Overlay Design Options for Campus Networks

Raj Kumar Goli
Technical Marketing Engineer
BRKENS-2501

Webex App

Questions?

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

How

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

Webex spaces will be moderated by the speaker until February 28, 2025.

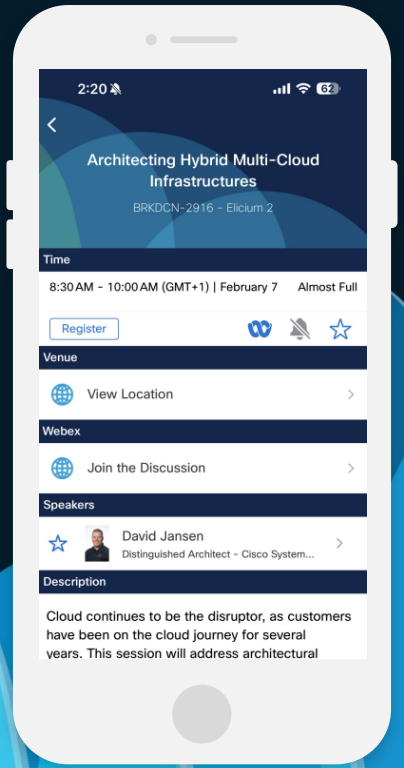

Introduction

Raj Kumar Goli
Technical Marketing Engineer

- I work as a **Technical Marketing Engineer** focusing on Catalyst 9000 switching platforms and Enterprise Campus Architectures. I primarily focus on Fabric Solutions like BGP-EVPN & MPLS, Security Solutions like IPsec, WAN MACsec, Cloud Security and Time Sensitive Solutions like Precision Time Protocol and Audio Vide Bridging.
- Started TME role for Catalyst 9000 in 2018
- Wichita State University Alumni
- CCIE (Data Center & Service Provider)

Sessions are available Online @ CiscoLive.com

Cisco Campus Architecture

Cisco Live Amsterdam 2025 - Session Map

You Are Here

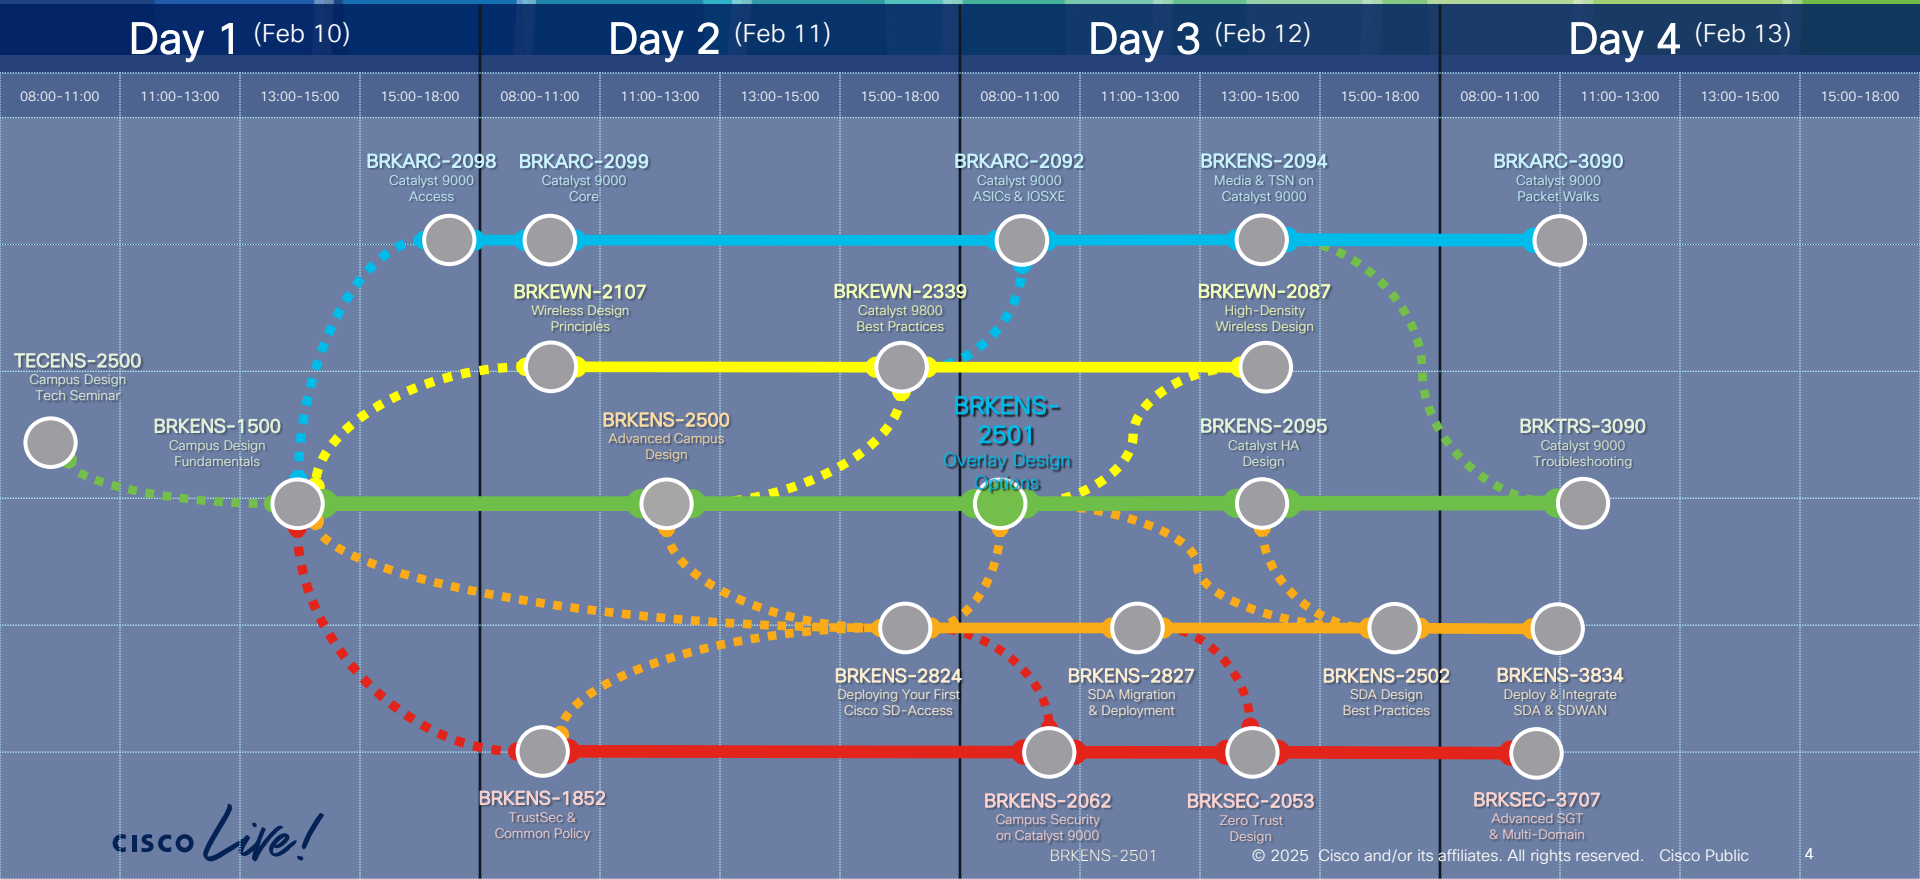

Agenda

- Campus Architecture
- What is Underlay
- What is Overlay
- Overlay Choices
- SD-Access with LISP
- Vxlan with BGP EVPN

What to Expect

- The session is designed to provide high level understanding of the overlay choices available on catalyst 9000 switches
- This session is not a deep dive session into any specific fabric.
- For deeper dive into each of the fabric offerings, below cisco live sessions are recommended - BRKENS-2810, BRKENS-3555, BRKENS-2502, BRKENS2-2092

Campus is focused on User Access

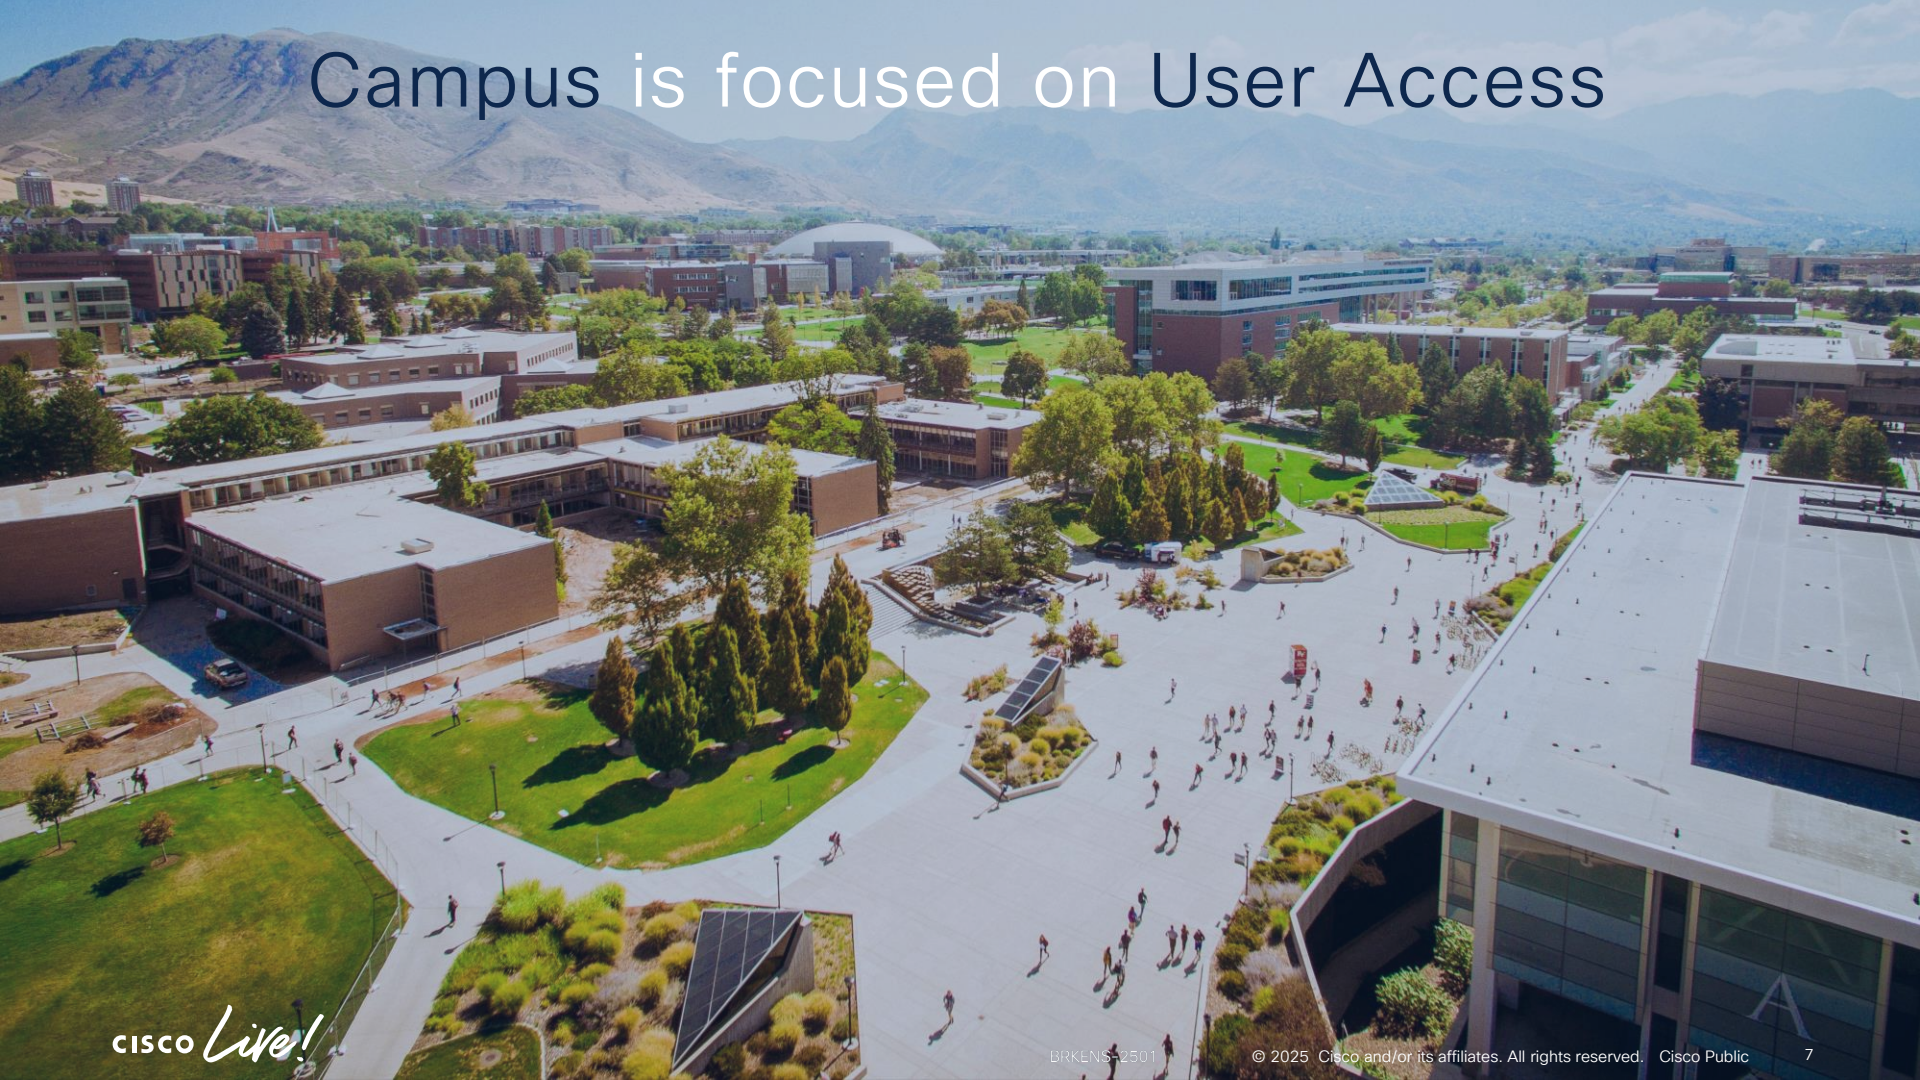

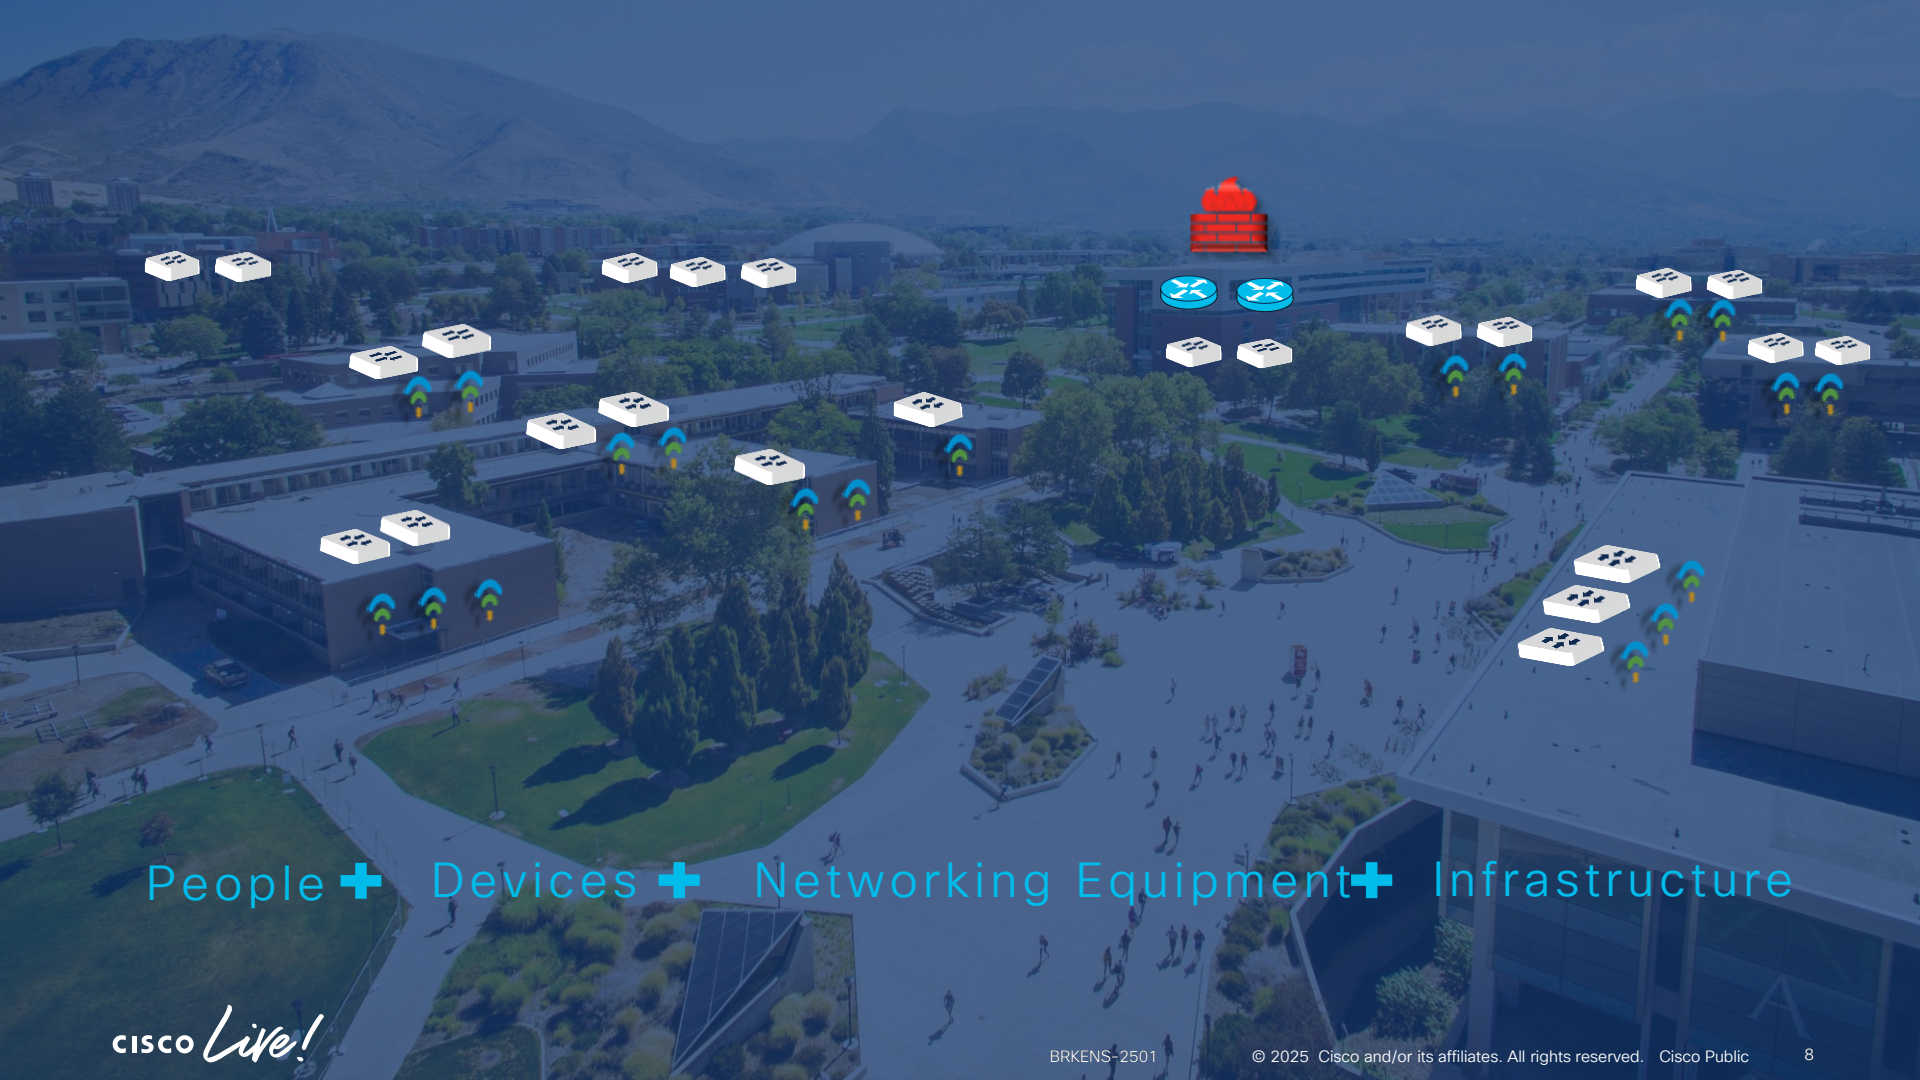

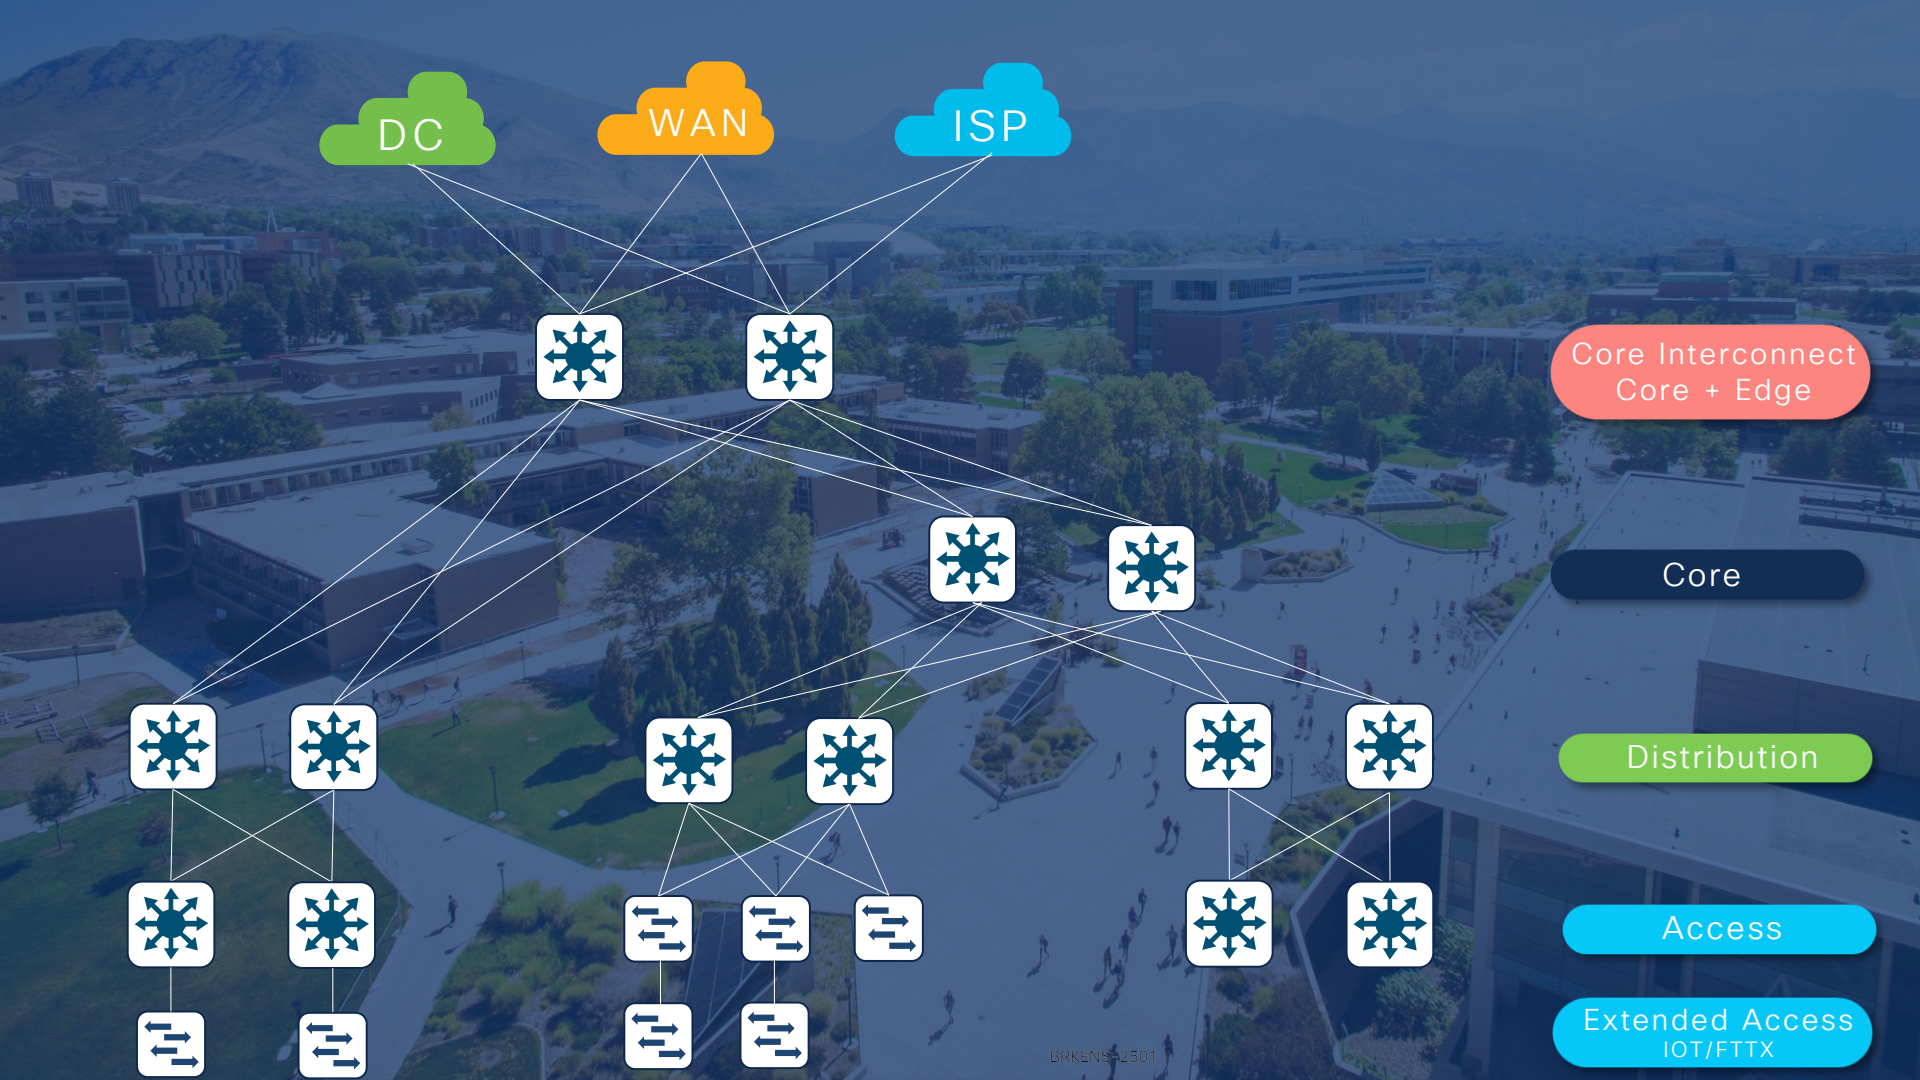
DC
WAN
ISP
Core Interconnect
Core + Edge
Core
Distribution
Access
Extended Access
IOT/FTTX

Campus Multi-Layer Model

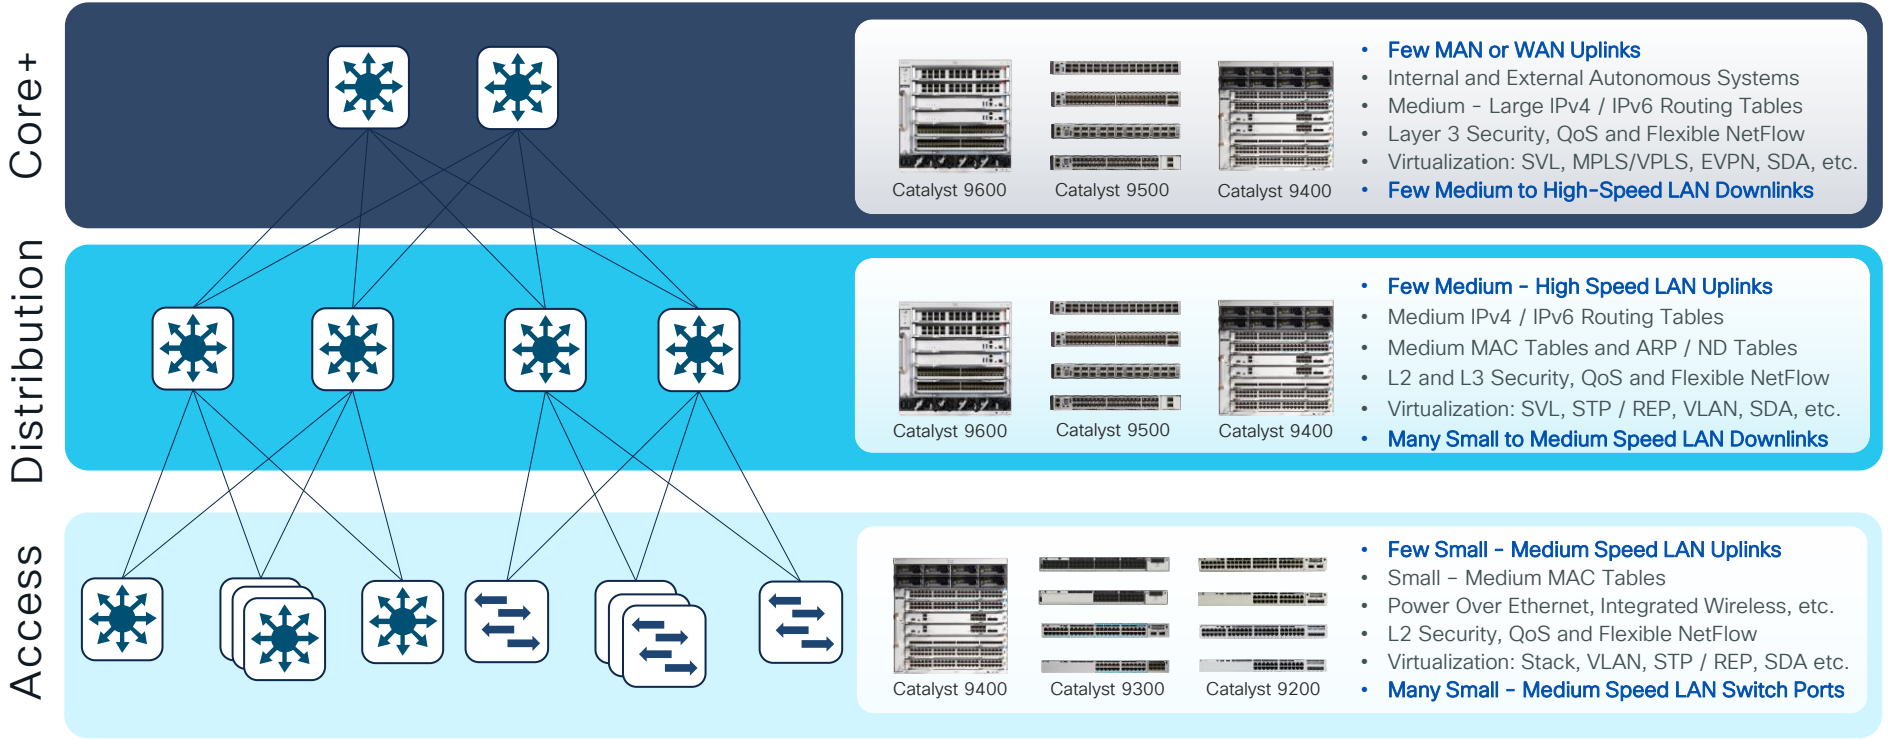

Logical Layers

Each layer has a specific set Each layer serves a specific set of functions

If you **collapse layers** – your multi-layer device needs to support **all "logical" functions**

Underlay Networks

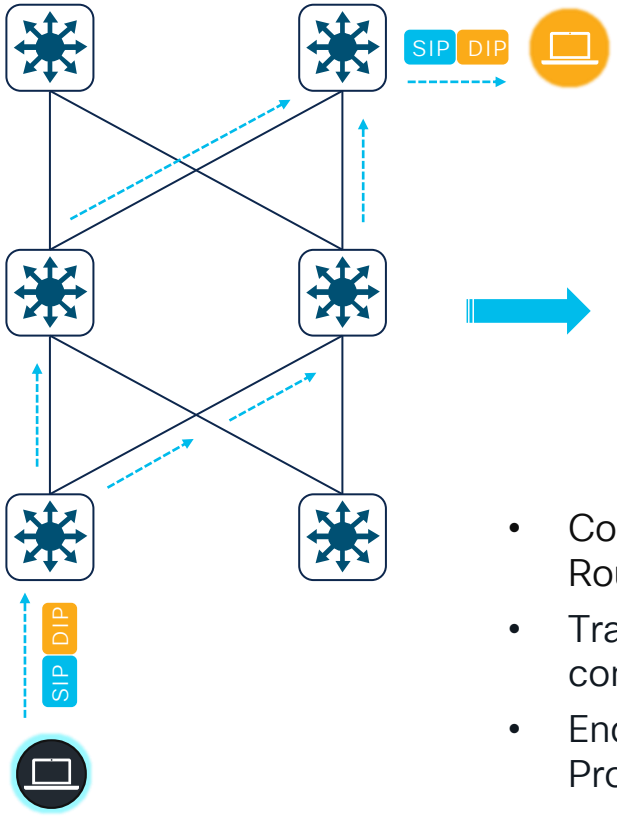

- Consists of the actual hardware (Routers/Switches), Routing, and physical Links (Fiber, copper, etc..)
- Traditional Layer 3 Routed Network and provides connectivity between network devices.
- End hosts communication relies on IGP or EGP Routing Protocols

Overlay Networks

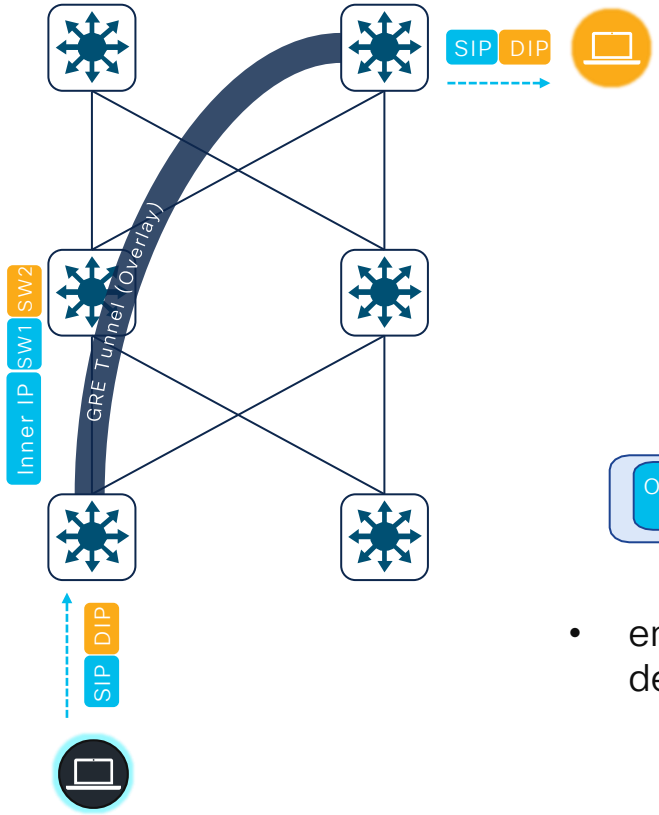

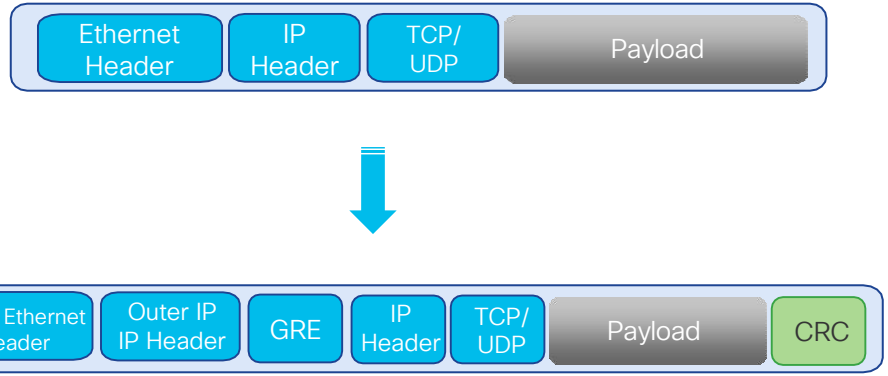

- encapsulate and transport data traffic between two devices over an existing underlay network.

Multiple Encapsulation choices

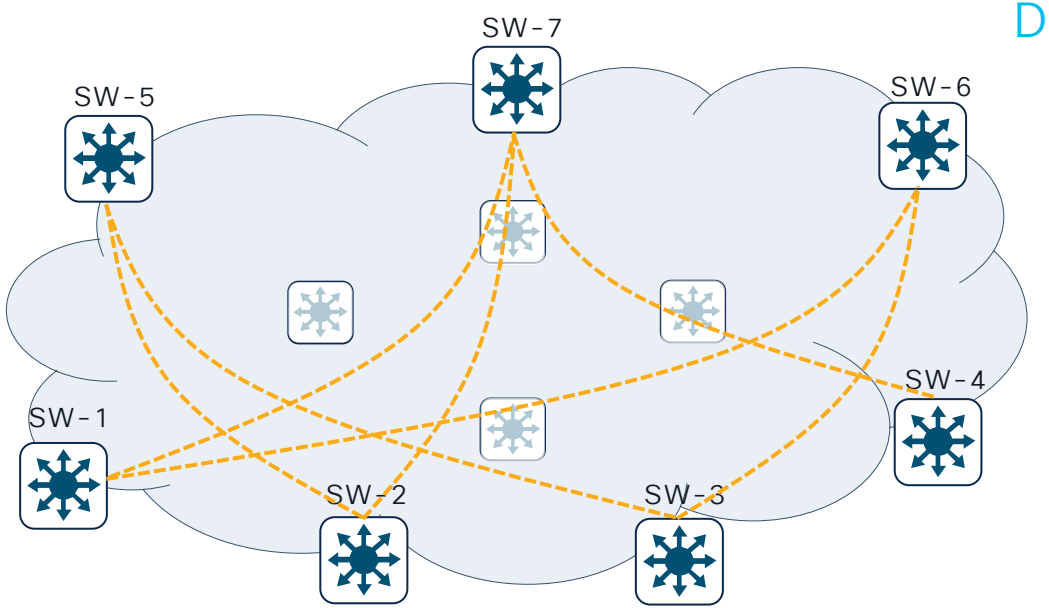

Different Encapsulation Types…

GRE (Generic Routing Encapsulation)

IPSec (Internet Protocol Security)

MPLS (Multiprotocol Label Switching)

Vxlan (Virtual Extensible LAN)

OTV (Overlay Transport Virtualization)

NVGRE (Network Virtualization using GRE)

L2TP (Layer 2 Tunneling Protocol)

Overlay Networks

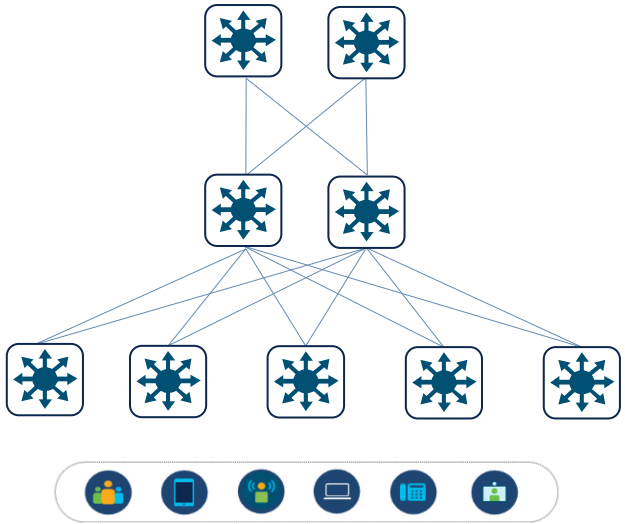

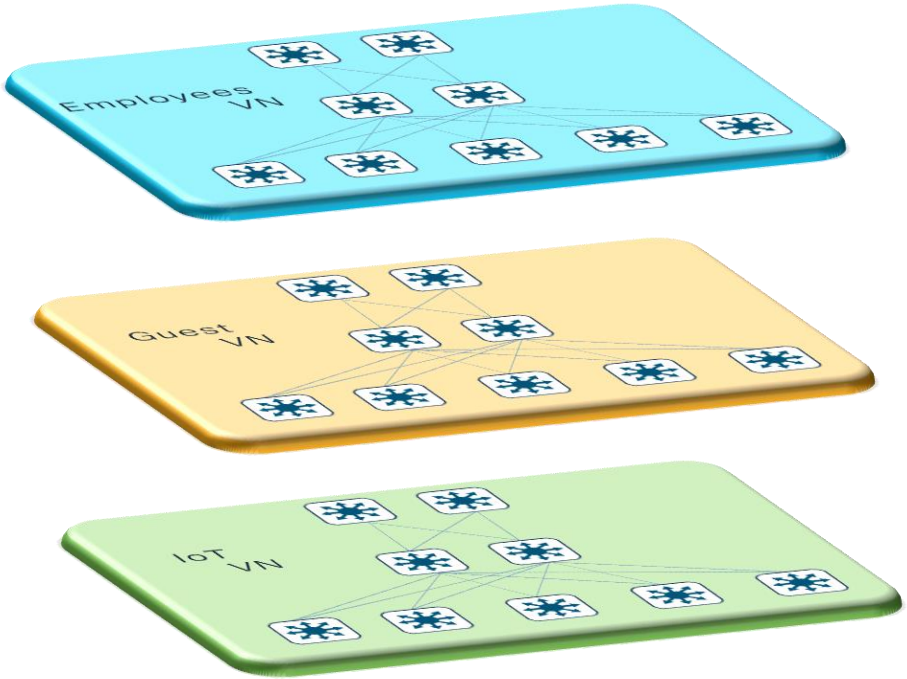

Virtualized network layer decoupled from physical underlay | Unique security policies per logical domain | Traffic isolation per application, group, service etc…

Virtual Private Network (VPN)

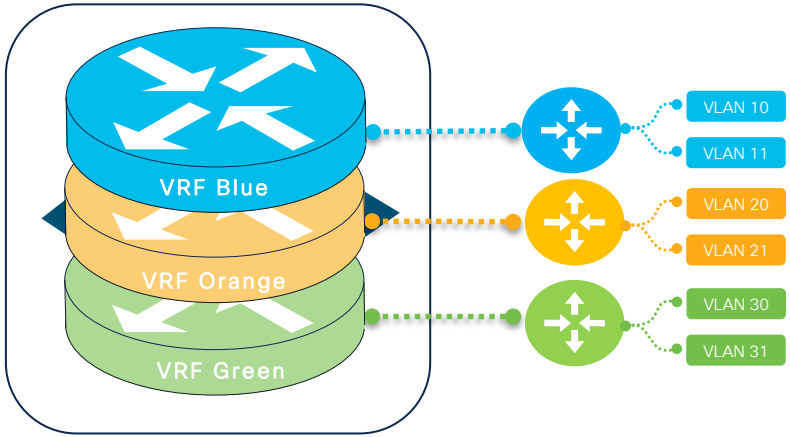

Device Virtualization

- Physically One device
- Logically many devices
  - Control Plane
  - Data Plane
- Creates independent & separate IPv4 & IPv6 address spaces
- Data traffic is not routed across VRF's with default configuration

Traditional Multilayer Design

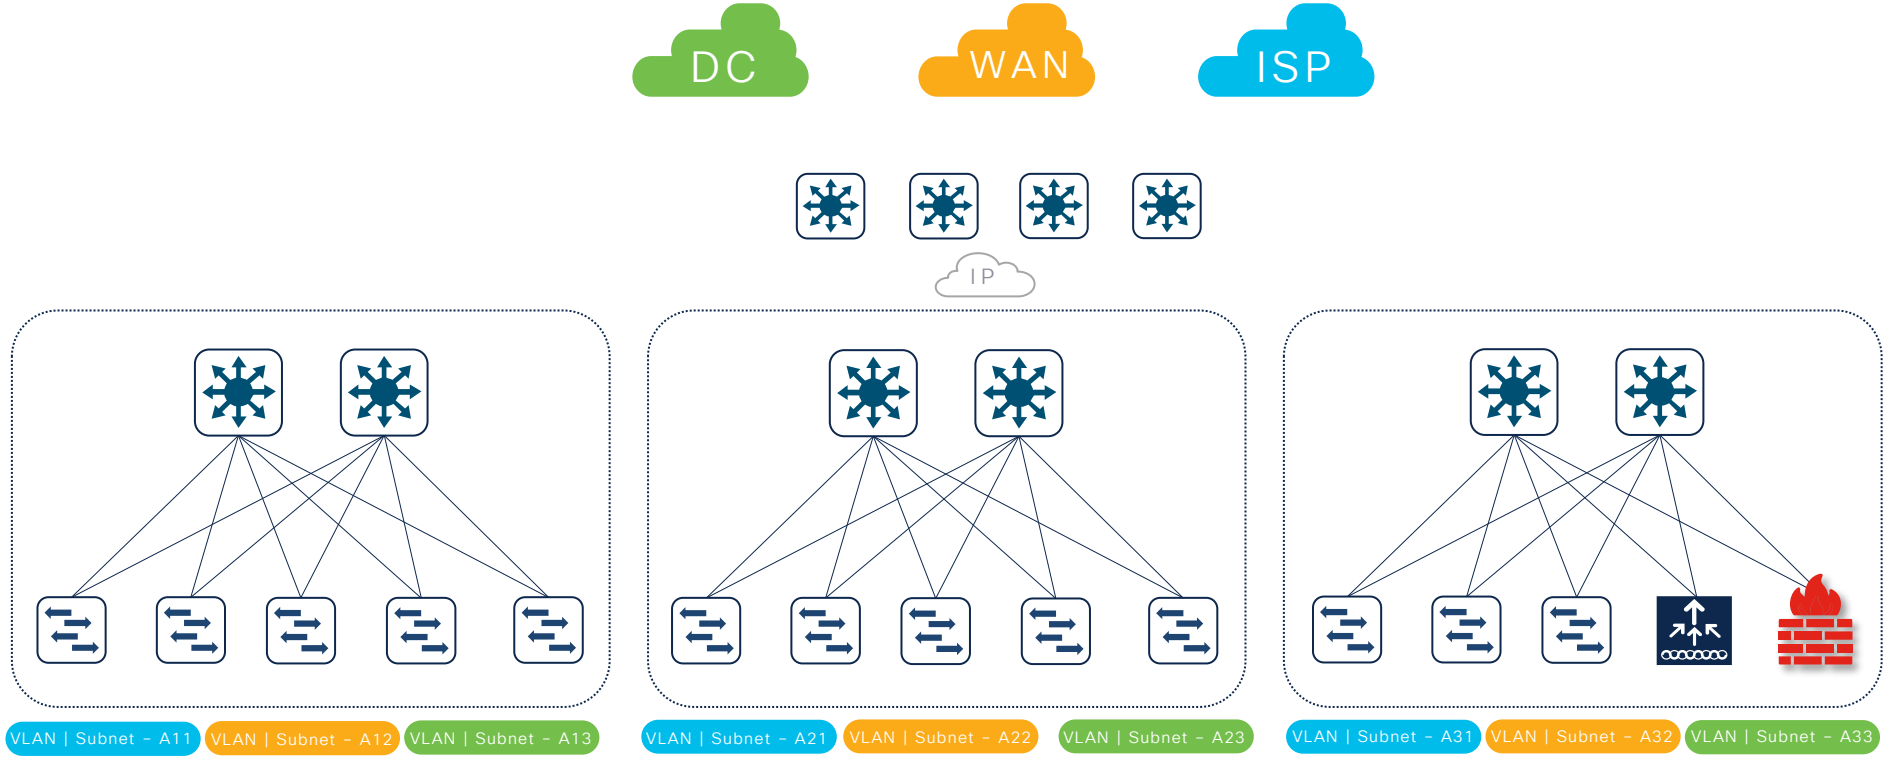

Traditional Multilayer Design

- Structured physical networking supporting distributed bridging and central routing function
- Hierarchical system role and networking function at each layer
- Deterministic Layer 2 and Layer 3 forwarding paths provides network scalability, resiliency in best practice network environment
- Limited traditional protocols flexibility and complex operational overhead

Traditional Layer 3 Access Design

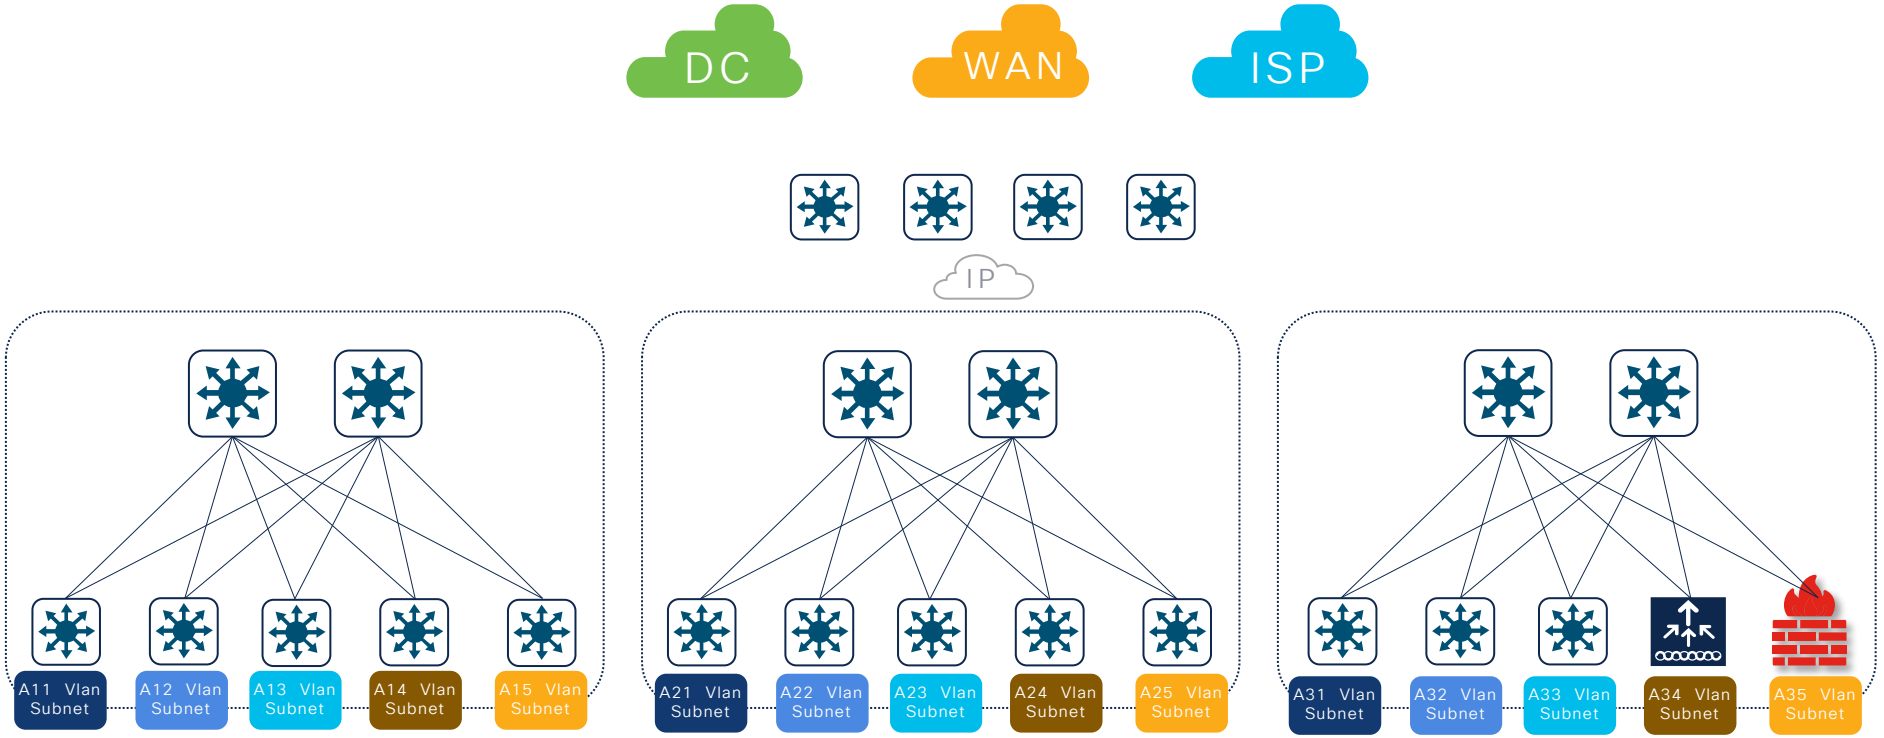

Traditional L3 Access Design

- Consistent structured physical networking supporting fully distributed bridging and routing function
- Common principal for hierarchical system role and networking function at each layer
- Simplified and deterministic non-blocking Layer 3 forwarding paths provides network scalability, resiliency in best practice network environment
- Limited mobility and traditional Layer 2 protocols flexibility, complex routing operational overhead and increase cost

Data Plane and Control Plane

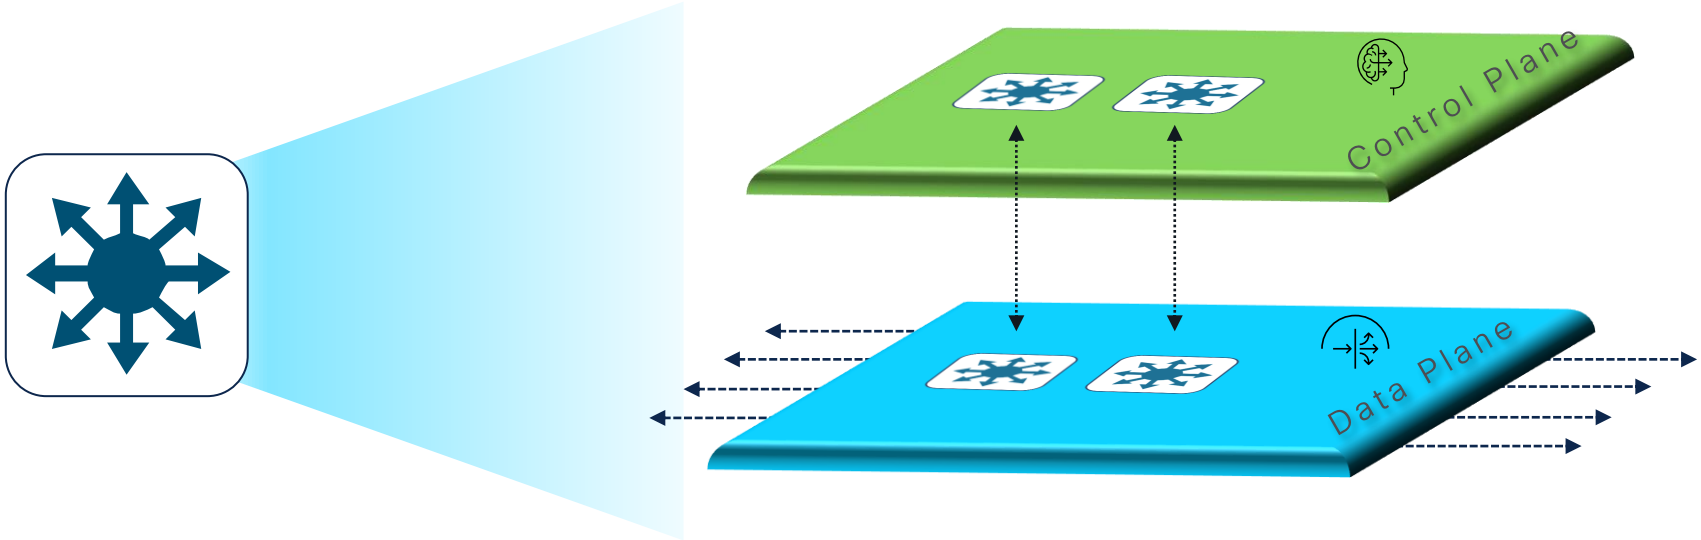

Cisco Enterprise Fabric Alternatives

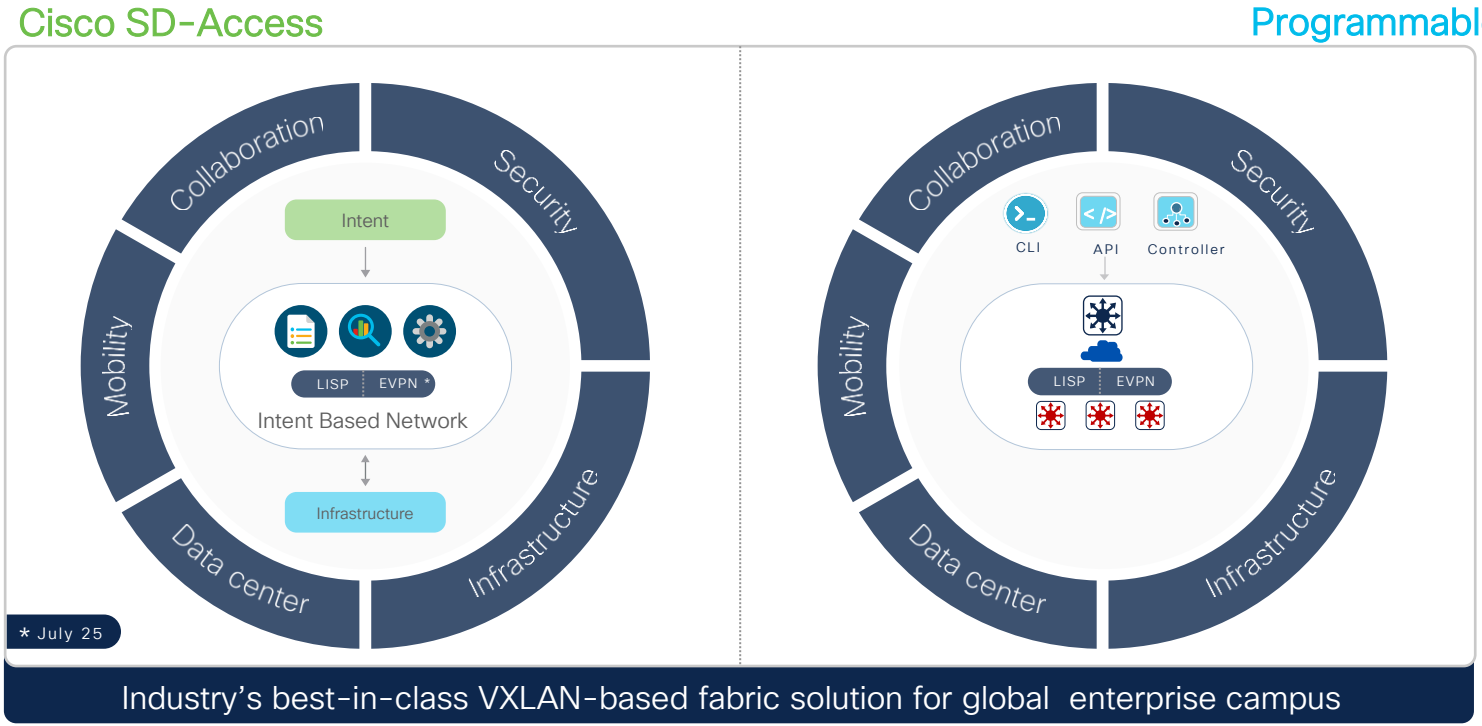

SDA-LISP – Industry-standard, light-weight purpose-built Wired + Wireless fabric control-plane for large scale distributed mobility.

SDA-EVPN – Multi-vendor, industry-standard unified control-plane for end-to-end Wired network fabric beyond enterprise campus boundary.

Vxlan Encapsulation

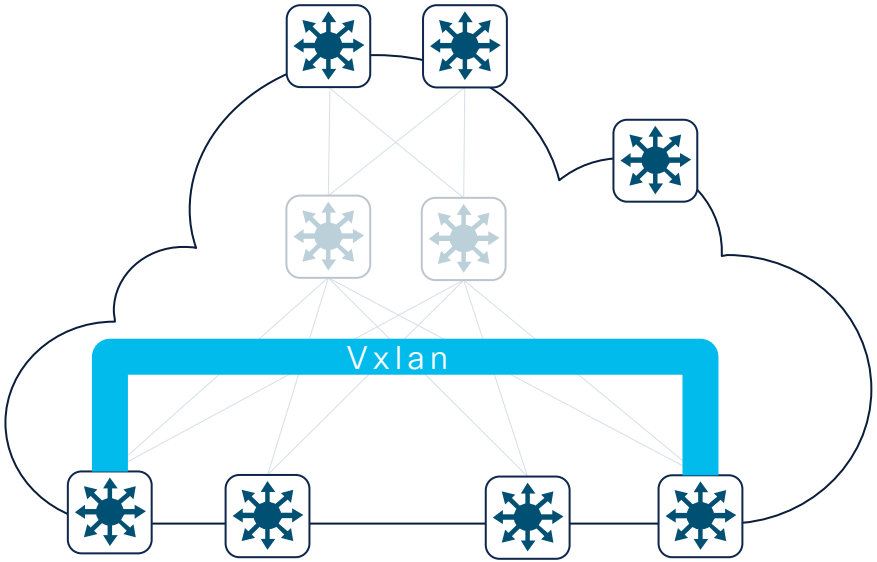

VXLAN

- Standards based Encapsulation
  - RFC 7348
  - Uses UDP-Encapsulation
- Transport Independent
  - Layer-3 Transport (Underlay)
- Flexible Namespace
  - 24-bit field (VNID) provides ~16M unique identifier
  - Allows Segmentations

Vxlan Packet Structure

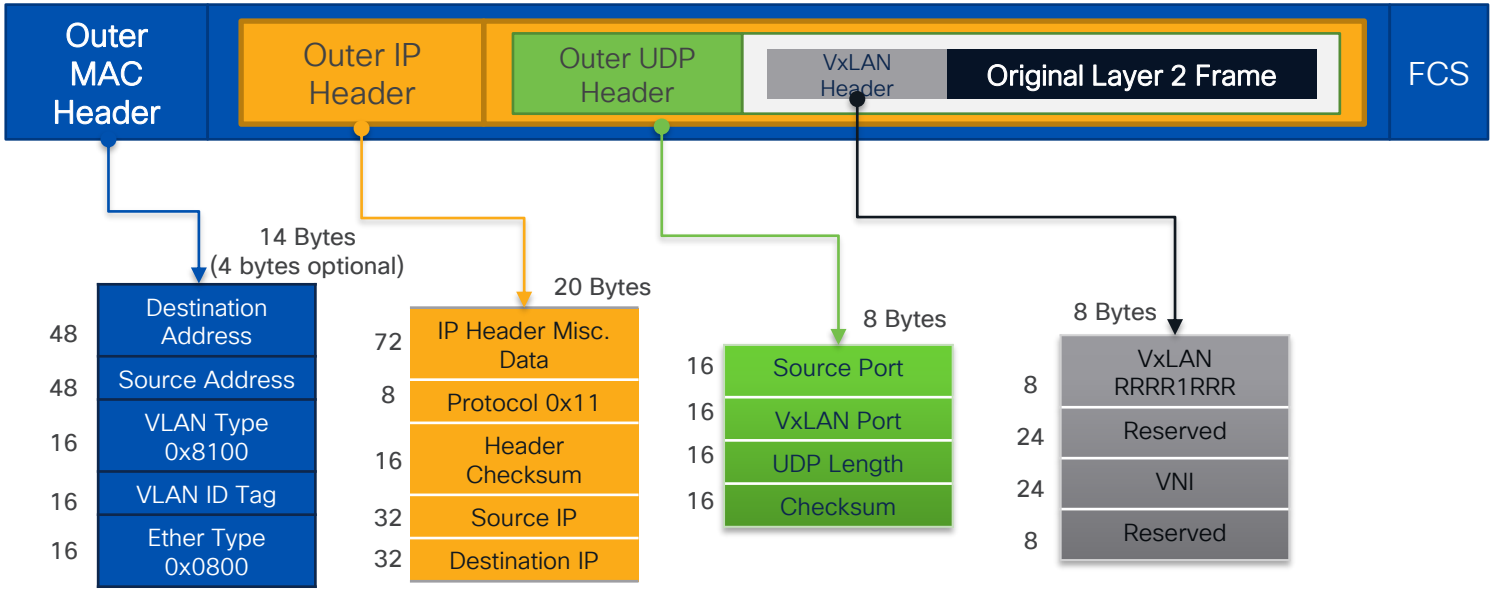

- VXLAN adds 50 Bytes (or 54 Bytes) to the Original Ethernet Frame
- Avoid Fragmentation by adjusting the IP Networks MTU
- Using a MTU of 9216* Bytes accommodates VXLAN Overhead plus other application MTU

Fabric Segmentation

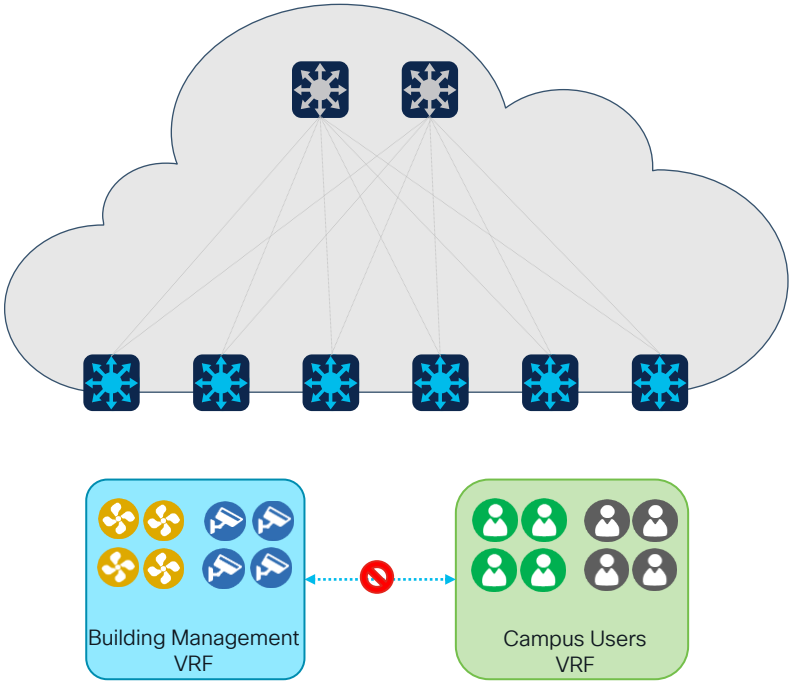

Macro Segmentation: No communication between VRF's

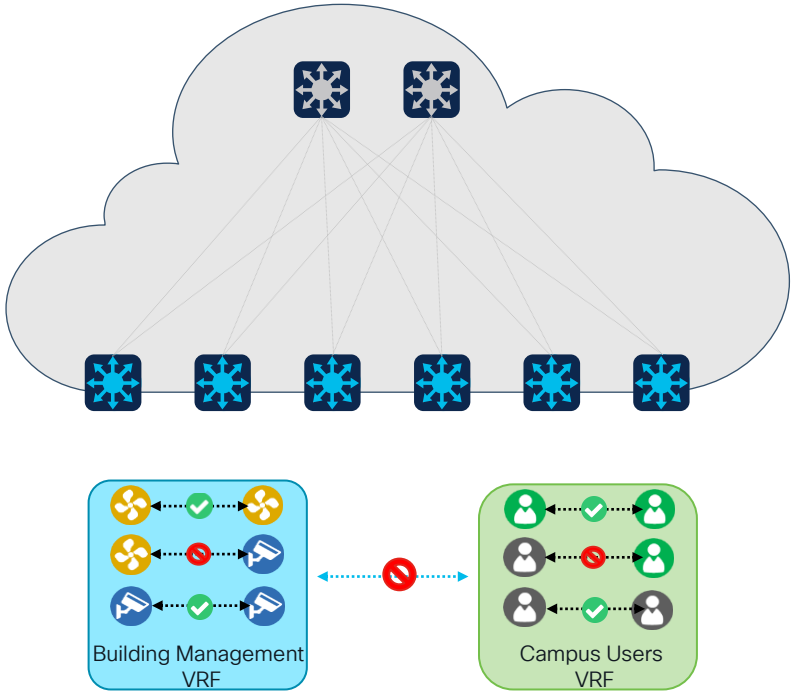

Micro Segmentation: Second level Segmentation between groups within a VRF

Micro Segmentation

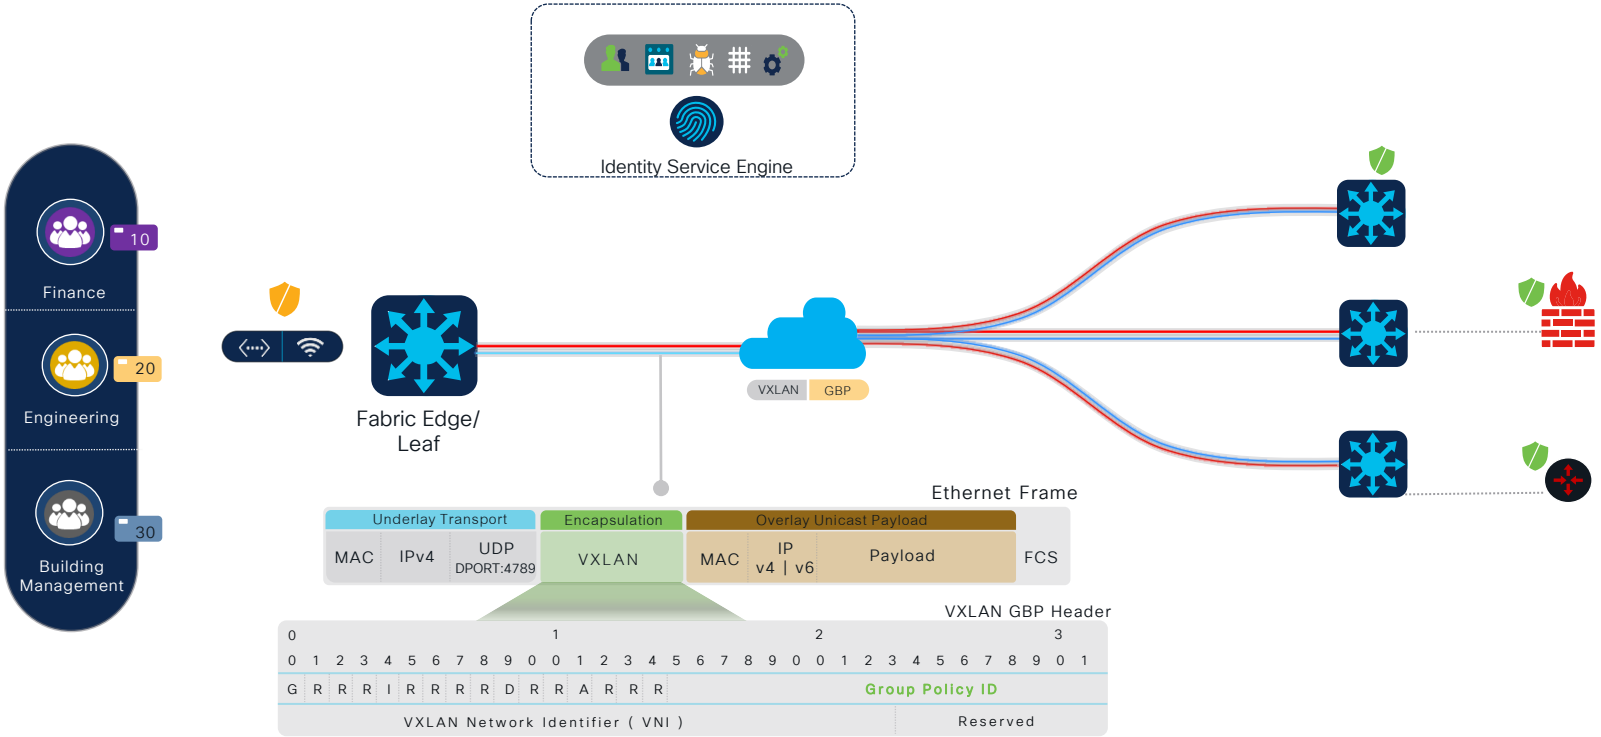

Classification

Enforcement

SD-Access with LISP

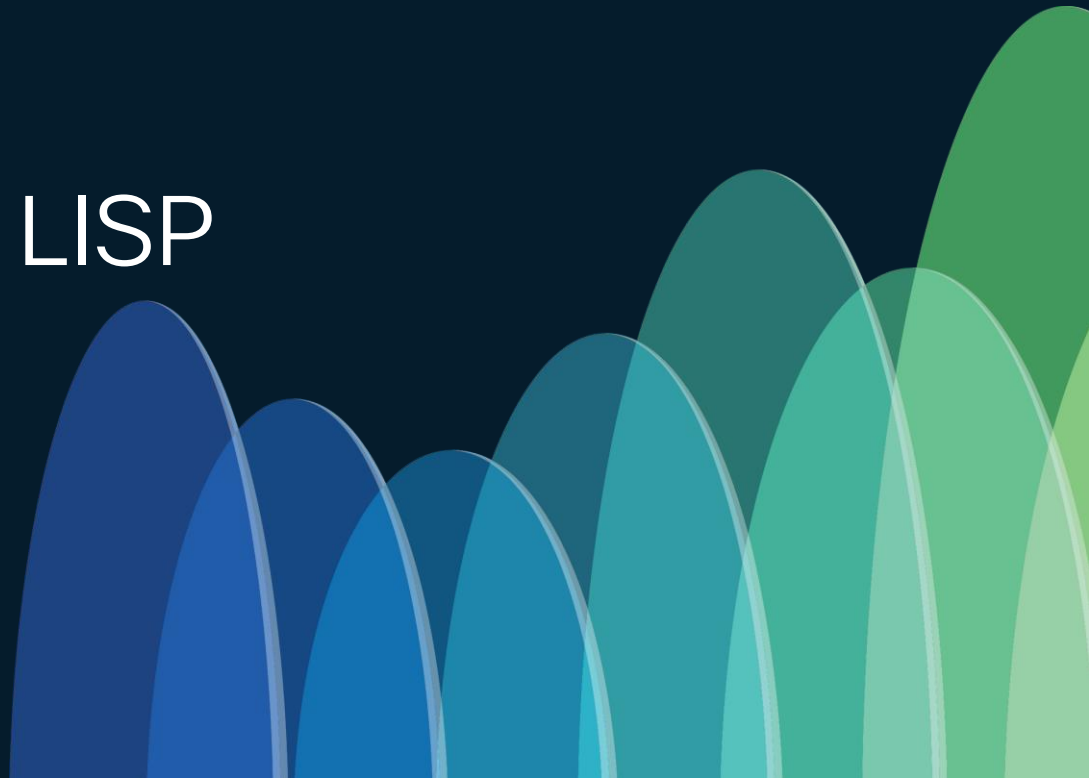

cisco Live!

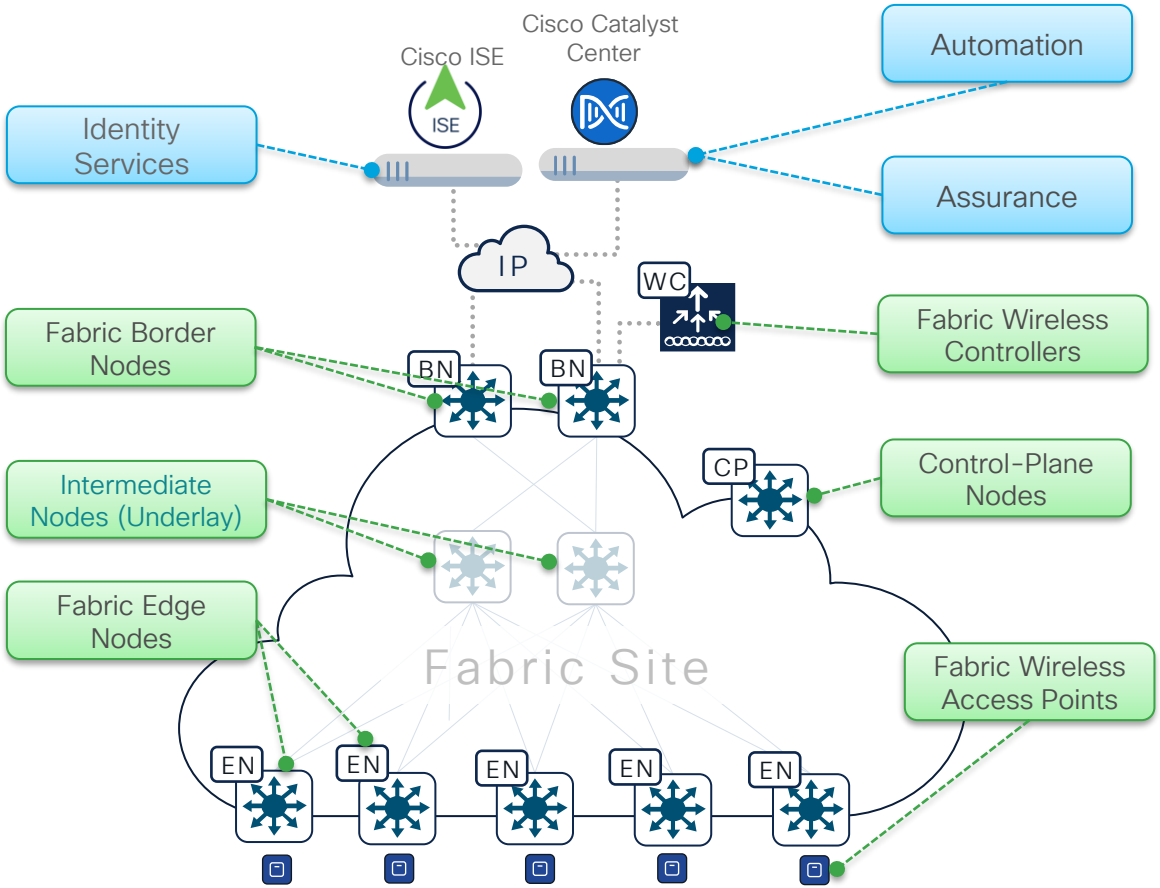

Compatibility Matrix:
https://www.cisco.com/c/dam/en/us/td/docs/Website/enterprise/catalyst_center_compatibility_matrix/index-sda.html

- **Network Automation** – Simple GUI and APIs for intent-based Automation of wired and wireless fabric devices
- **Network Assurance** – Data Collectors analyze Endpoint to Application flows and monitor fabric device status
- **Identity Services** – NAC & ID Services (e.g. ISE) for dynamic Endpoint to Group mapping and Policy definition
- **Control-Plane Nodes** – Map System that manages Endpoint to Device relationships
- **Fabric Border Nodes** – A fabric device (e.g. Core) that connects External L3 network(s) to the SD-Access fabric
- **Fabric Edge Nodes** – A fabric device (e.g. Access or Distribution) that connects Wired Endpoints to the SD-Access fabric
- **Intermediate Nodes** – A Layer 2 or Layer 3 Underlay network system providing basic transport and forwarding plane.
- **Fabric Wireless Controller** – A fabric device (WLC) that connects Fabric APs and Wireless Endpoints to the SD-Access fabric

Locator/ID Separation Protocol (LISP)

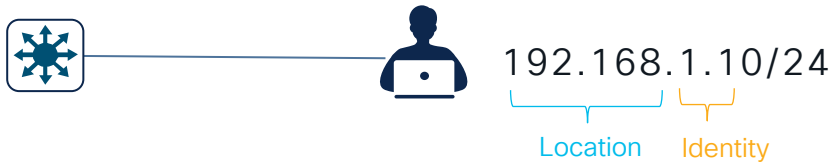

Resolving Locators for Identities

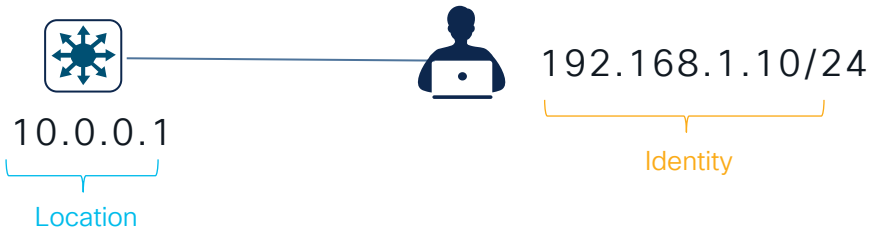

Locator/ID Separation Protocol (LISP)

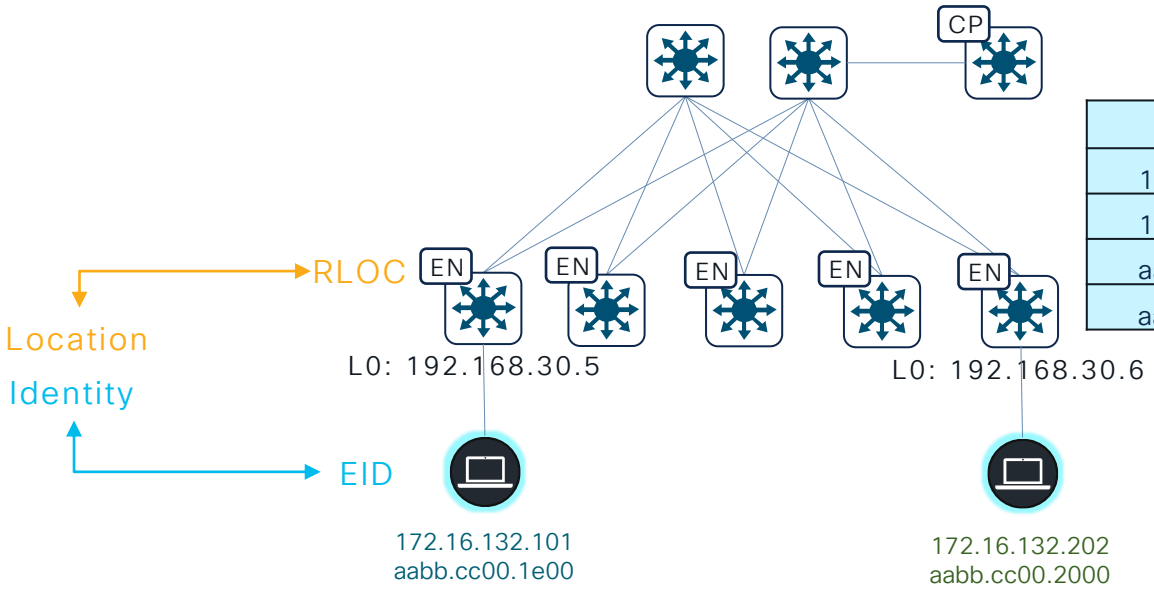

EID-to-RLOC Mapping

EID	RLOC
172.16.132.101	192.168.30.5
172.16.132.202	192.168.30.6
aabb.cc00.1e00	192.168.30.5
aabb.cc00.2000	192.168.30.6

EID – Endpoint ID
RCLO – Routing Locator

Locator/ID Separation Protocol (LISP)

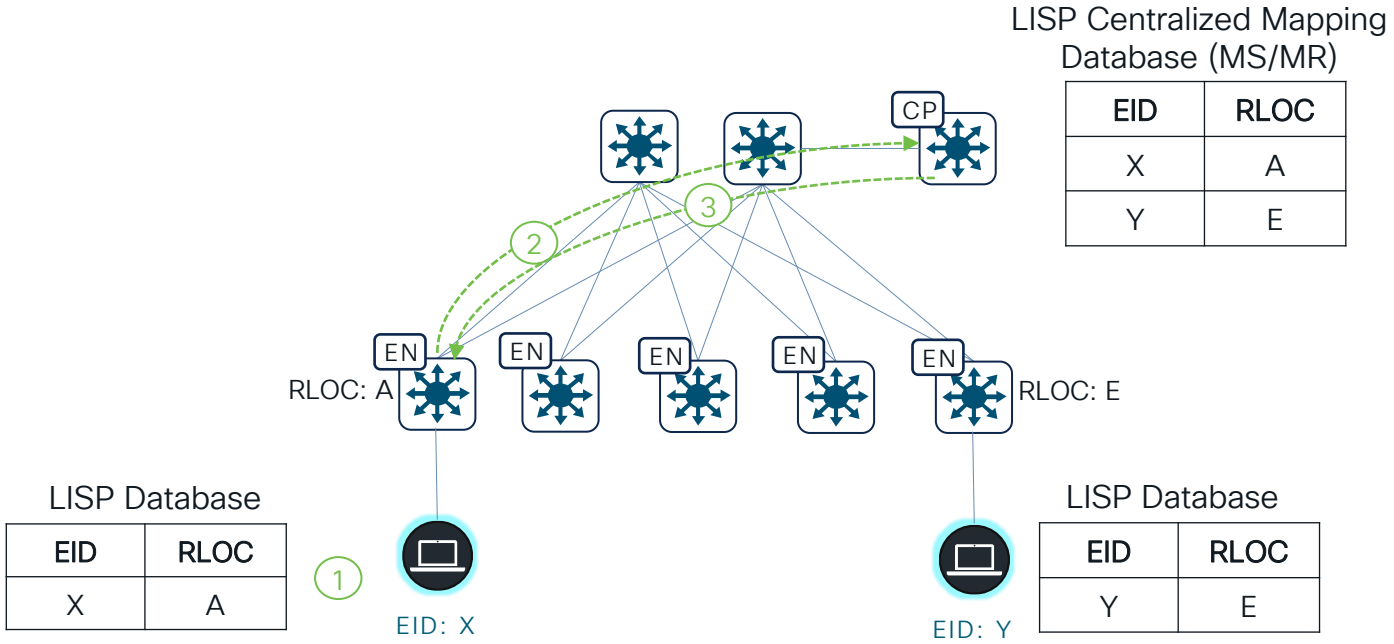

Locator/ID Separation Protocol (LISP)

Seamless Roaming

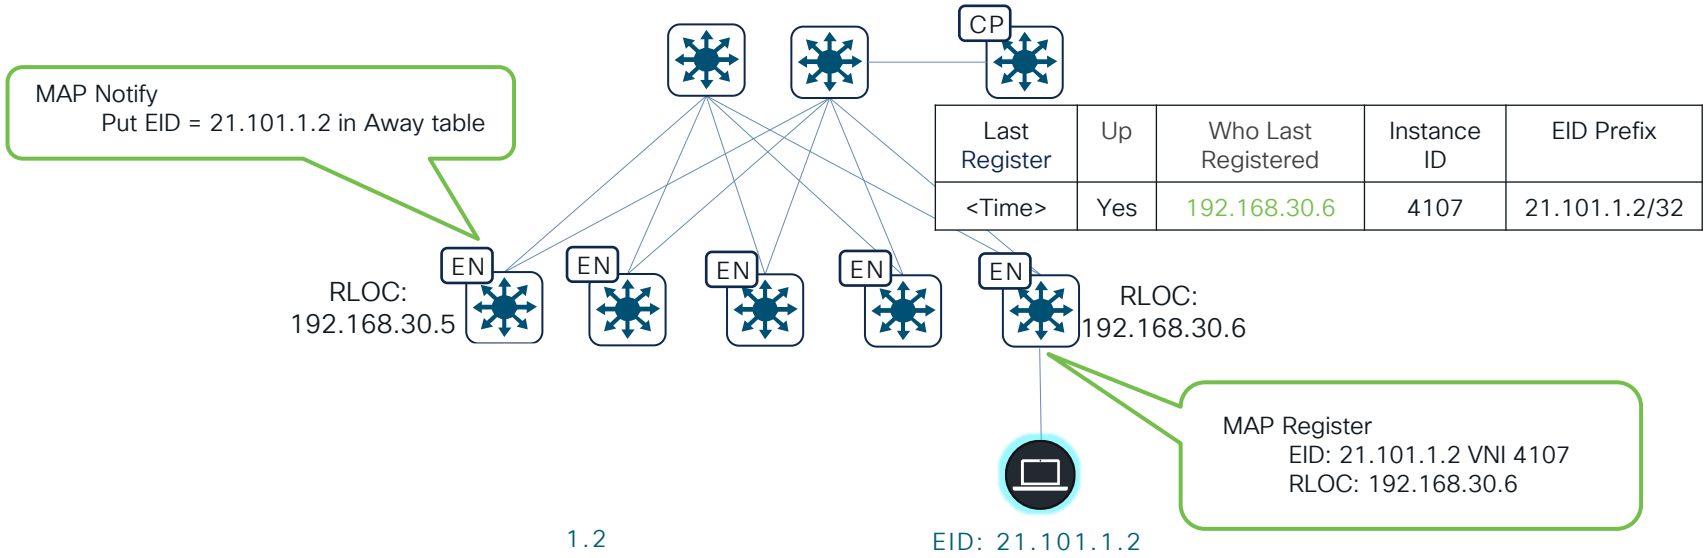

Cisco SD-Access Fabric Site Design Options

Fabric Site Design Options

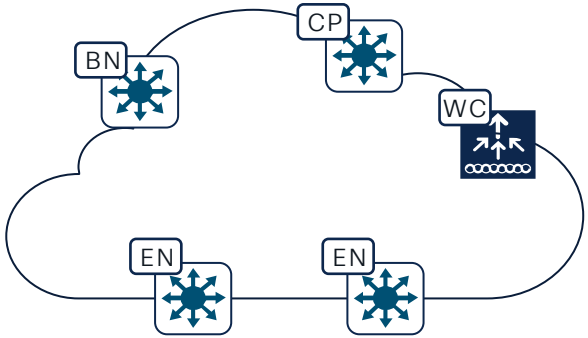

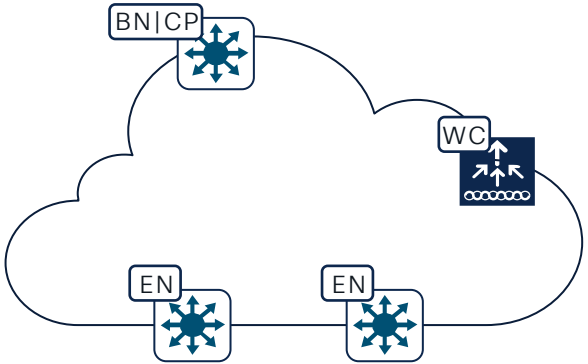

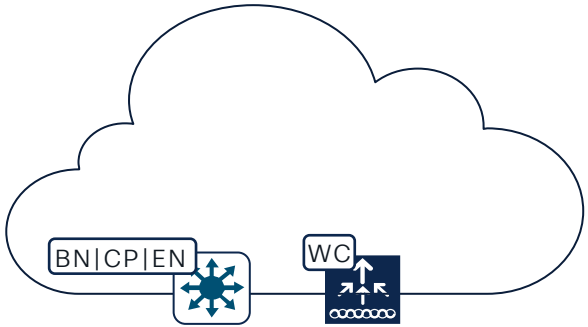

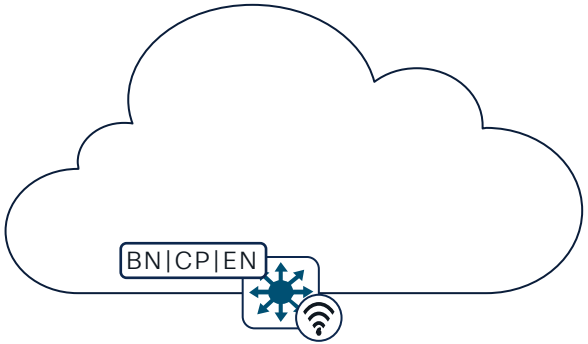

SD-Access with LISP – Overlay Types

Distributed Anycast Gateway

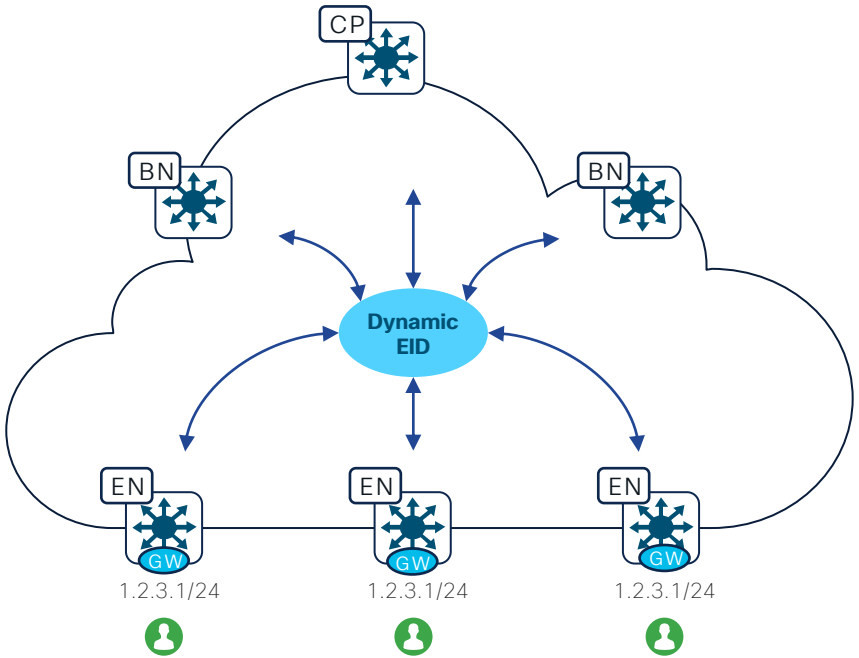

Layer 2 Overlay

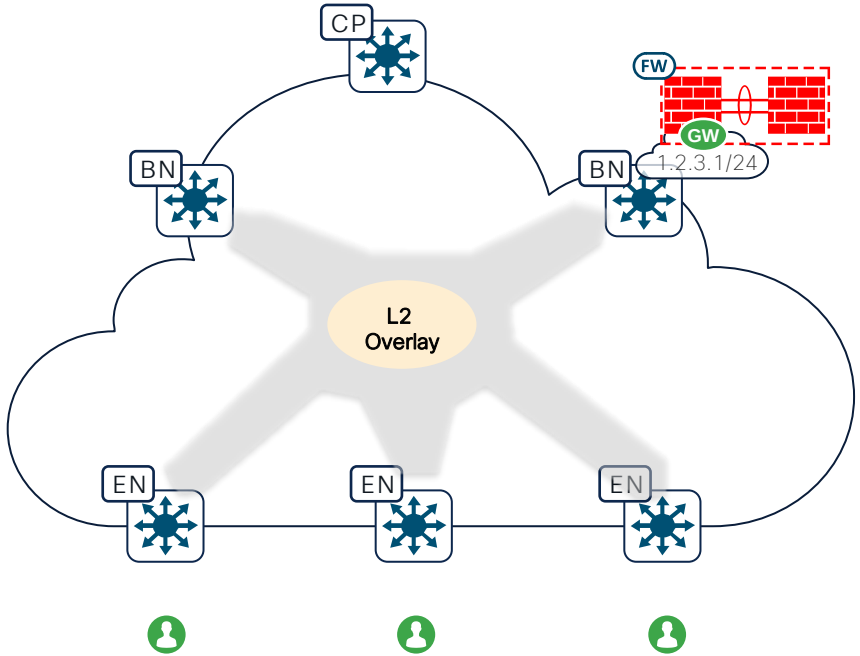

Cisco SD-Access Architecture
Single-Site Architecture

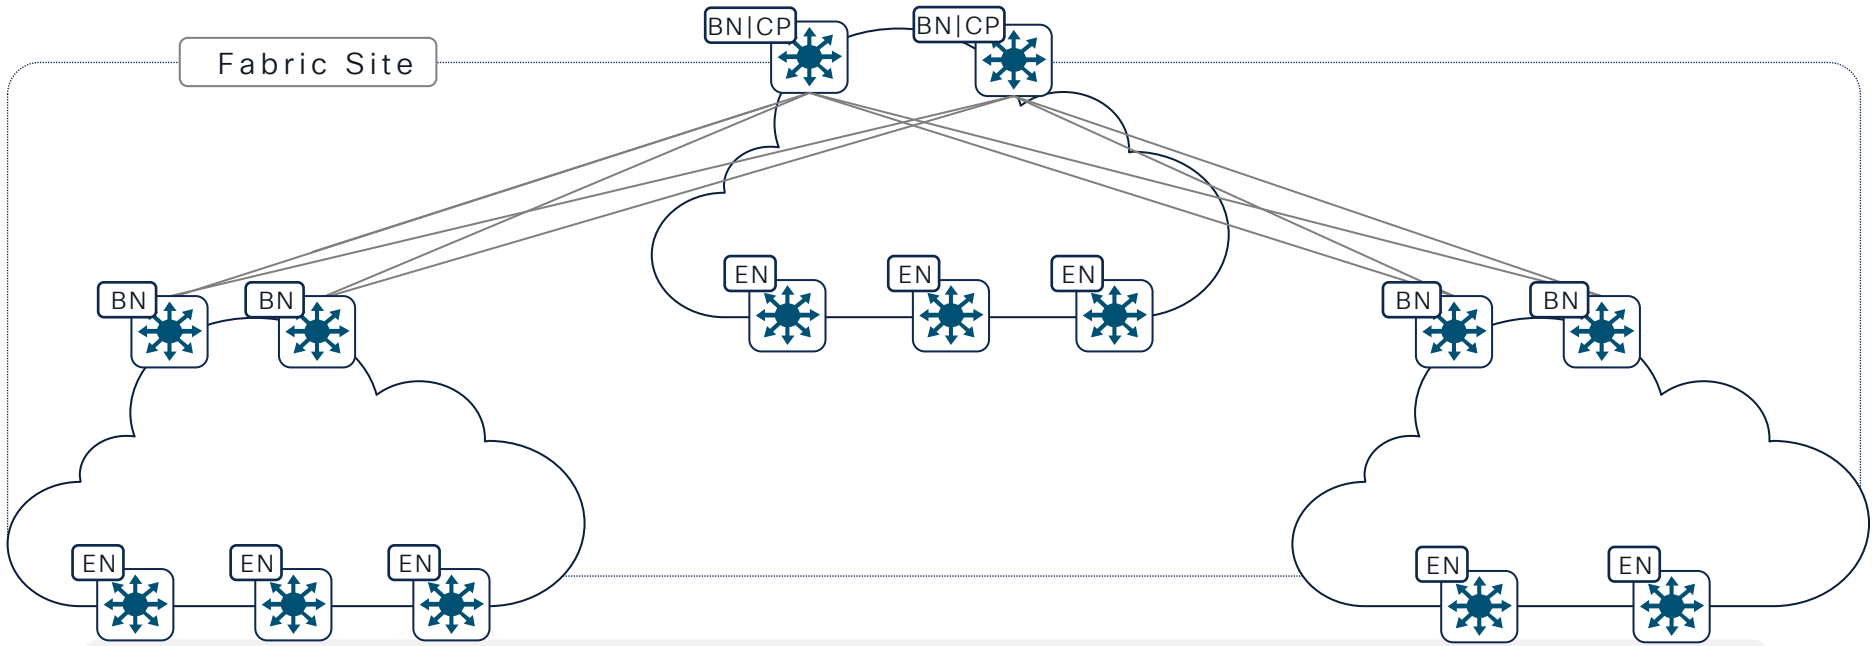

One Subnet available across all buildings/Sites

Scale Limitations – IP Pools supported per site or Border/Control plane Scale

Cisco SD-Access Architecture

Multisite Architecture with IP Transit

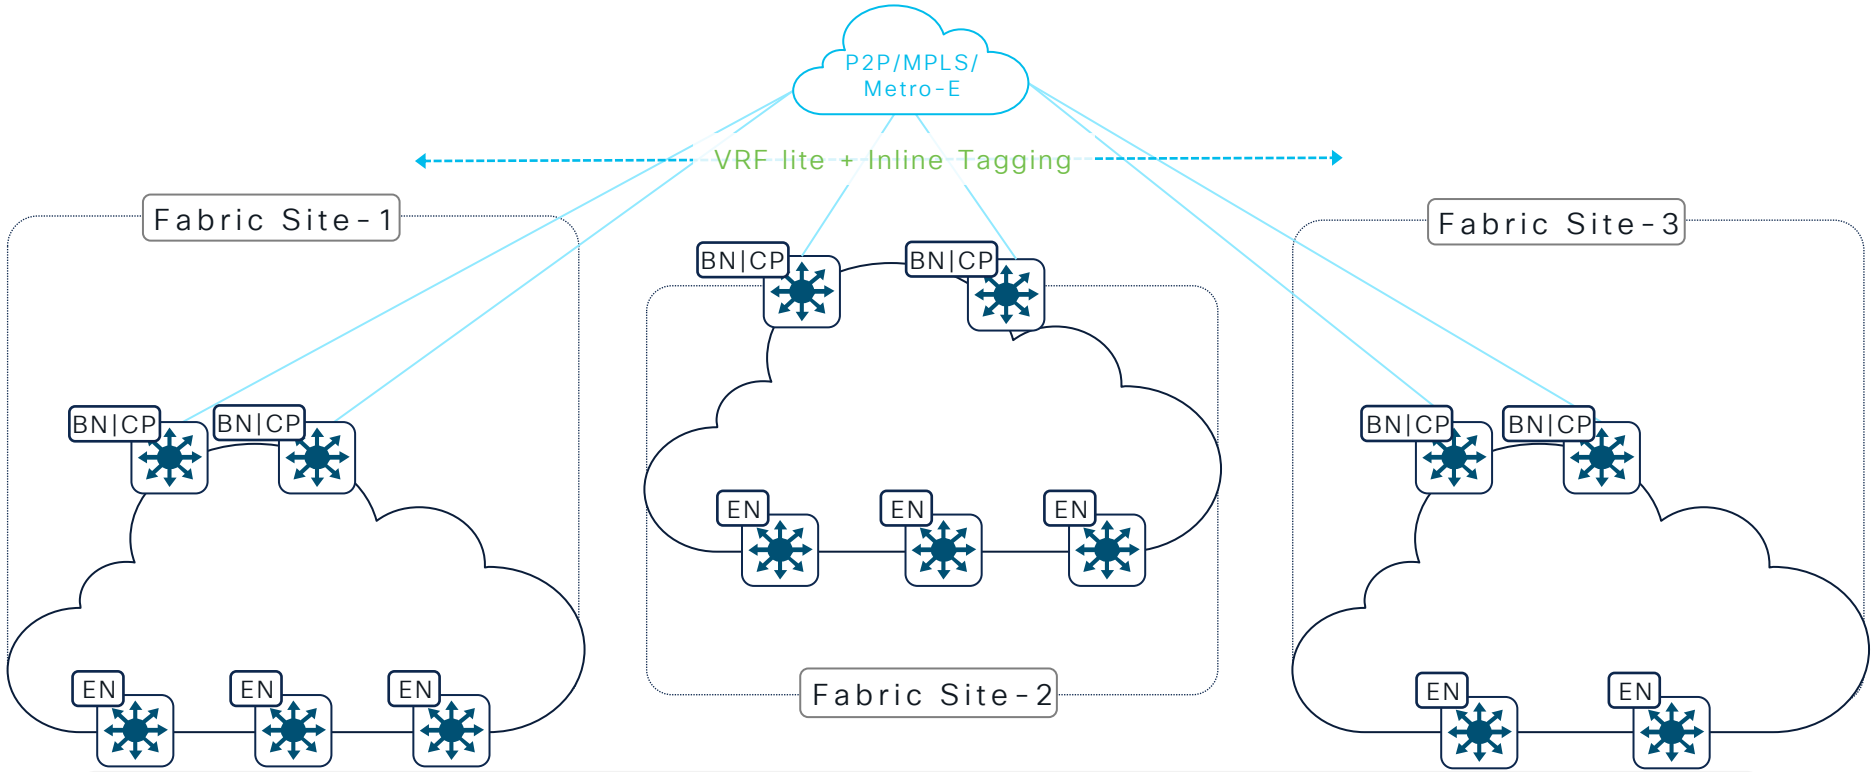

IP handoff at each fabric site

vrf lite handoff and Inline tagging to carry context between sites

Cisco SD-Access Architecture

Multisite Architecture with SD-Access Transit

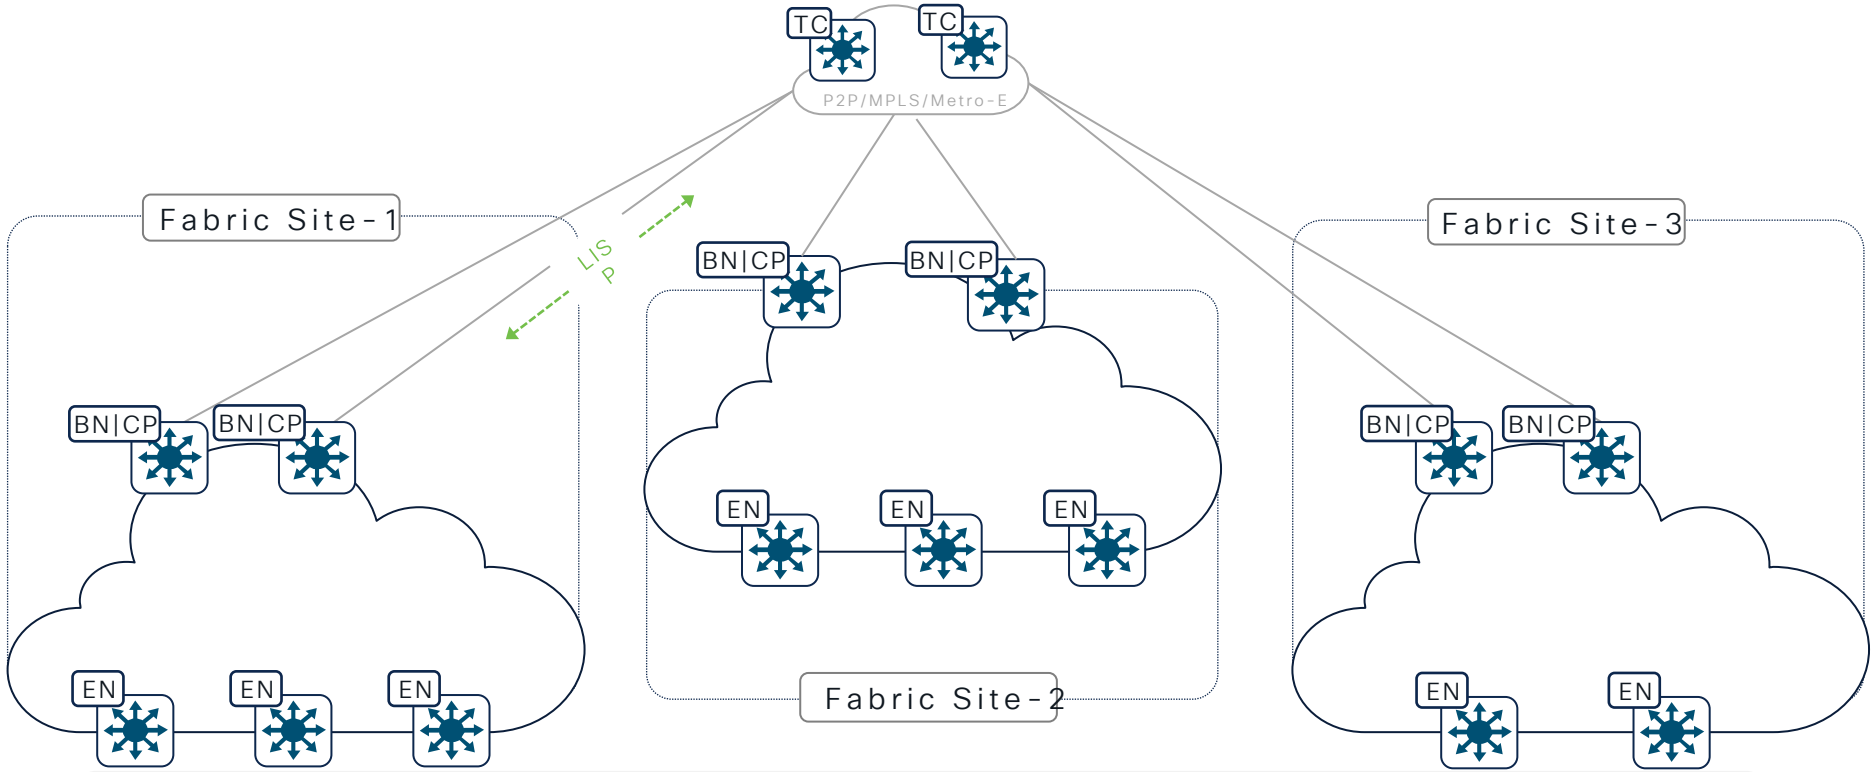

Higher MTU Support on WAN
Else use TCP Adjust-MSS (Available 2.3.7.0)

End to End Segmentation.
Automated configuration

Cisco SD-Access with LISP Architecture

Fabric Enabled Wireless

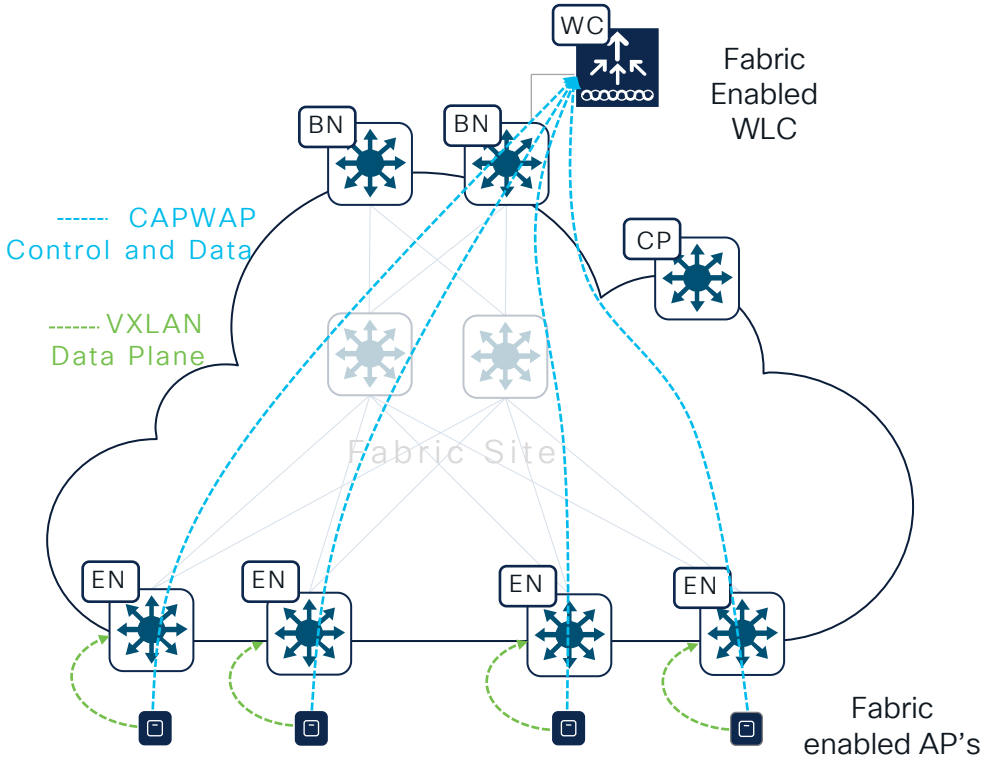

- Unified Wired and Wireless Architecture
- Simplified Management
- Consistent Wired and Wireless Policy
- End-to-End Automation for Wired and Wireless Deployment
- CAPWAP Control Plane, VXLAN Data Plane
- WLC/APs integrated in Fabric, LISP advantages

Cisco SD-Access with LISP Architecture
Wireless Options

Central Switching

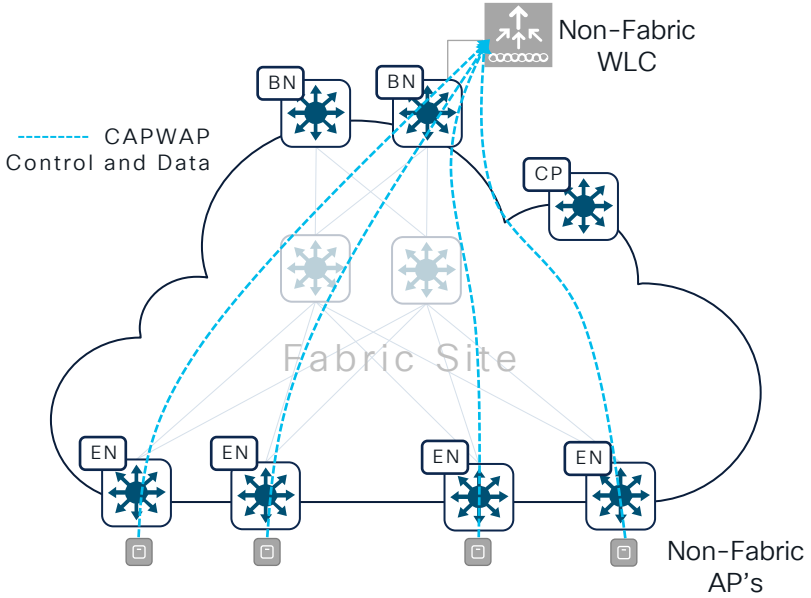

- CAPWAP for Control Plane and Data Plane
- SDA Fabric is just a transport
- Supported on any WLC/AP software and hardware

Flex Connect Mode

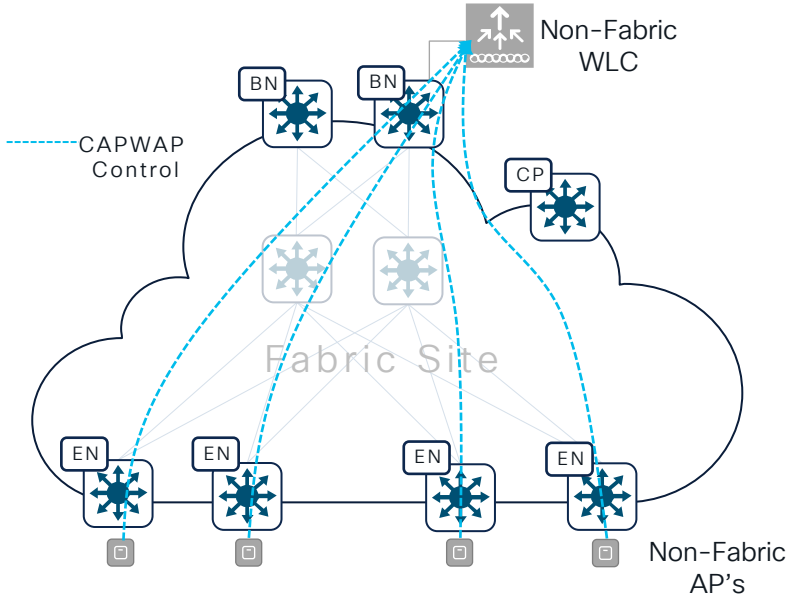

- CAPWAP for Control Plane
- Data plane is locally switched.
- Wireless traffic is treated like wired traffic.

Vxlan BGP EVPN

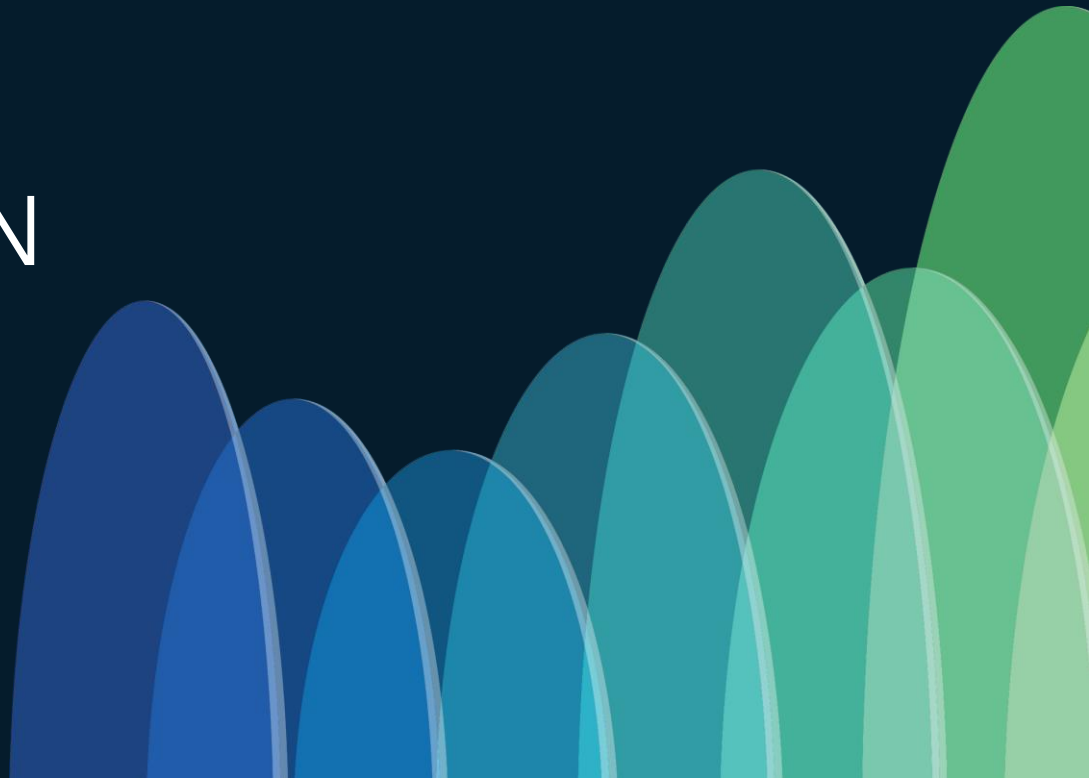

Vxlan BGP EVPN

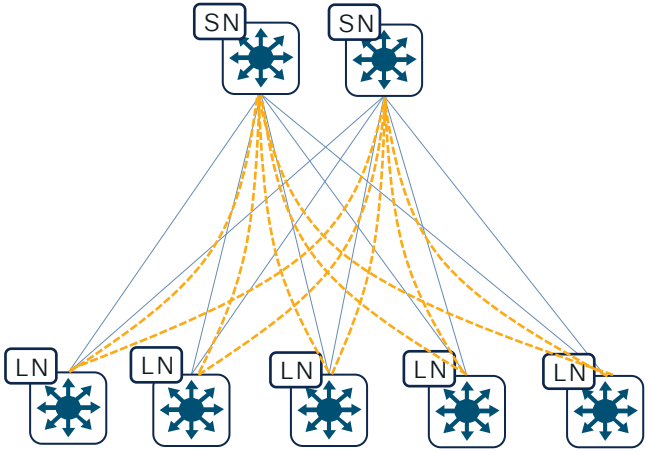

- Standards based Control-Plane
  - RFC 8365 (and RFC 7432)
  - Uses Multiprotocol BGP
- Uses Various Data-Planes
  - VXLAN (EVPN-Overlay), MPLS, Provider Backbone (PBB)
- Many Use-Cases
  - Integrated Routing and Bridging, MAC Mobility, Multi-Tenancy (VPN)

BGP EVPN System Role

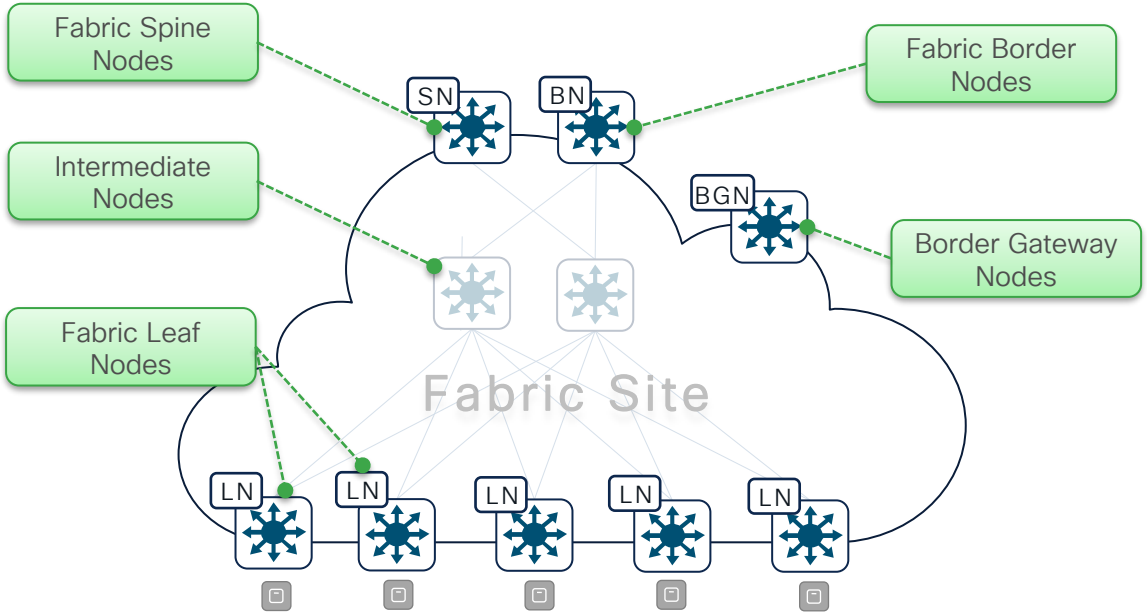

- **Leaf Nodes (VTEP)** – An origination and termination point of VXLAN enabled overlay network.
- **Spine Nodes**– An BGP EVPN reflects the L2/L3 VPN prefixes providing hierarchical neighbor peering, learning and distribution point.
- **Intermediate Nodes**– A Layer 2 or Layer 3 (IP/MPLS) Underlay network system providing basic transport and forwarding plane.
- **Border Nodes**– A gateway point of between EVPN fabric and external network domain.
- **Border Gateway**– A gateway point of between two or more BGP EVPN administrative domain boundary.

Catalyst EVPN Scale and Performance Matrix

Cisco Catalyst BGP EVPN Configuration Guide
Scale and Performance Chapter

Vxlan with BGP EVPN

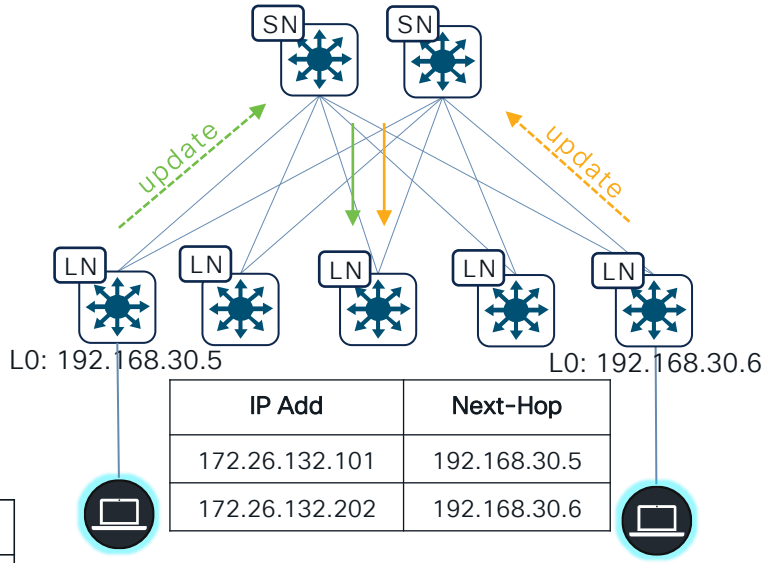

l2vpn table

IP Add	Next-Hop
172.26.132.101	G1/0/0
172.26.132.202	192.168.30.6

l2vpn table

IP Add	Next-Hop
172.26.132.202	G1/0/0
172.26.132.101	192.168.30.5

Vxlan with BGP EVPN

- MAC or MAC/IP moves (local/local, local/remote, remote/local, remote/remote)
- Duplicate detection timers can be adjusted

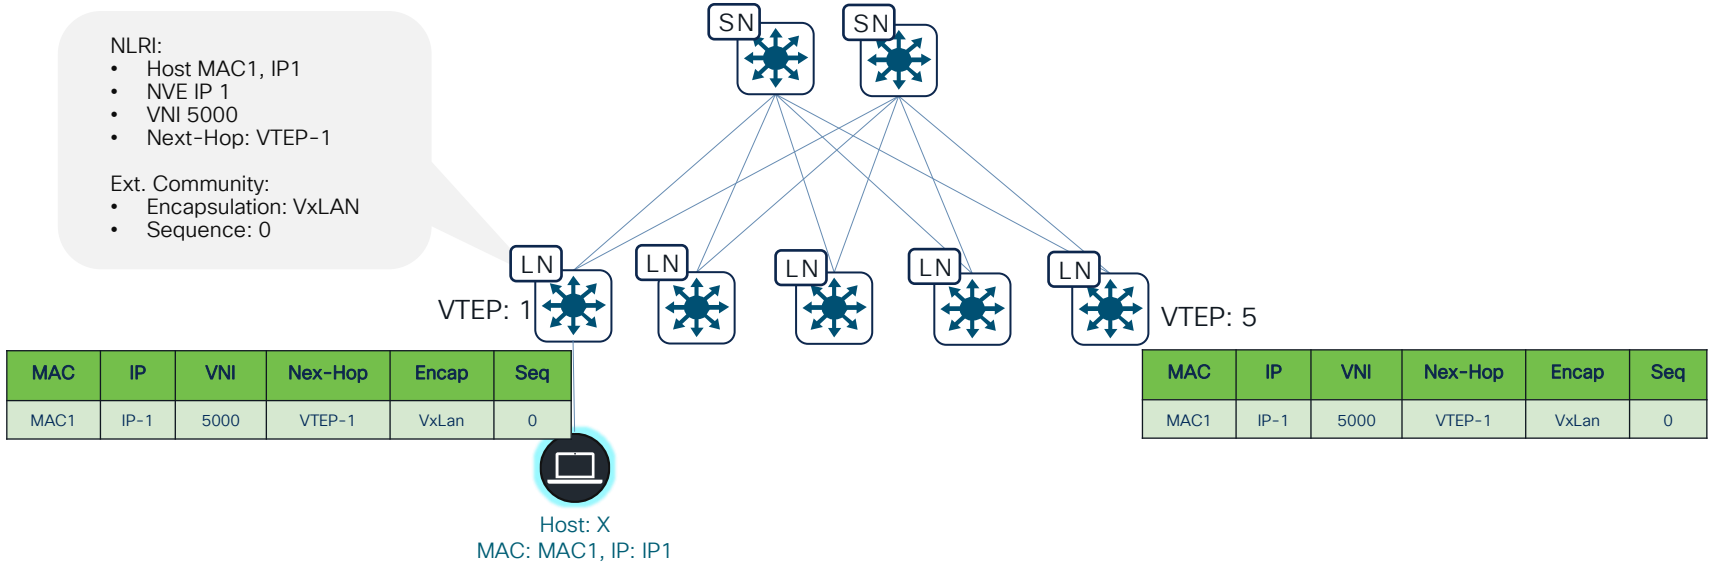

1. VTEP-1 detects Host1 and advertise an EVPN route for Host1 with Seq# 0

Vxlan with BGP EVPN

- MAC or MAC/IP moves (local/local, local/remote, remote/local, remote/remote)
- Duplicate detection timers can be adjusted

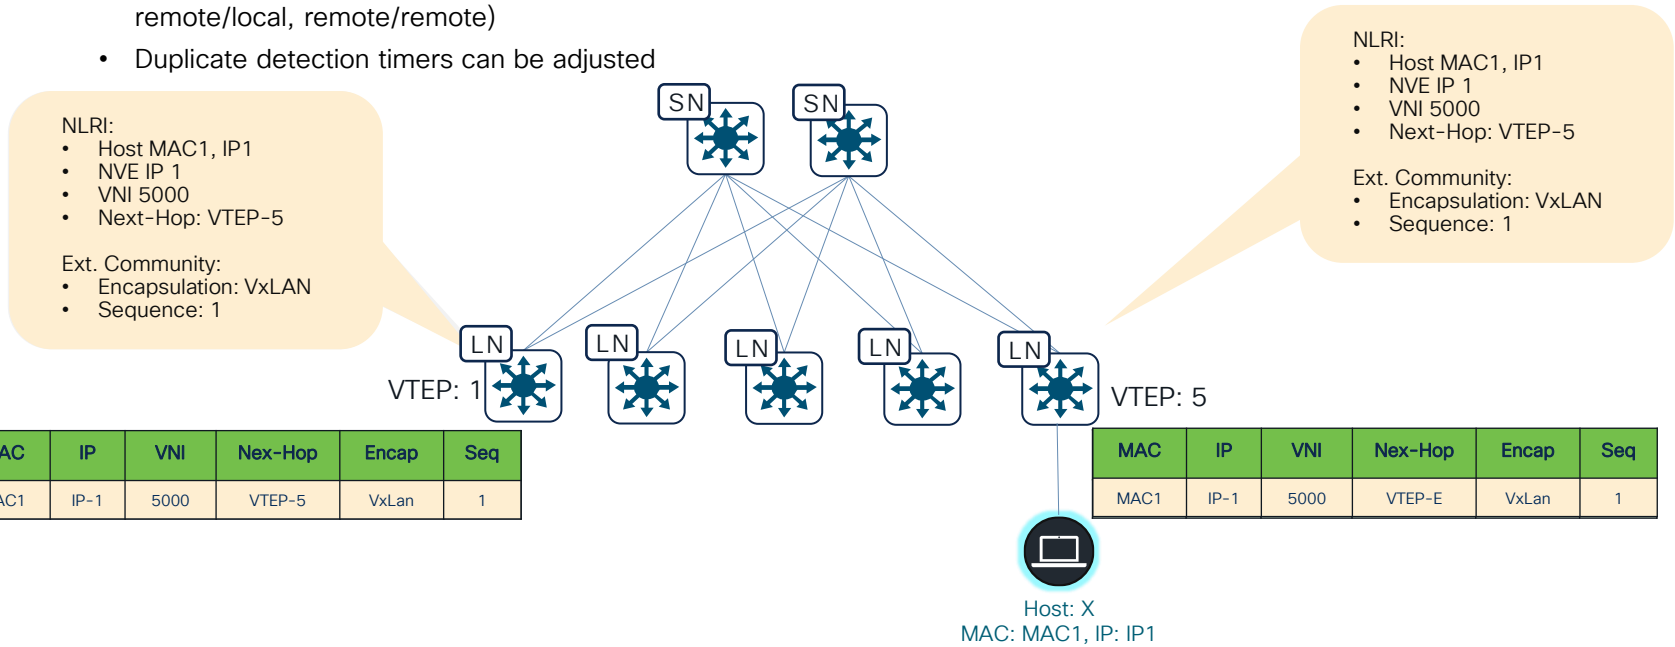

MAC	IP	VNI	Nex-Hop	Encap	Seq
MAC1	IP-1	5000	VTEP-5	VxLan	1

MAC	IP	VNI	Nex-Hop	Encap	Seq
MAC1	IP-1	5000	VTEP-E	VxLan	1

1. VTEP-1 detects Host1 and advertise an EVPN route for Host1 with Seq# 0
2. Host1 Moves behind VTEP-5
3. VTEP-5 detects Host1 and advertises an EVPN route for Host1 with Seq #1
4. VTEP-1 sees more recent route and withdraws its advertisement

Routing and Bridging Overlay Types

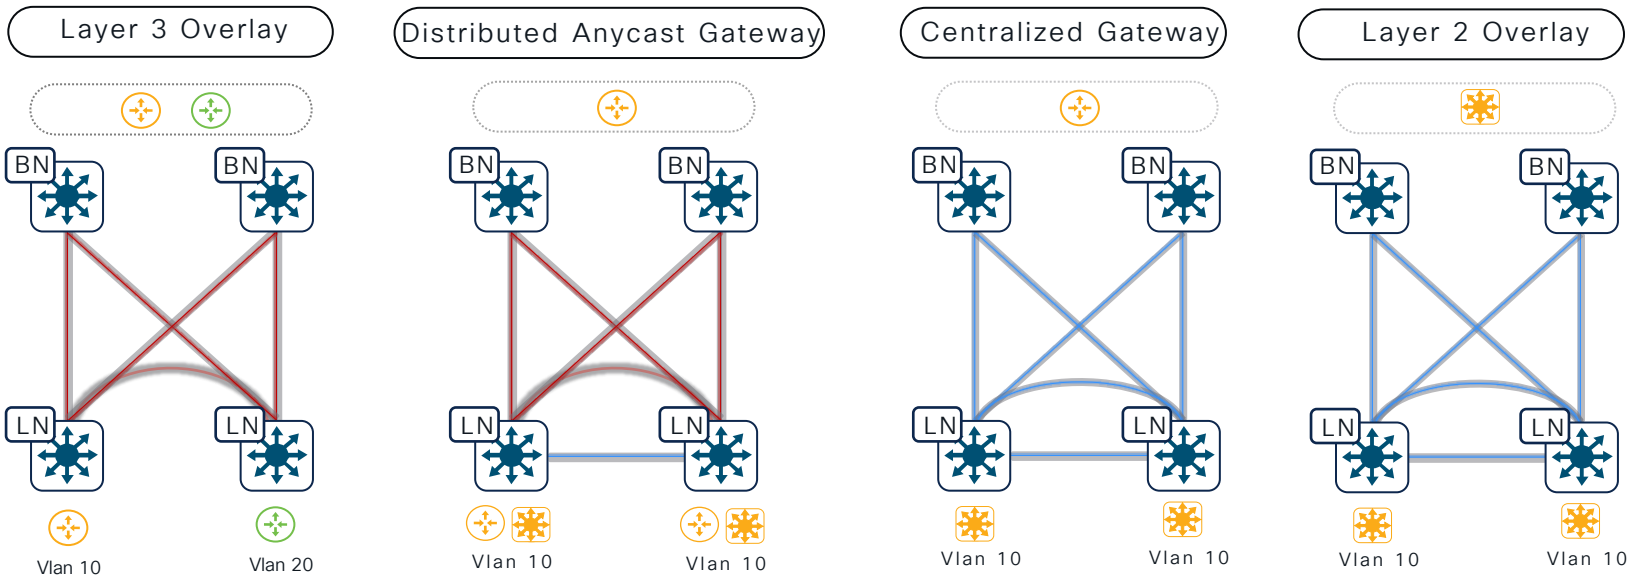

Overlay Types

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

Vxlan with BGP EVPN
Wireless Options

Central Switching

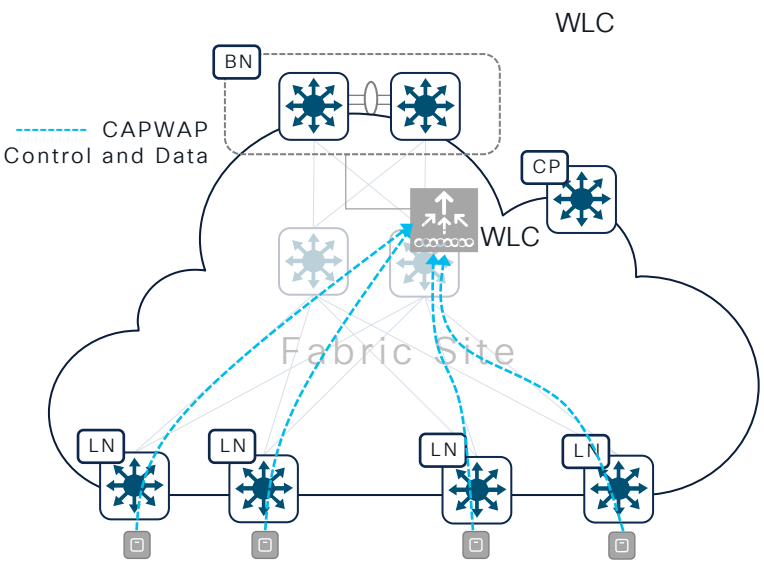

Flex Connect Mode

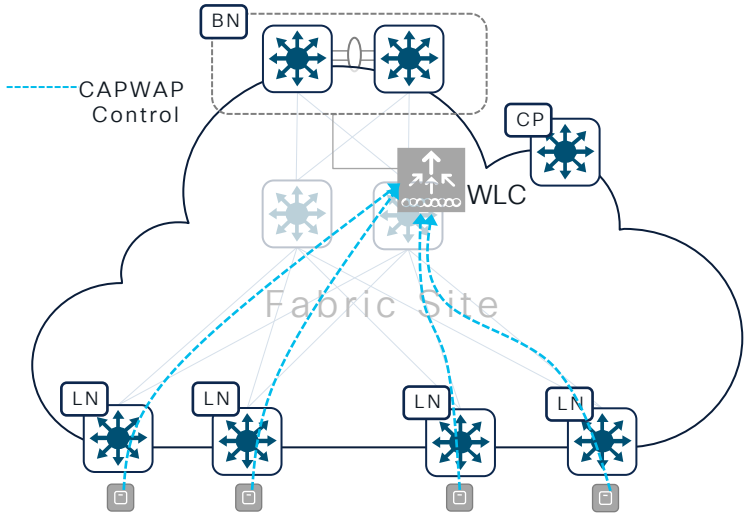

Wireless

Over the Top Wireless. Intact WLC and AP communication in Underlay

Flexible SSID alternatives – Central Switching, Local Switching, Central + Local Switching

Fabric boundary initiates from Wireless Client IP gateway.

Single Cluster Fabric Architecture

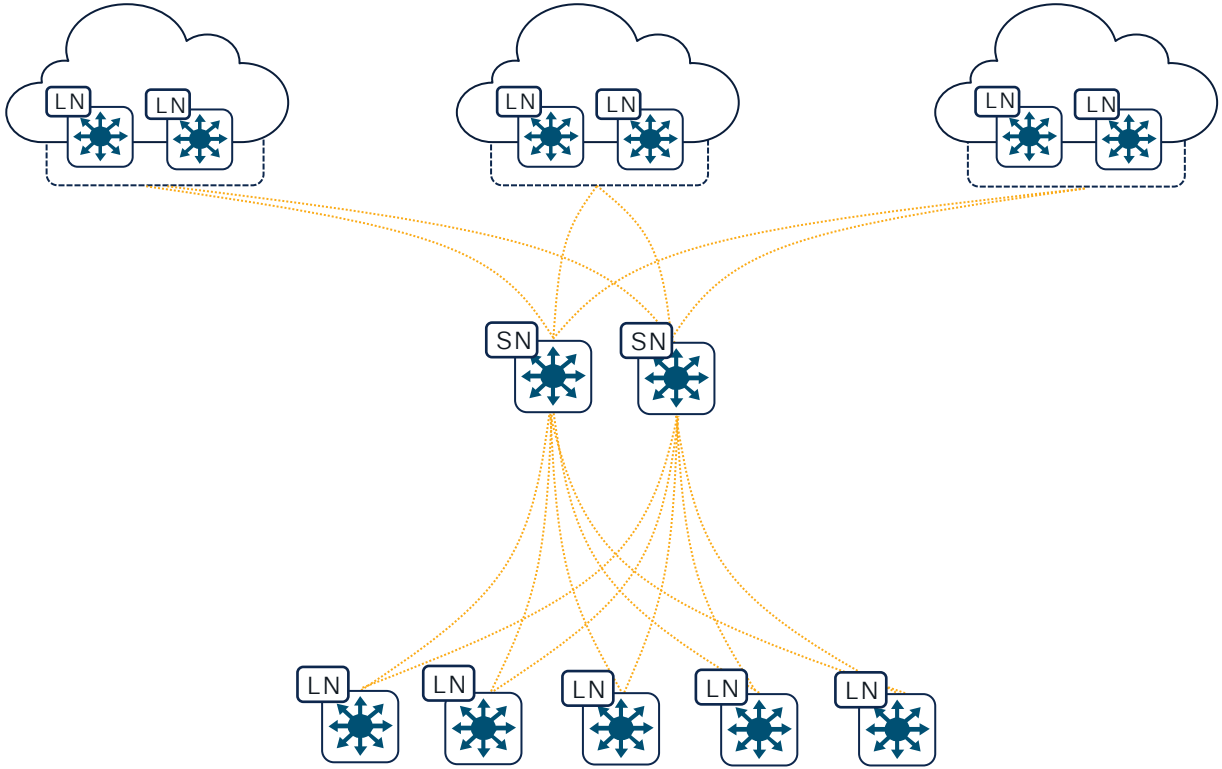

Shared Spine

- 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

Multi Cluster Fabric Architecture

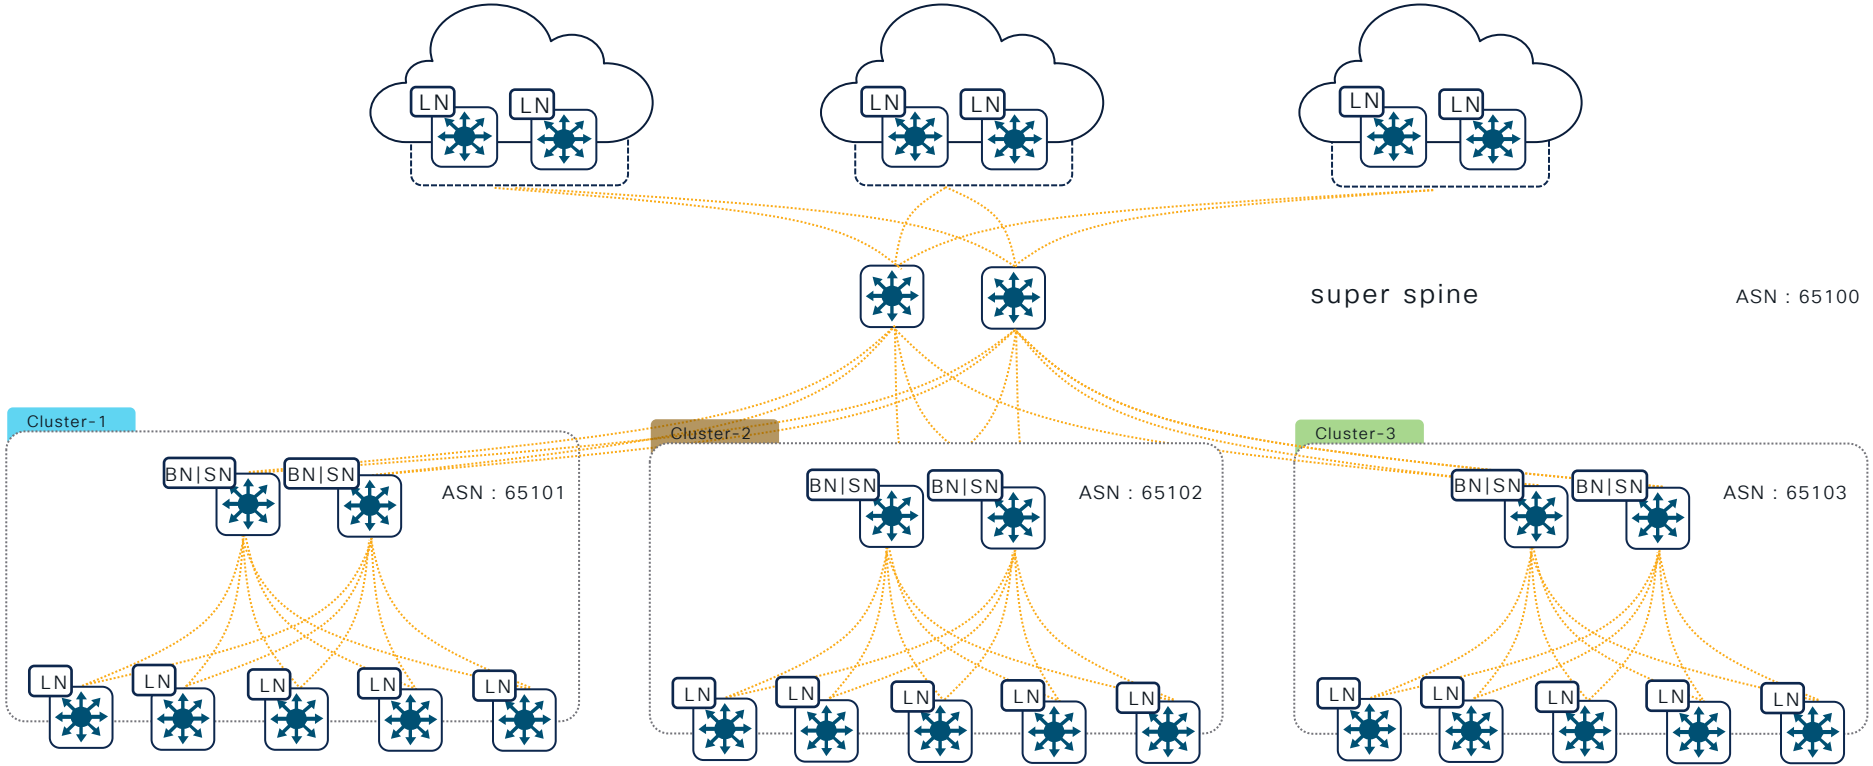

Scalable Fabric Cluster

- Increase fabric domain scale with hierarchical dynamic overlay VXLAN tunnels per fabric cluster
- Consistent Layer 2 domain scale size as traditional non-fabric networks
- Scalable overlay routing with per-VN prefix summarization and re-origination by each Border Spine cluster
- End-to-End Unicast IPv4/IPv6 support. Layer 2 Mobility and overlay Multicast limited per cluster

Multi-Site Fabric Architecture

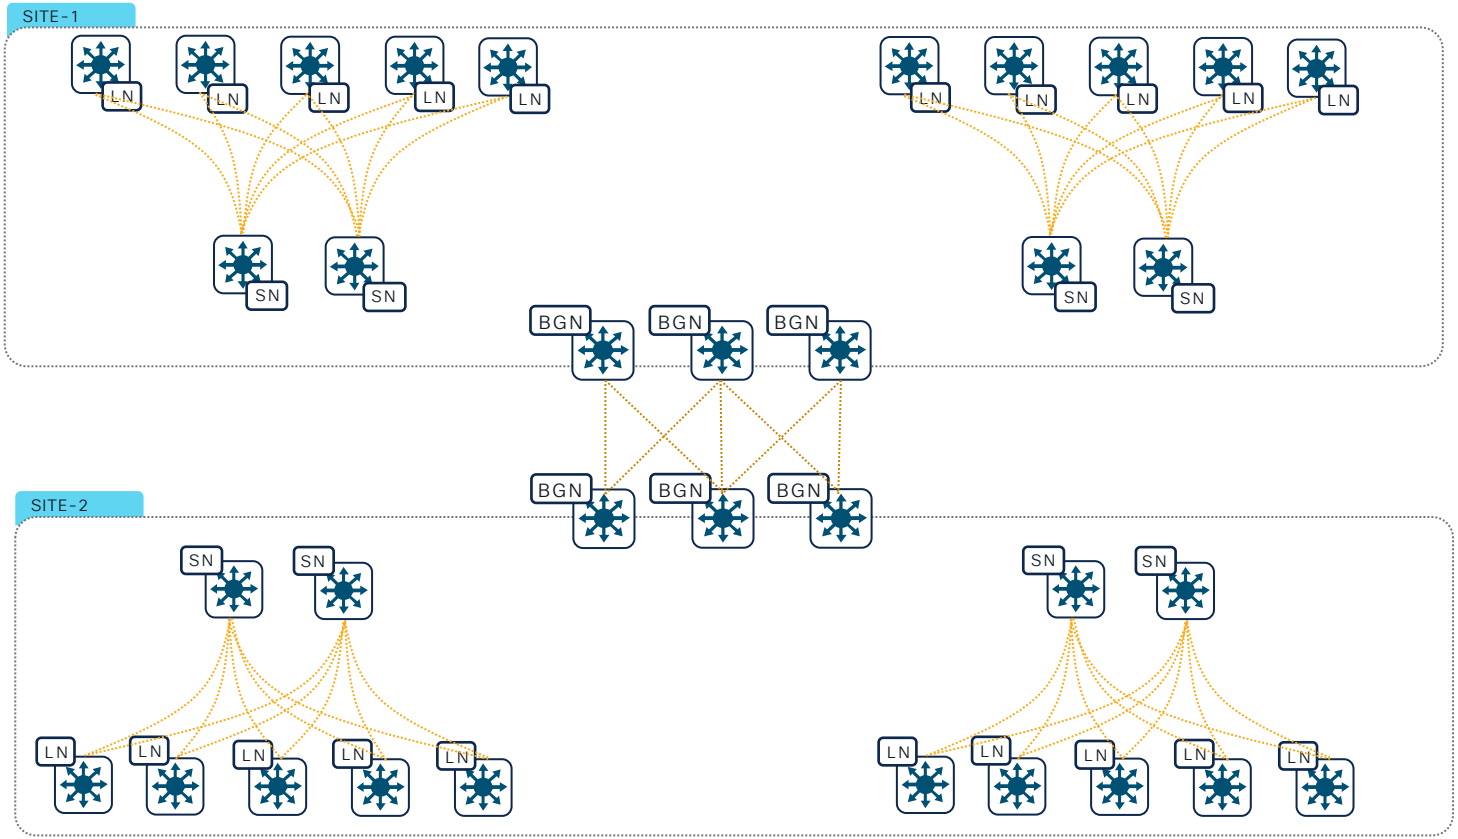

Multisite Fabric

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

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

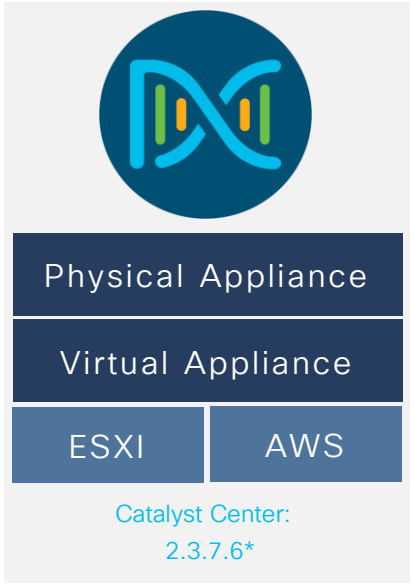

IOS-XE: 17.12

Catalyst Advantage

Underlay	Layer 3 Access | Cisco Stack-Wise Virtual
Overlay	Layer 3 | Distributed Anycast Gateway
Multicast	TRM | Default/DATA MDT
Assurance	Fabric Site | Infrastructure

Catalyst 9300	Catalyst 9400	Catalyst 9500	Catalyst 9600
9300/L/LM 9300B | 9300X	9400 SUP1/1XL/Y 9400X SUP2/XL	9500 | 9500H	9600 SUP1
Stackwise |	Stackwise-Virtual		

Overlay	Greenfield Deployments
Underlay	Greenfield or Brownfield Deployments

SD-Access with BGP EVPN Catalyst Center

Beta Signup

* Private beta

Catalyst 9000 Provides
Segmentation with Architecture Flexibility

Cisco Preferred

Campus optimized LISP Fabric

- ✓ Industry Standard based and Optimized for enterprise campus
- ✓ Fully Automated end to end Fabric
- ✓ Integrated Wireless with Fabric for faster convergence
- ✓ Single box fabric for Branch-in-a-box use case
- ✓ Lightweight, massive scale with rapid convergence and highly extensible to address newer use cases and drive innovations like Pub-Sub , Multi-Site, Extranet etc.

BGP-EVPN Fabric

- ✓ Multi Vendor Interoperability
- ✓ One Fabric Architecture across Campus and Datacenter
- ✓ Proven and Scalable leveraging BGP Control-plane
- ✓ Multi-Tier and Hierarchical Overlay Network Architecture
- ✓ Use-case driven customizable overlay network types and topologies.

Fill Out Your Session Surveys

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

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

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

Content Catalog

Continue your education

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

Contact me at: ragoli@cisco.com


cisco
Thank you
cisco Live!

cisco Live!

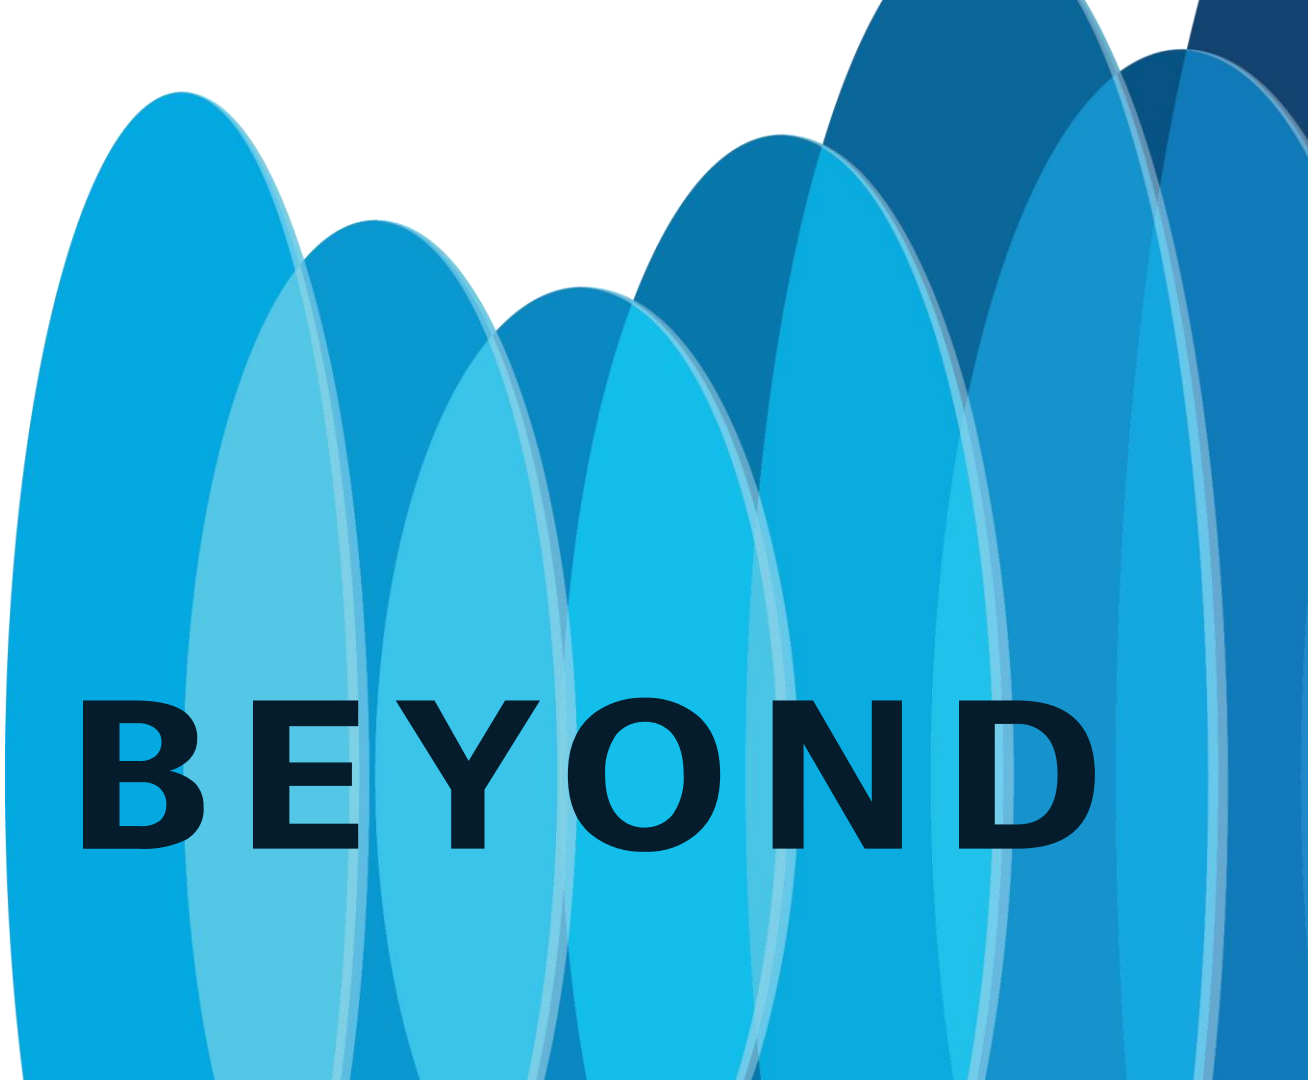

GO BEYOND