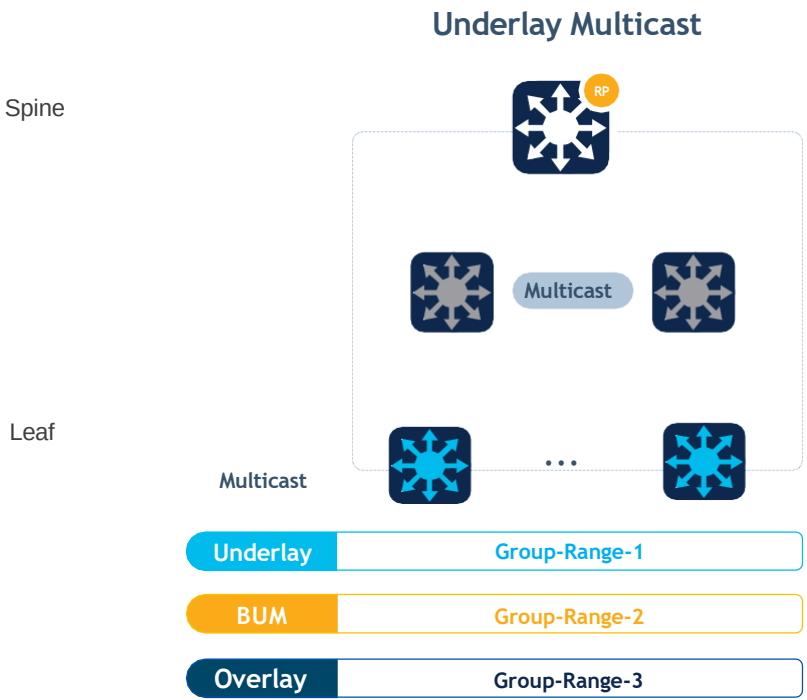

- Multicast RP integrated on Spine or separate system
- Non-overlapping Multicast Group for different purpose
- Recommended to large scale EVPN deployments
- Default MDT Group Range for Overlay TRM Multicast

Layer 3 Overlay

Distributed Anycast RP

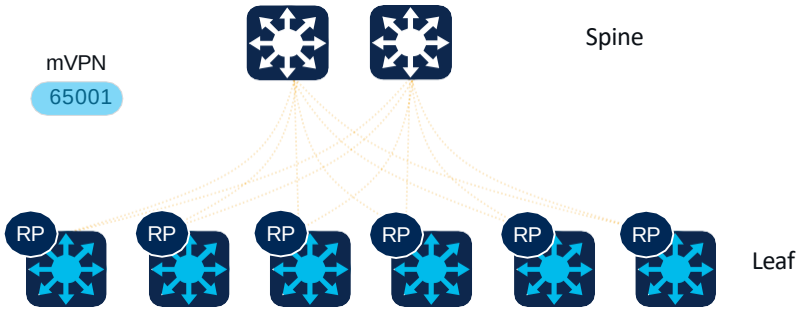

VRF	RP	IP	MDT
Blue	Anycast	Lo1: 10.1.1.101	239.1.1.101
Yellow	Anycast	Lo2: 10.2.1.101	239.2.1.101
Green	Anycast	Lo3: 10.3.1.101	239.3.1.101

Layer 3 Overlay

Fabric Border RP

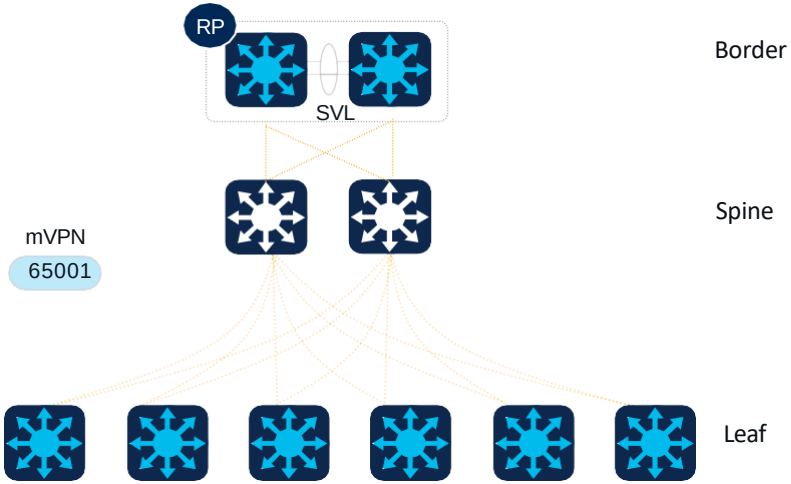

VRF	RP	IP	MDT
Blue	Anycast	Lo1: 10.1.1.101	239.1.1.101
Yellow	Anycast	Lo2: 10.2.1.101	239.2.1.101
Green	Anycast	Lo3: 10.3.1.101	239.3.1.101

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

TRM Default MDT

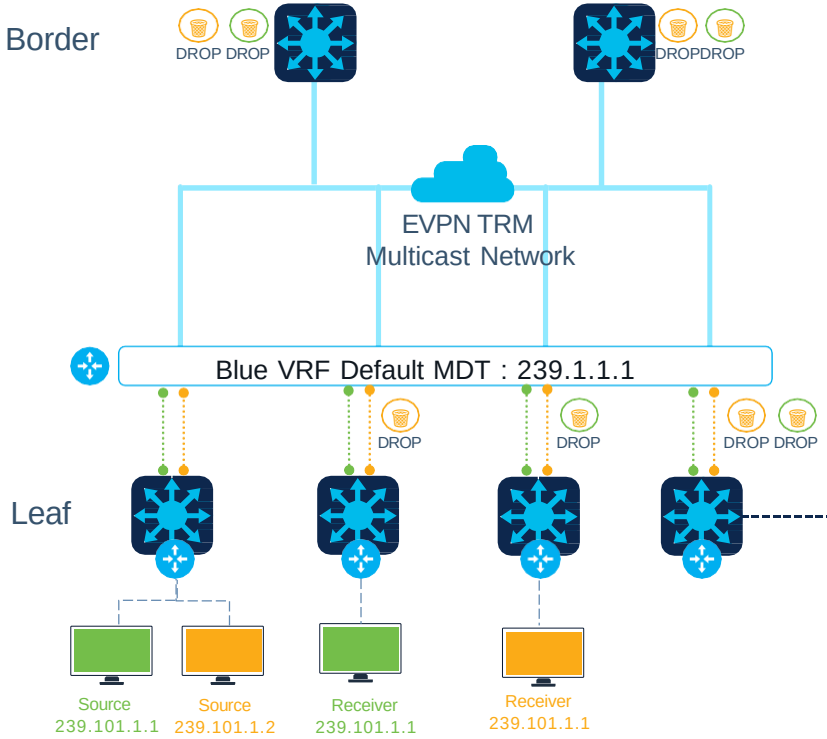


```
vrf definition S1-EVPN
 rd 10:10
 !
 address-family ipv4
  mdt auto-discovery vxlan inter-as
  mdt default vxlan 239.1.1.1 → MDT Default
  mdt overlay use-bgp spt-only
```

Challenges

Non-selective overlay Multicast replication

Inessential core network bandwidth utilization

Redundant system resources utilization

Limited scale for dense network environment

TRM Data MDT

Border

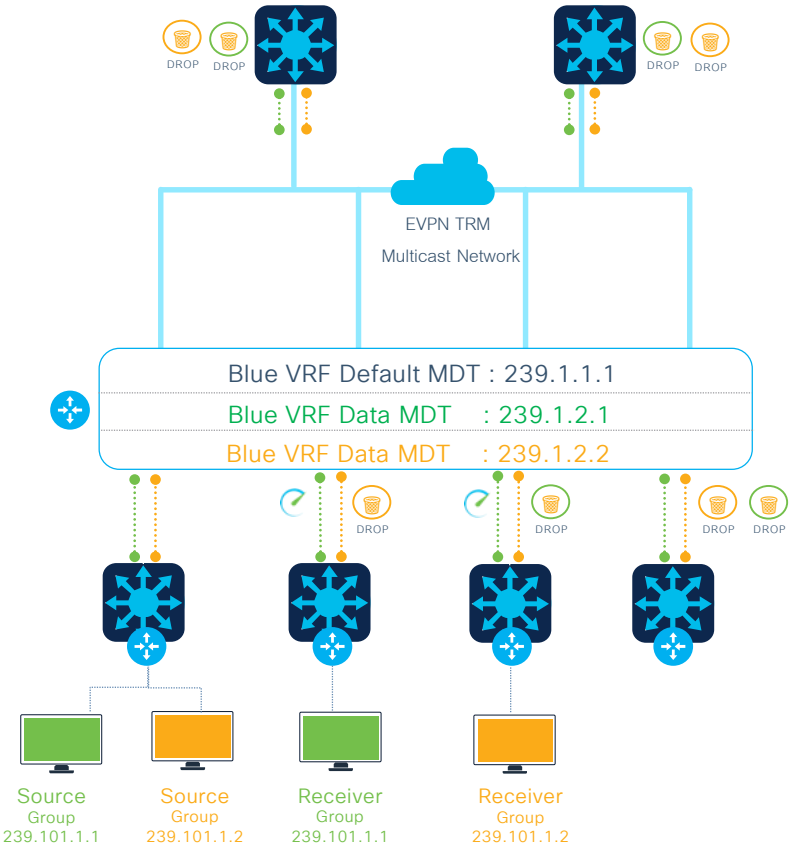

Leaf

Key Benefits

Stateful L2 Multicast Overlay network

Industry-standard based control-plane

Applicable to Centralized Gateway or Cross-Connect Overlay networks

Scale. Performance. Security.

TRM Data MDT

Border

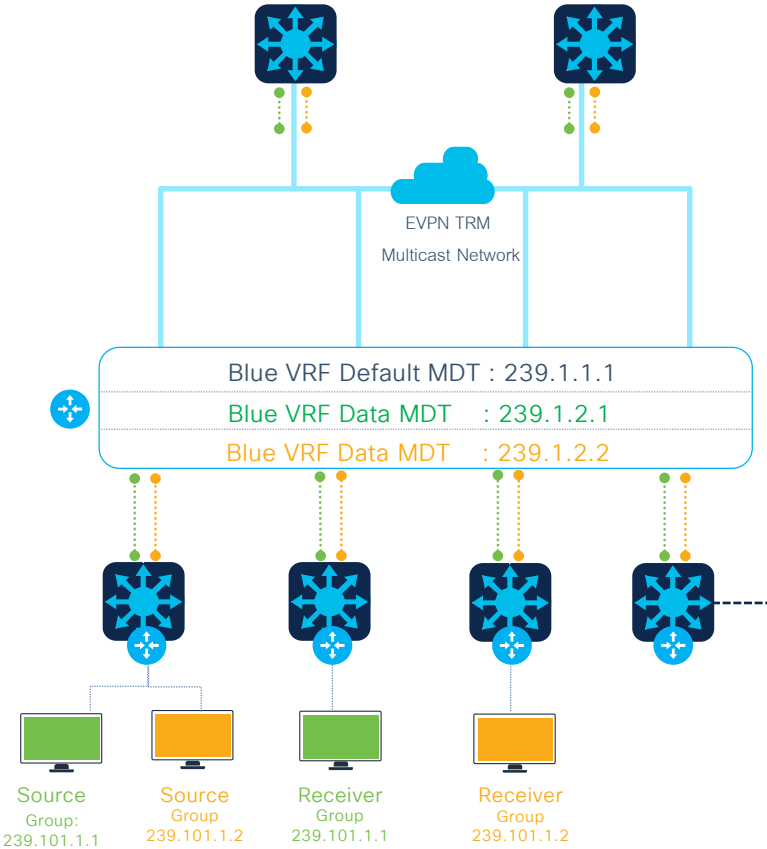

Leaf

```
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
```

EVPN Fabric Interworking

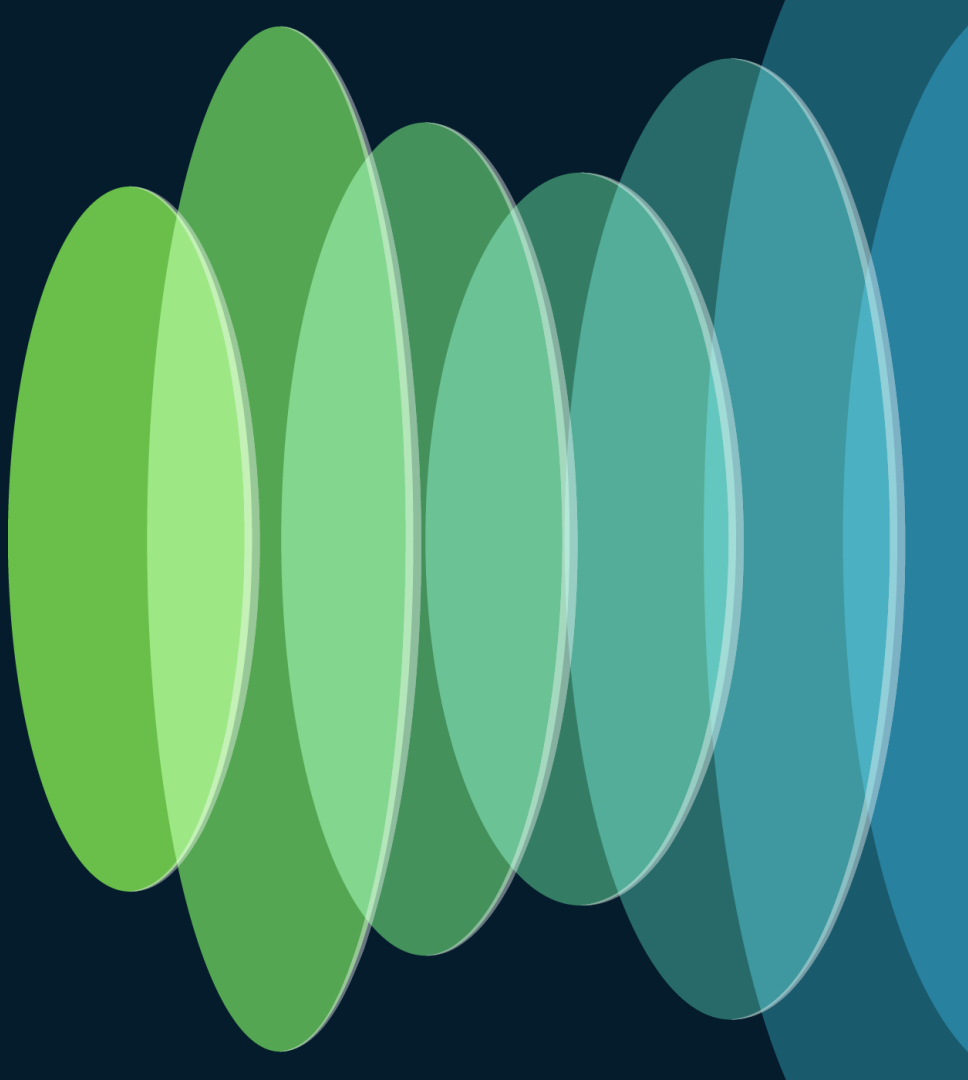

cisco Live!

External Domain Handoff Types

Terminate

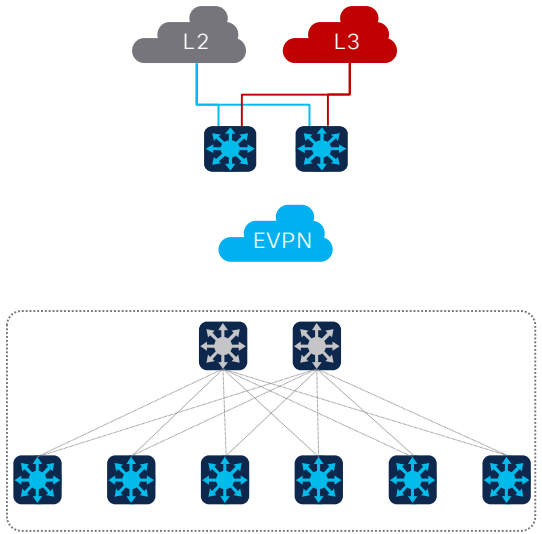

BGP EVPN fabric termination at Border

Simple Layer 2 / Layer 3 hand off

Layer 3 VRF segmentation to L3 system

L2 extension handoff, only if needed.

Interworking

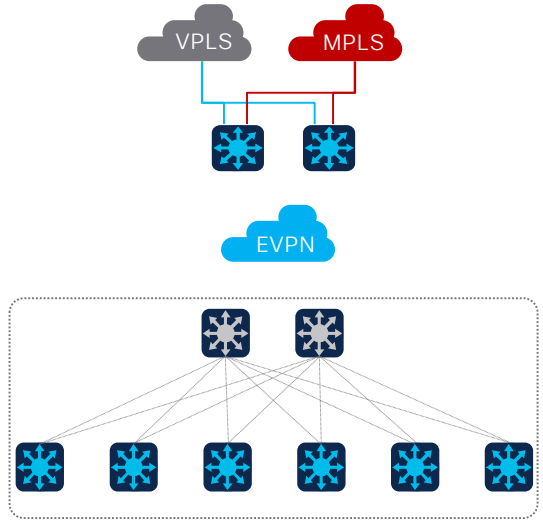

Integrated fabric interworking at Border

Seamless EVPN & classic overlay ”stitching”

End-to-End network segmentation

Loop-free Layer 2 overlays across domains

Re-Originate

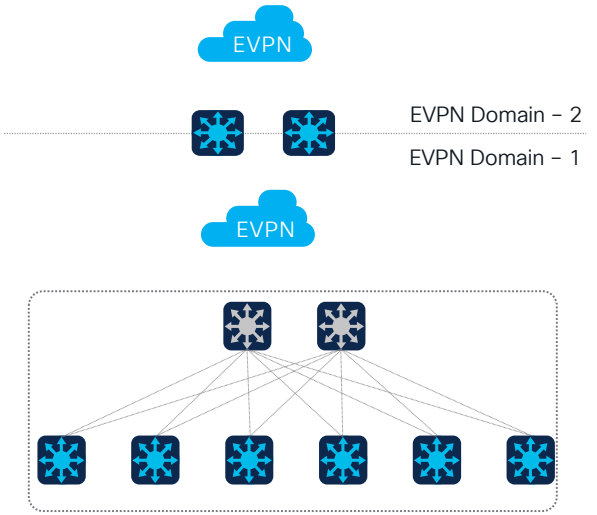

BGP EVPN fabric re-origination at Border

L3 segmentation between fabric domains

Can collapse with Border/Spine role

L2 and Multicast in overlay unsupported

Layer – 2 Handoffs Alternatives

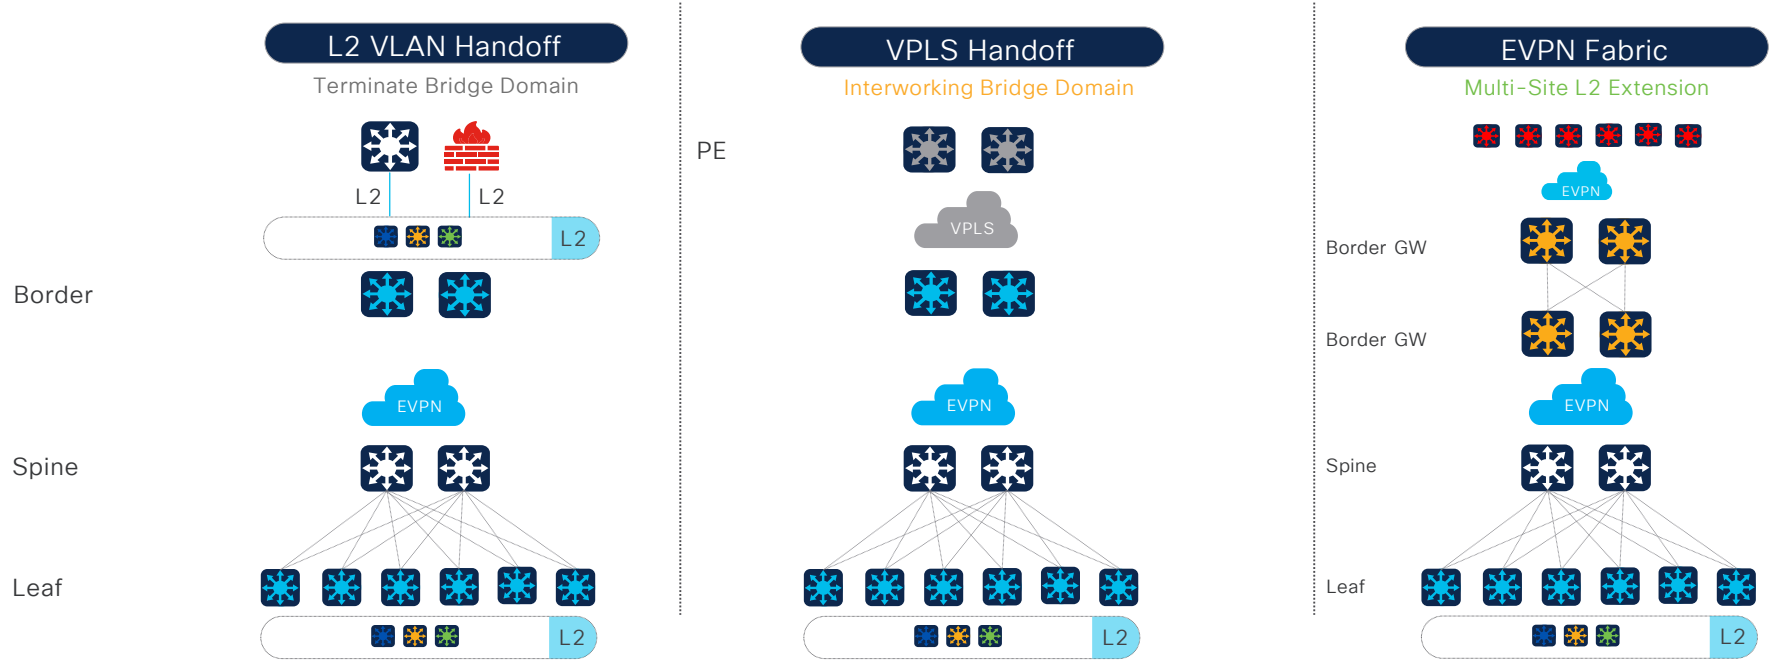

Seamless Layer 2 Handoff

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 Alternatives

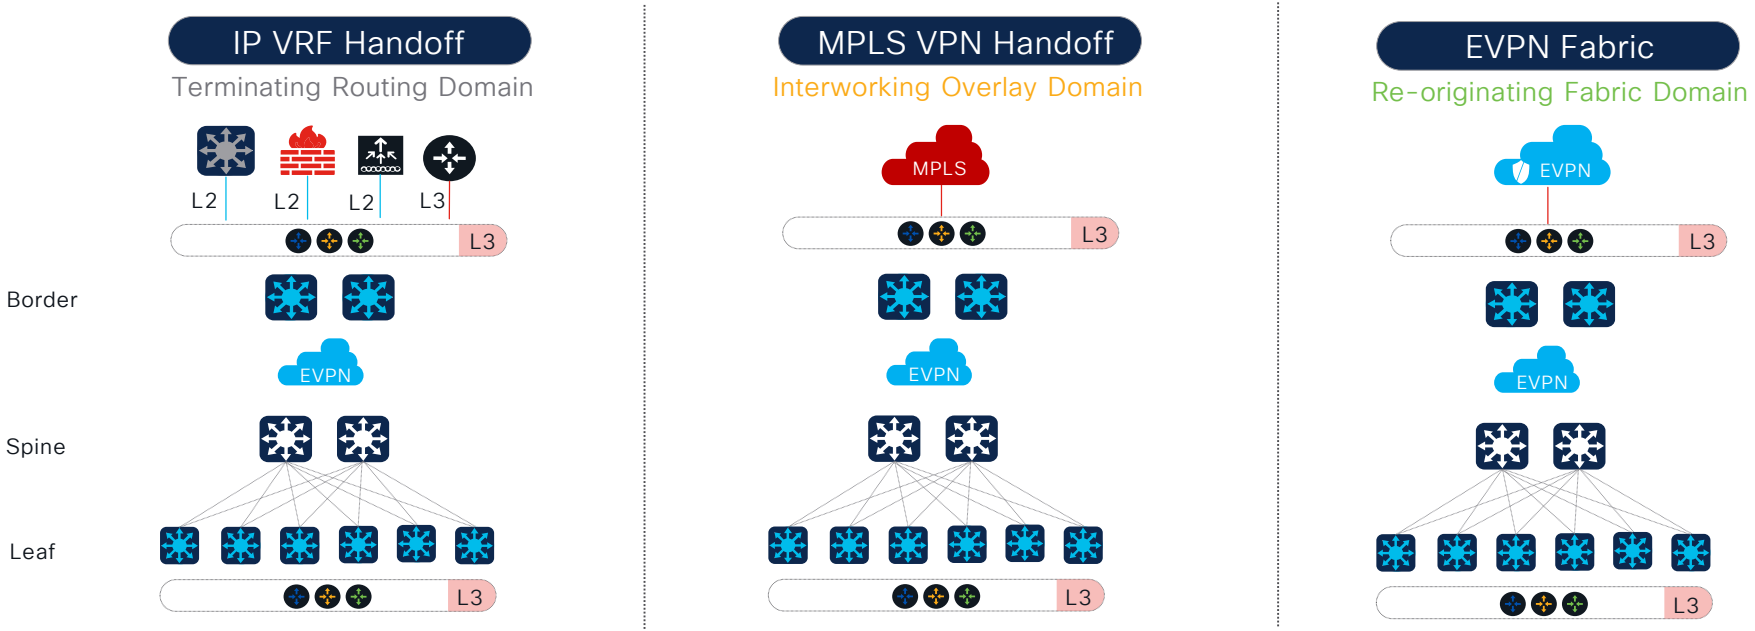

Integrated Extranet

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

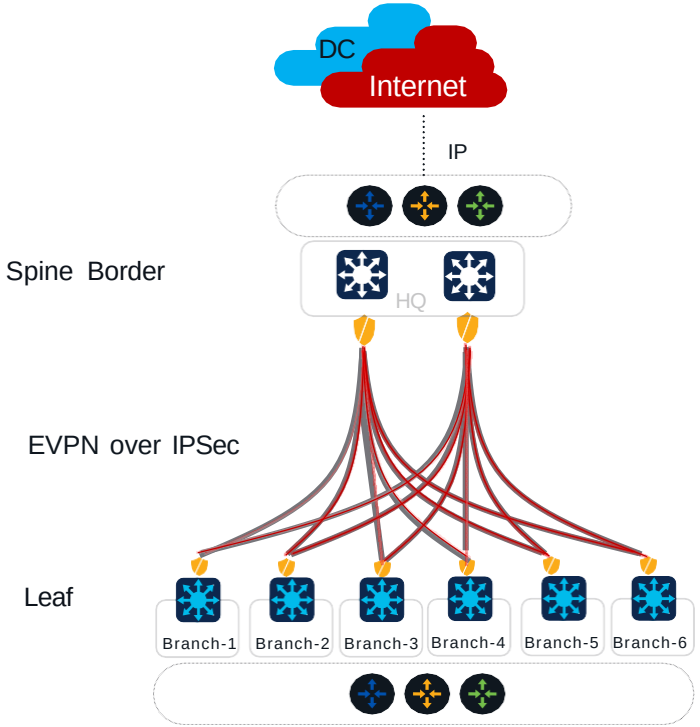

Secure Encrypted Fabric

Multi Cluster – Hierarchical Fabric

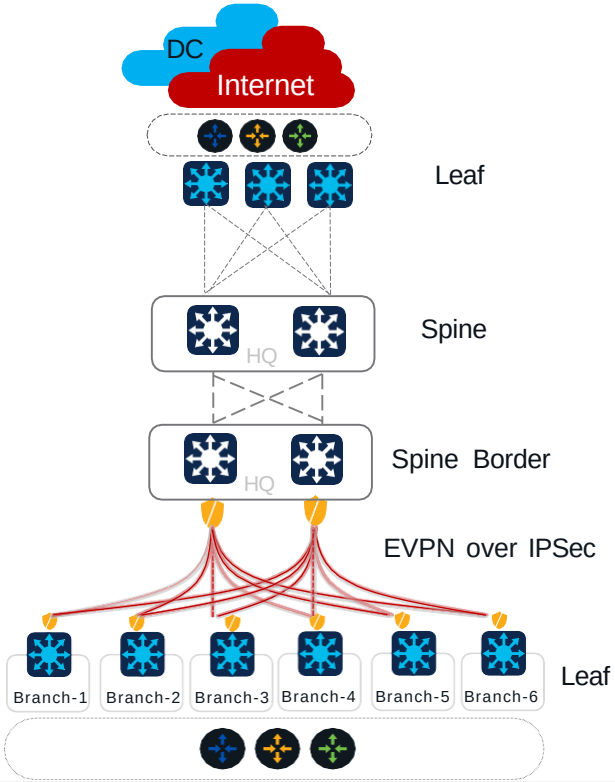

Encrypted EVPN Fabric

- 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

Fabric Deployment Options

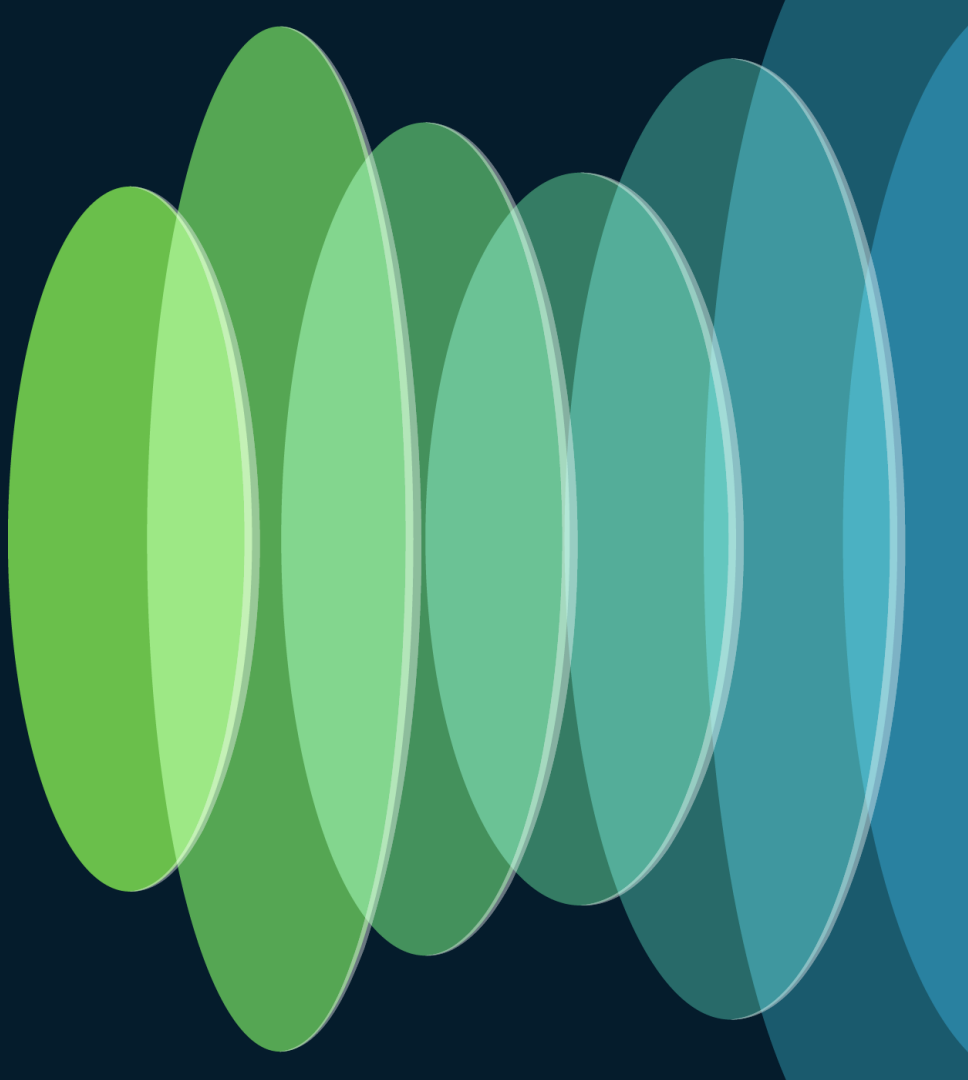

Cisco Enterprise BGP EVPN Solution

Do-It-Yourself

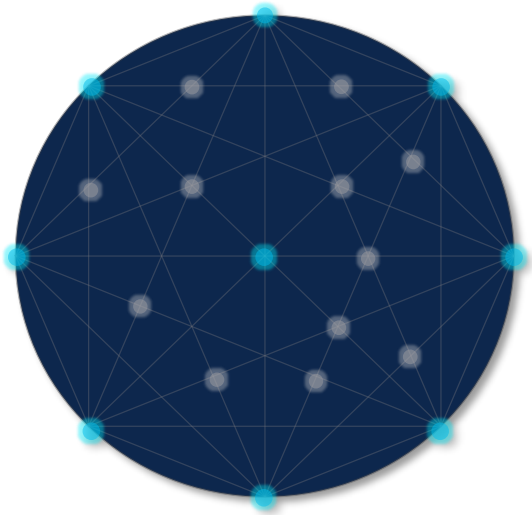

Programmable

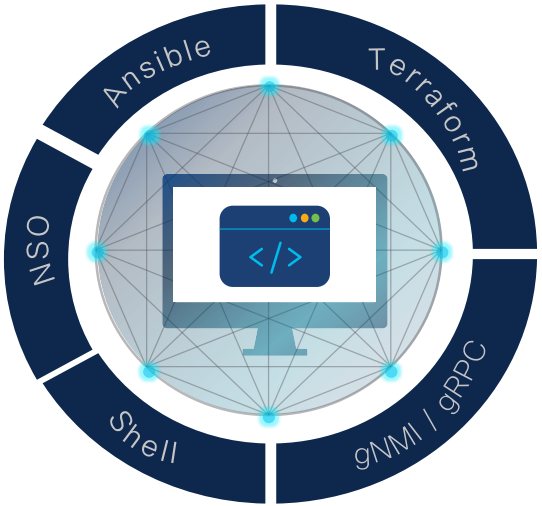

Intent-Based

Q2CY24

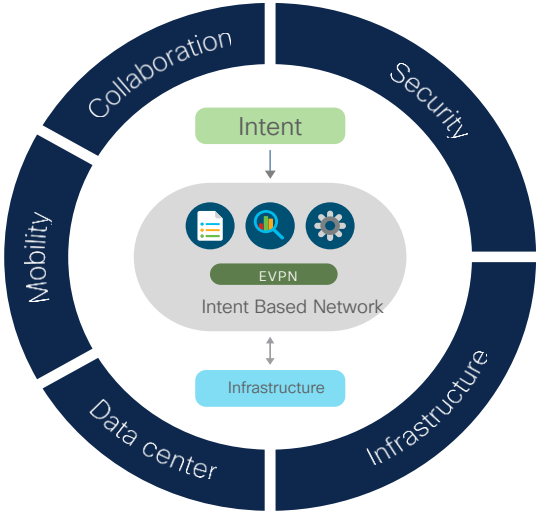

BGP EVPN – CLI Simplification

L2 Overlay Provisioning

```
l2vpn evpn
 replication-type ingress

interface nve1
 source-interface Loopback0
 host-reachability protocol bgp
```
Global

```
!
l2vpn evpn instance 10 vlan-based
 encapsulation vxlan
 replication-type ingress
l2vpn evpn instance 12 vlan-based
 encapsulation vxlan
 replication-type ingress
```
MAC VRF

```
!
interface nve1
 member vni 20010 ingress-replication
 member vni 20012 ingress-replication
```
L2vni – nve

```
!
vlan configuration 10
 member evpn-instance 10 vni 20010
vlan configuration 12
 member evpn-instance 12 vni 20012
```
vlan – evpn l2 service

current config

```
l2vpn evpn
 route-target auto vni
 profile l2vpn default
    encapsulation vxlan
 exit
!
interface nve1
 source-interface Loopback0
 host-reachability protocol bgp
!
vlan configuration 10,12
 member evpn-instance
```

simplified config

- auto rt and profile – global config
- MAC vrf is auto created
- L2 VNI is auto allocated
- Auto allocate l2VN's to evpn enabled Vlan's

- EVI Auto-RD: Router-ID:EVI | EVI Auto-RT: ASN:EVI
- evi-base and l2vni-base to modify the default evi and l2vni values

Inter-operable with existing cli model

BGP EVPN – CLI Simplification

L3 Overlay Provisioning

```
vrf definition purple                                   ip vrf
 rd 172.168.1.1:1111
 !
 address-family ipv4
  route-target export 64512:30500
  route-target import 64512:30500
  route-target export 64512:30500 stitching
  route-target import 64512:30500 stitching
 !
 address-family ipv6
  route-target export 64512:30500
  route-target import 64512:30500
  route-target export 64512:30500 stitching
  route-target import 64512:30500 stitching
!
interface nve1                                          l3vni - nve
 source-interface Loopback0
 host reachability bgp
 member vni 30500 vrf BLUE
!
vlan configuration 500                                  core vlan
 member vni 30500
!
interface Vlan500                                       core svi
 description Core Vlan for VRF BLUE
 vrf forwarding BLUE
 ip unnumbered Loopback0
 no autostate
```

current config

```
vrf rd auto → global command
!
vrf definition purple
  vnid 30500 evpn-instance vxlan core-vlan 500
!
 address-family ipv4
 address-family ipv6
!
interface nve1 → one time command
 source-interface Loopback0
 host reachability bgp
```

simplified config

- auto rd (One time Global command for all VRF)
- L3 VRF RT's and RD are auto allocated
- L3 VNI auto created under NVE without any explicit config
- Core VLAN and SVI are auto created

Inter-operable with existing cli model

BGP EVPN Automation – Ansible & Terraform

Feature + Solution Playbooks

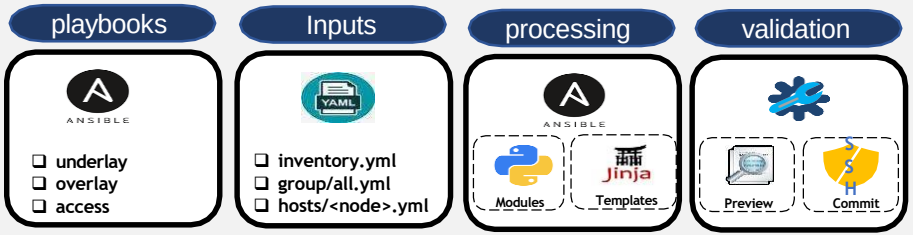

Day 0: Leaf/Spine/Border Design, DAG, VRF's, VN
Day 1: Incremental changes – add/delete IP/MAC VRF/VTEP/Access int
Day N: L3TRM, IPv6 etc

Solution Playbooks

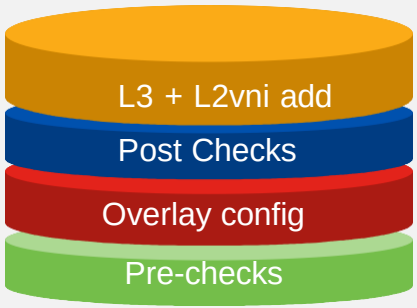

ANSIBLE

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

EVPN Terraform Provider

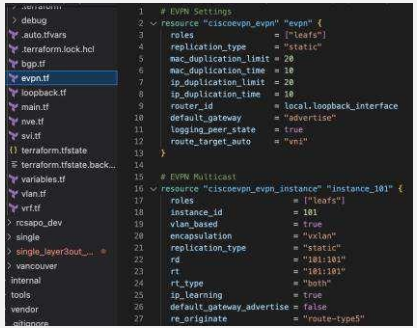

Choose Your Fabric Control-Plane

(LISP or BGP EVPN)

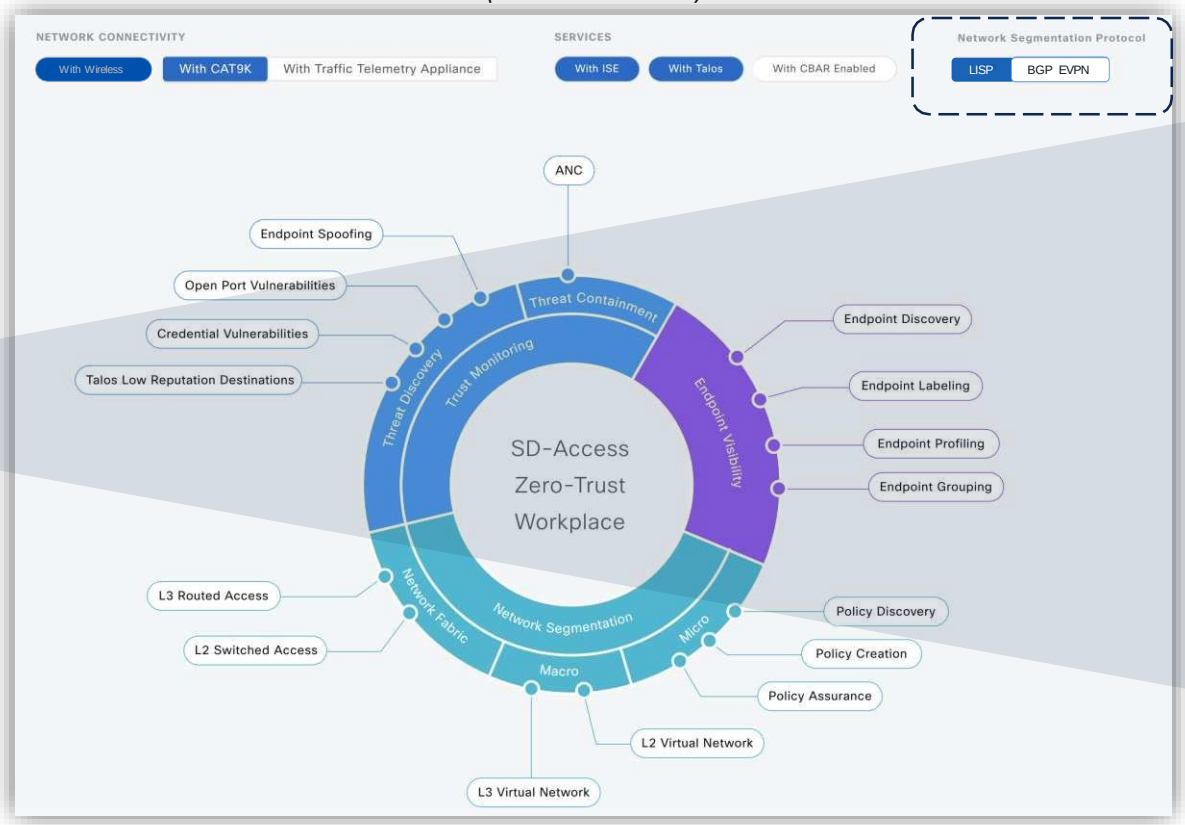

Single Data-Plane Vxlan

IBN based workflows

Automation

Assurance

Macro/Micro Segmentation

Seamless experience irrespective of choice of protocol

- Either LISP or BGP EVPN Control Plane

SD-Access with BGP EVPN: Catalyst Center and IOS-XE

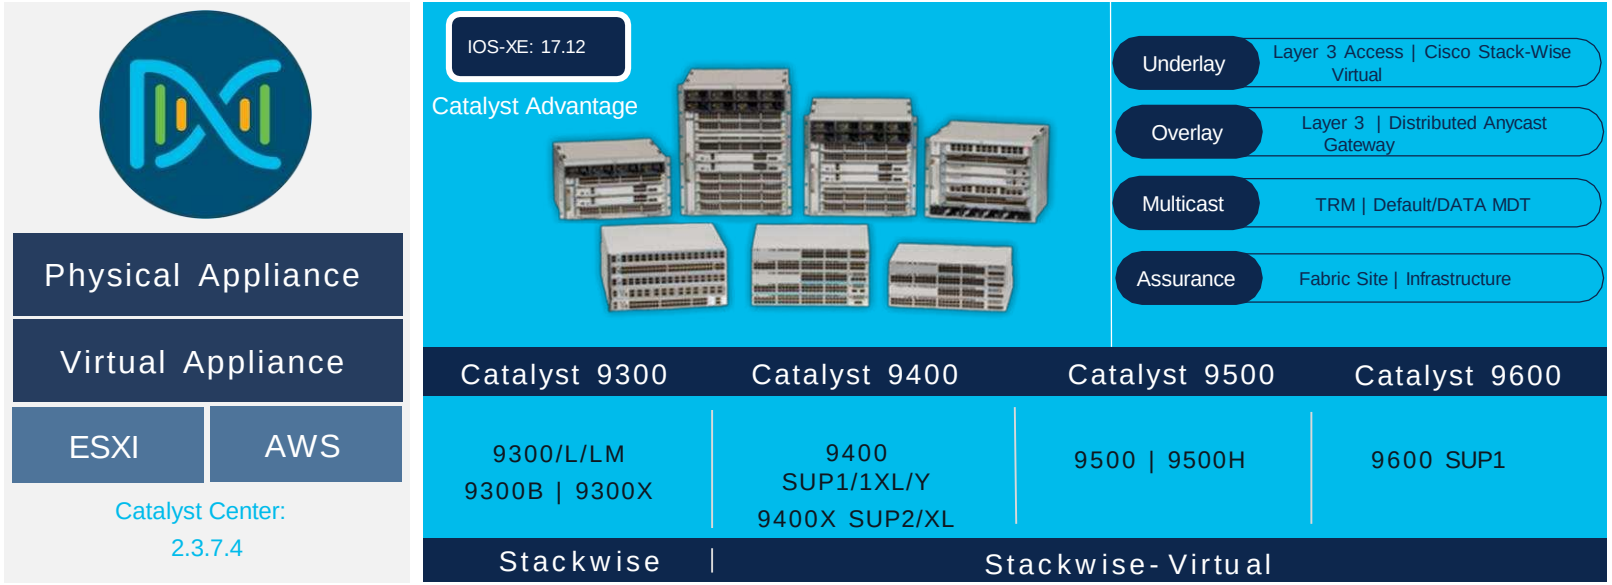

SD-Access with BGP EVPN Catalyst Center and IOS-XE

Early Field Trial

Catalyst 9000 EVPN Reference

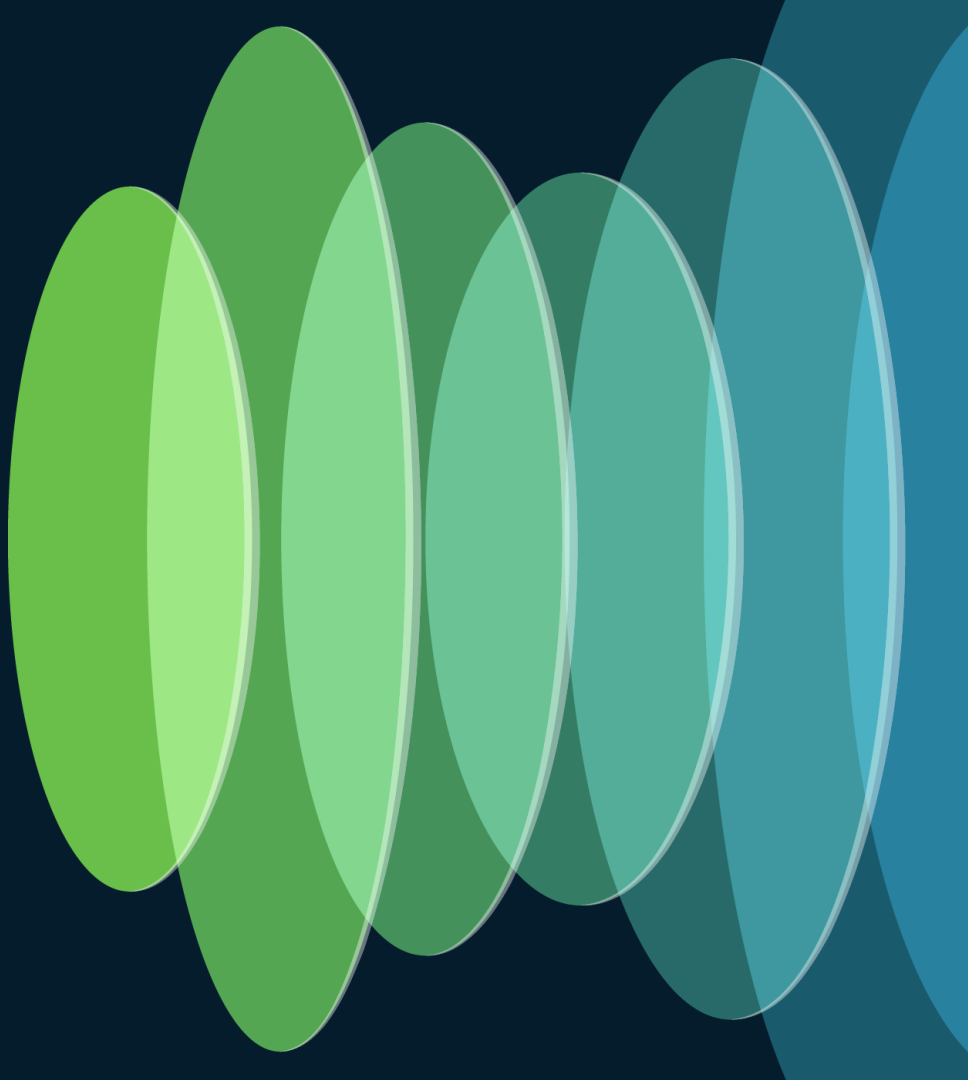

cisco Live!

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

More Coming Soon …

Reference
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/17-13/configuration_guide/vxlan/b_1713_bgp_evpn_vxlan_9300_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

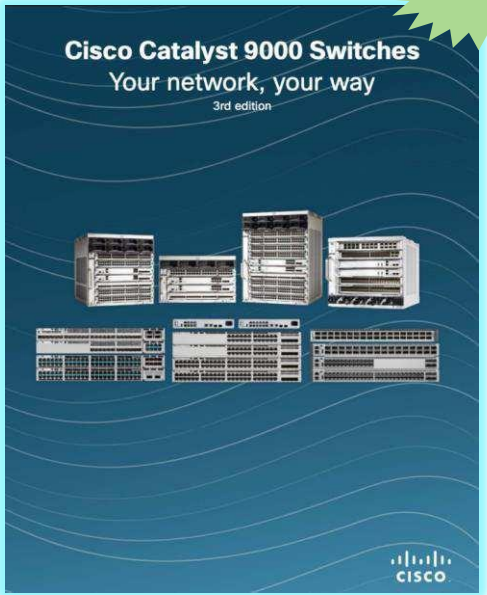

cs.co/cat9kbook

Cisco Catalyst TV @ YouTube

SUBSCRIBE

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.

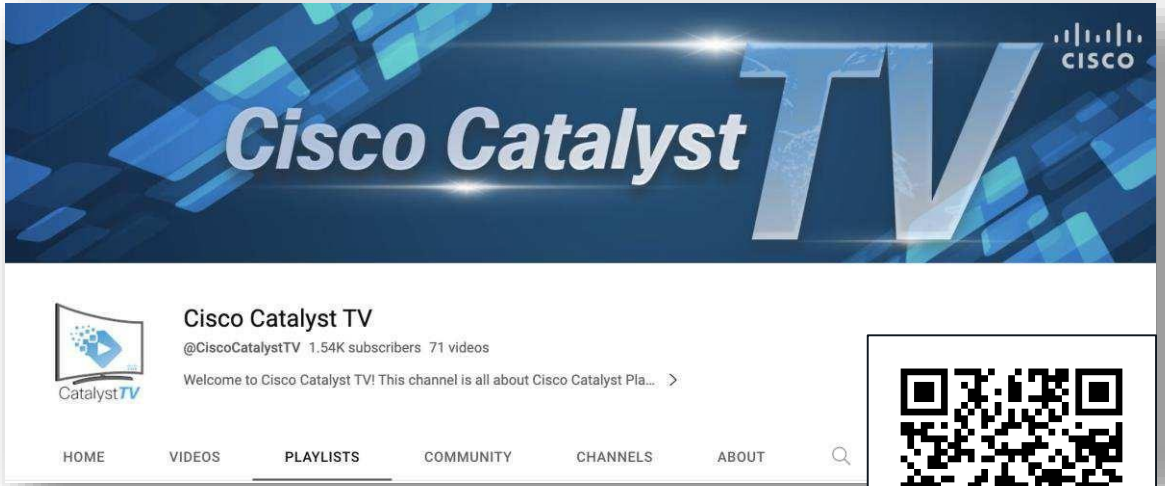

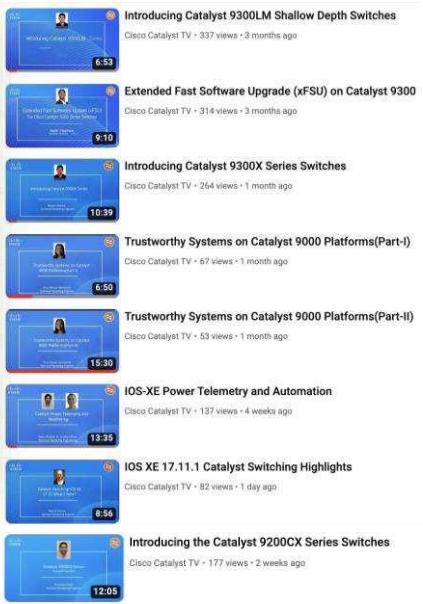

2100+ Subscribers

https://www.youtube.com/c/CiscoCatalystTV

Complete Your Session Evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to **win 1 of 5 full conference passes** to Cisco Live 2025.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn **exclusive prizes!**

Complete your surveys in the **Cisco Live mobile app.**

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 www.CiscoLive.com/on-demand

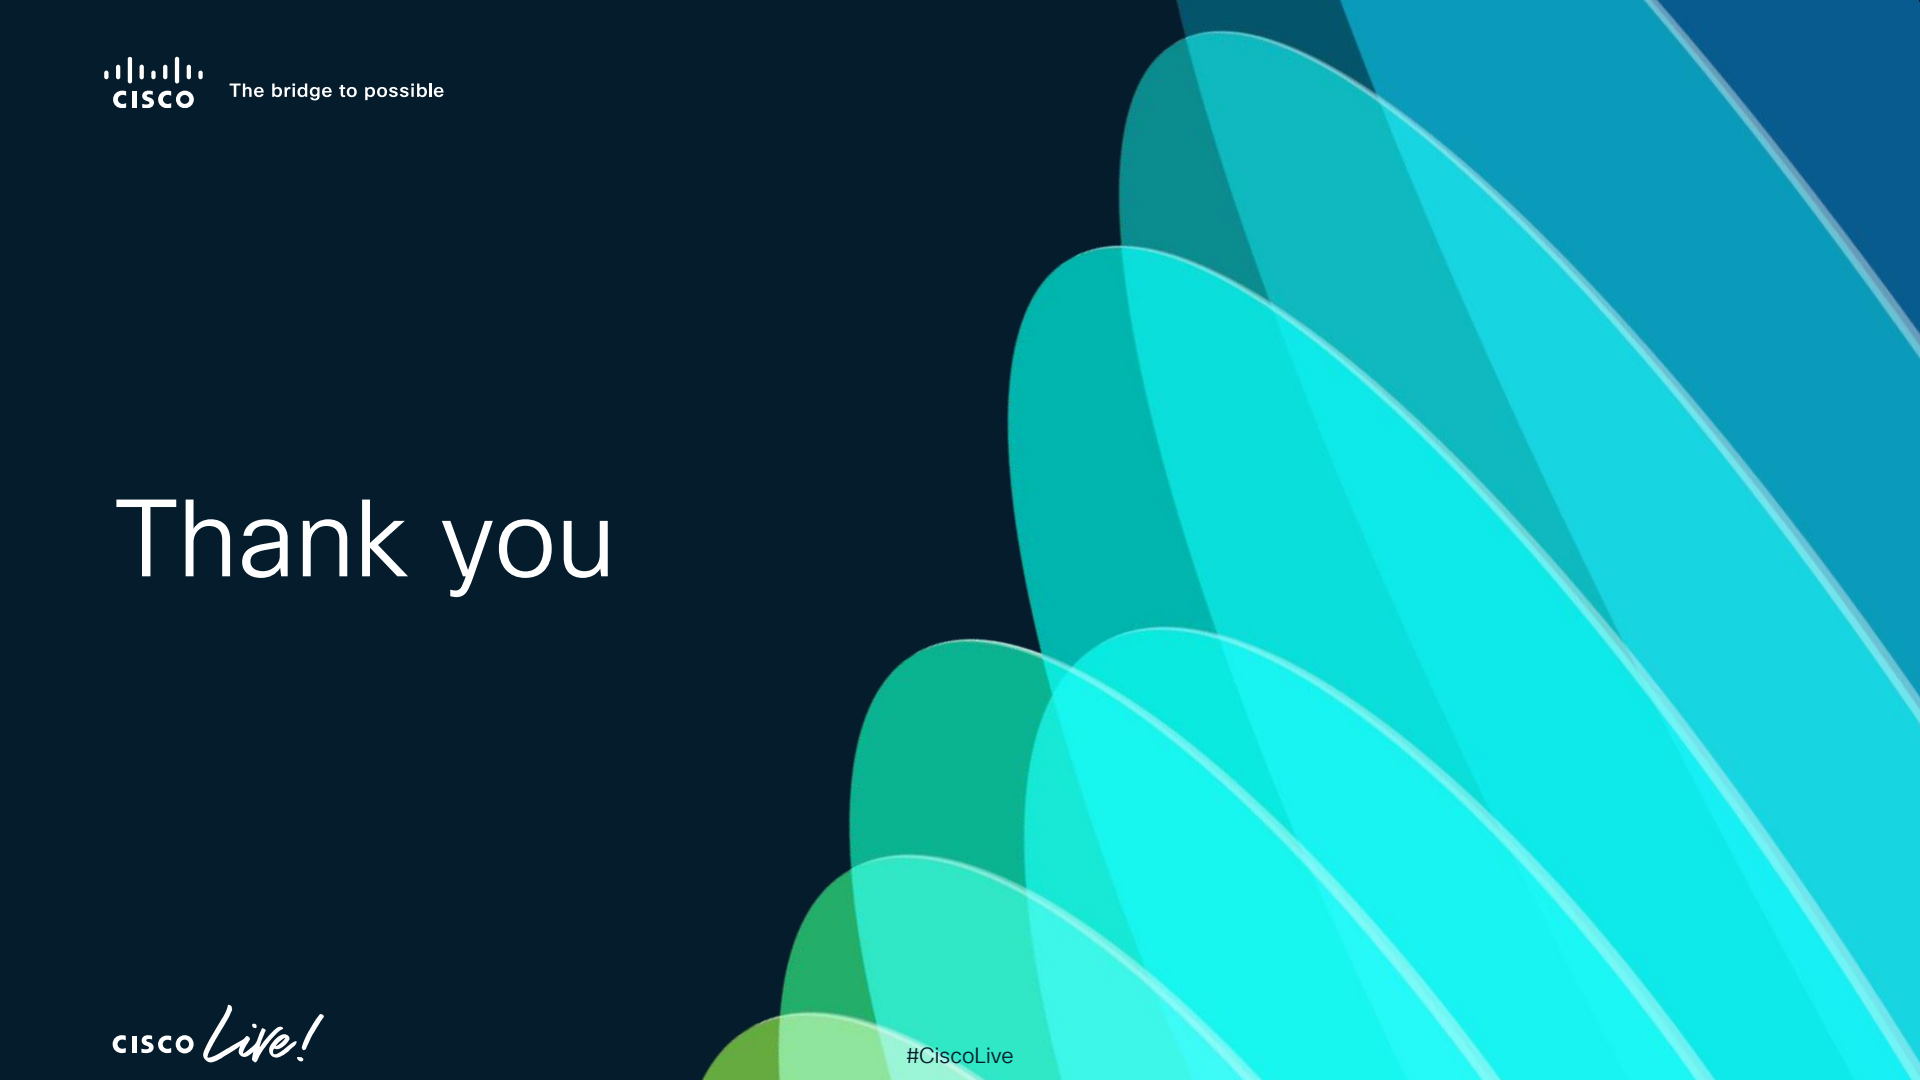
The bridge to possible
Thank you
cisco Live!
#CiscoLive