SDA Introduction

NtwrkPeople DET

Jesse Schmidt
Solutions Engineer

21-05-2026

Agenda

01 Why Cisco SD-Access LISP?

02 Roles & Terminology

03 Fabric Fundamentals

04 Multiple Fabric Sites

05 Design

Why Cisco SD-Access LISP?

Traditional Networking Challenges

Network Deployment Challenges

Network Infrastructure

Switching Routers Wireless

Network Security Challenges

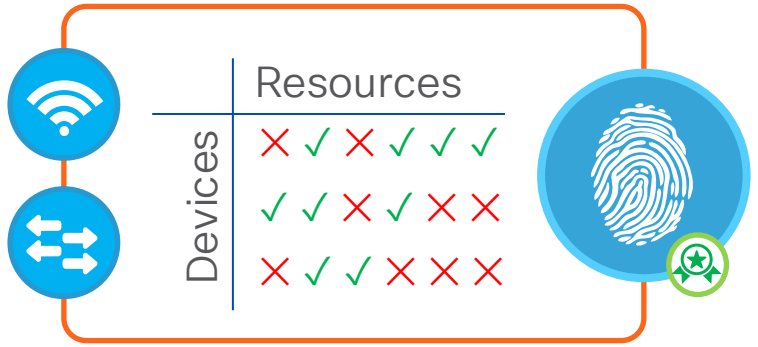

Wireless and Wired Challenges

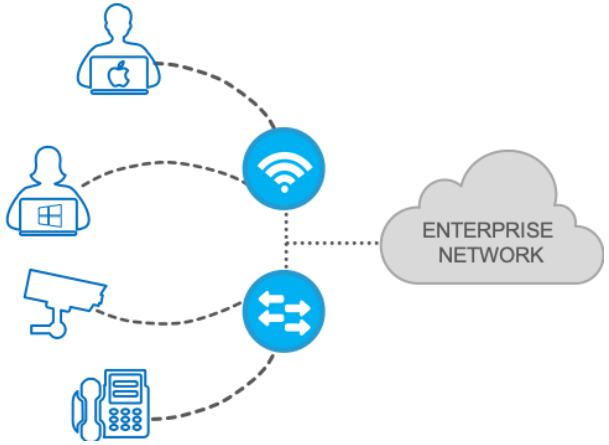

Network Operations Challenges

SDA Dramatically Simplifies the Network

Before

The *Traditional* Enterprise **Protocol** Stack

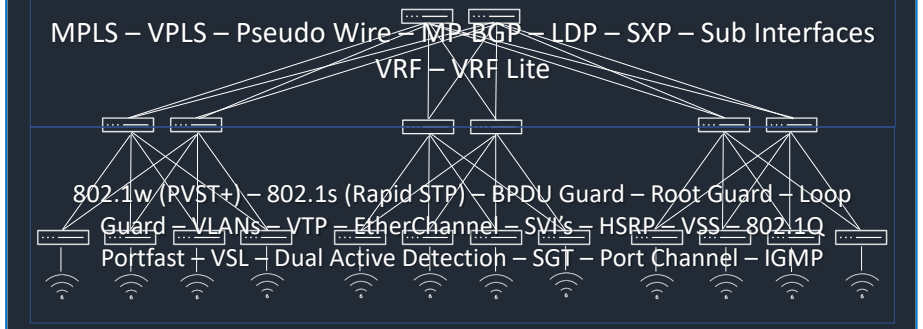

Unique **Configurations** Drive Up Complexity

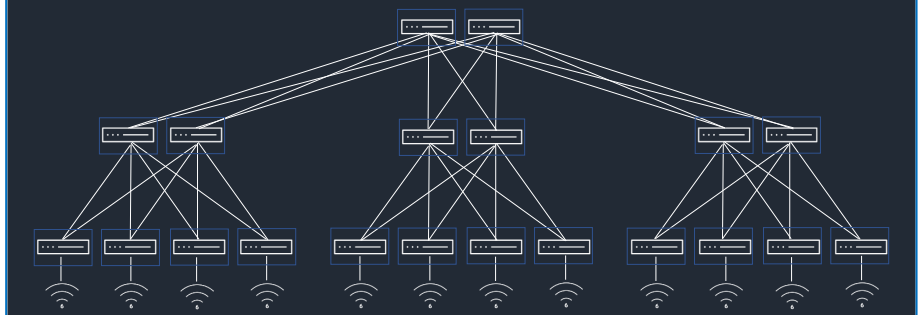

After

NEW Enterprise Protocol *Stack* with **SDA**

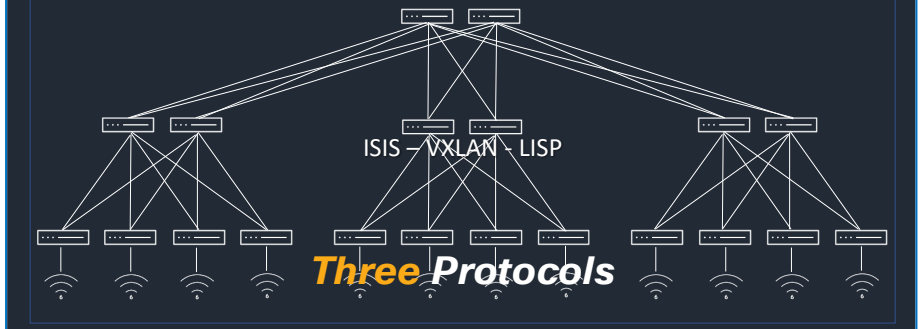

Common **Configurations** with **SDA**

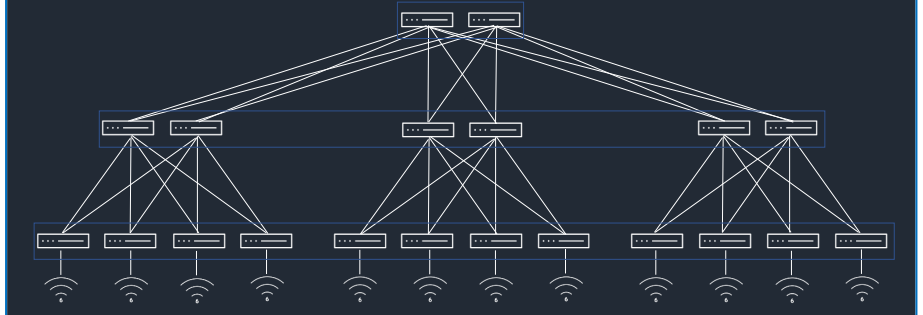

***Eliminates* Layer 2** and ***Simplifies* Network Protocol Stack**

SDA Dramatically *Simplifies* the Network

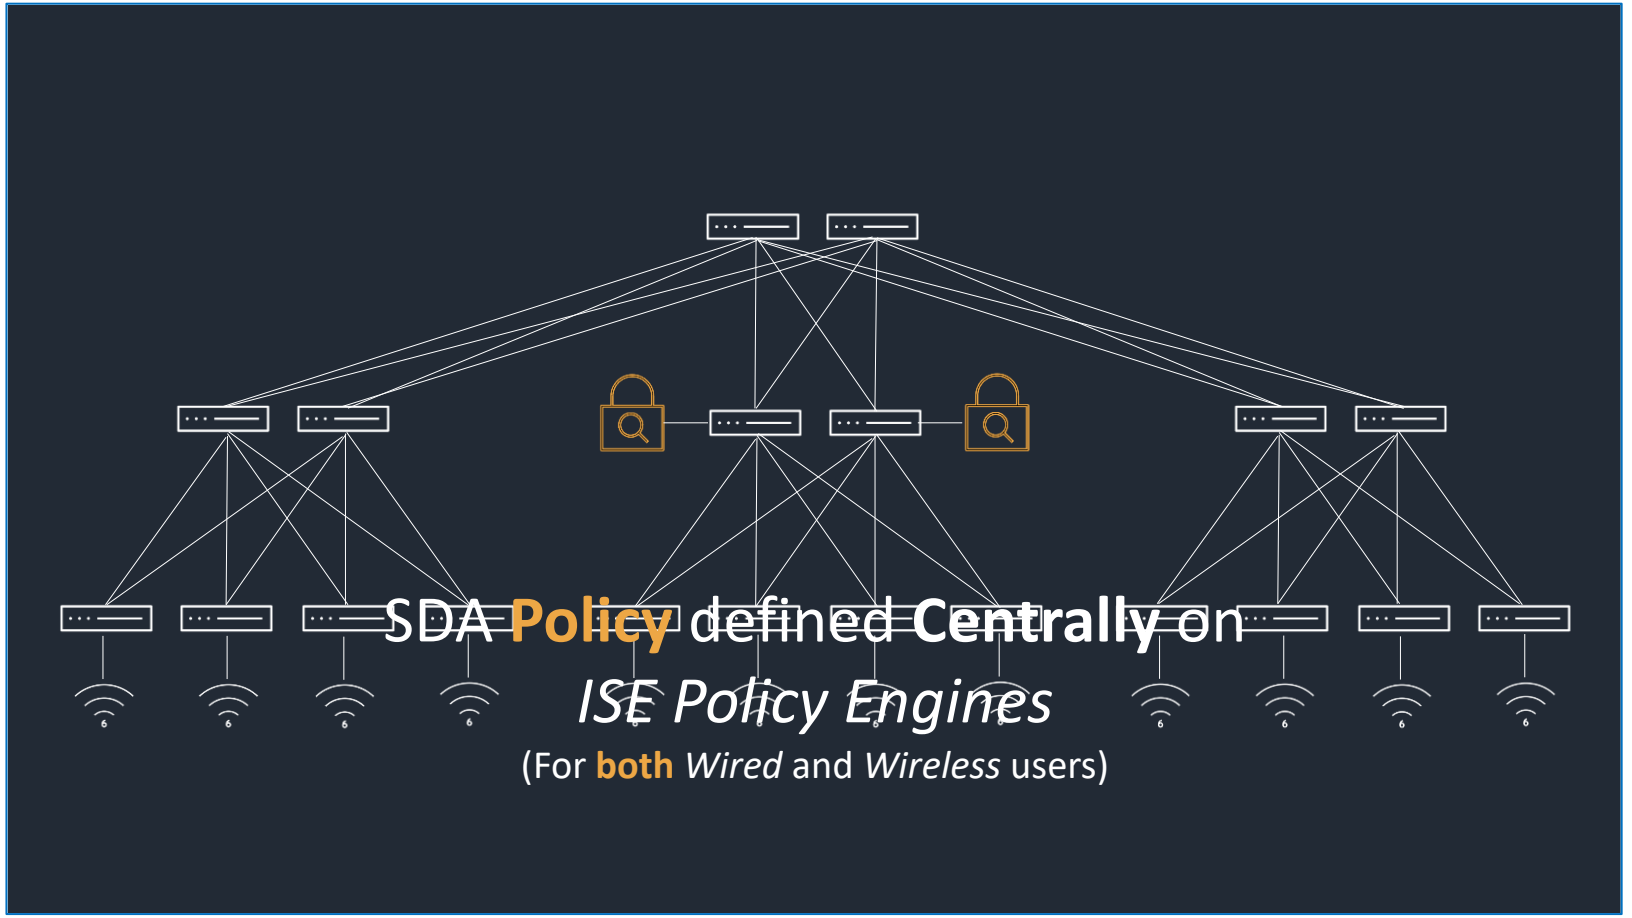

Single Policy Store for
Wired and ***Wireless Users***

Cisco Software-Defined Access

Intent-Based Networking

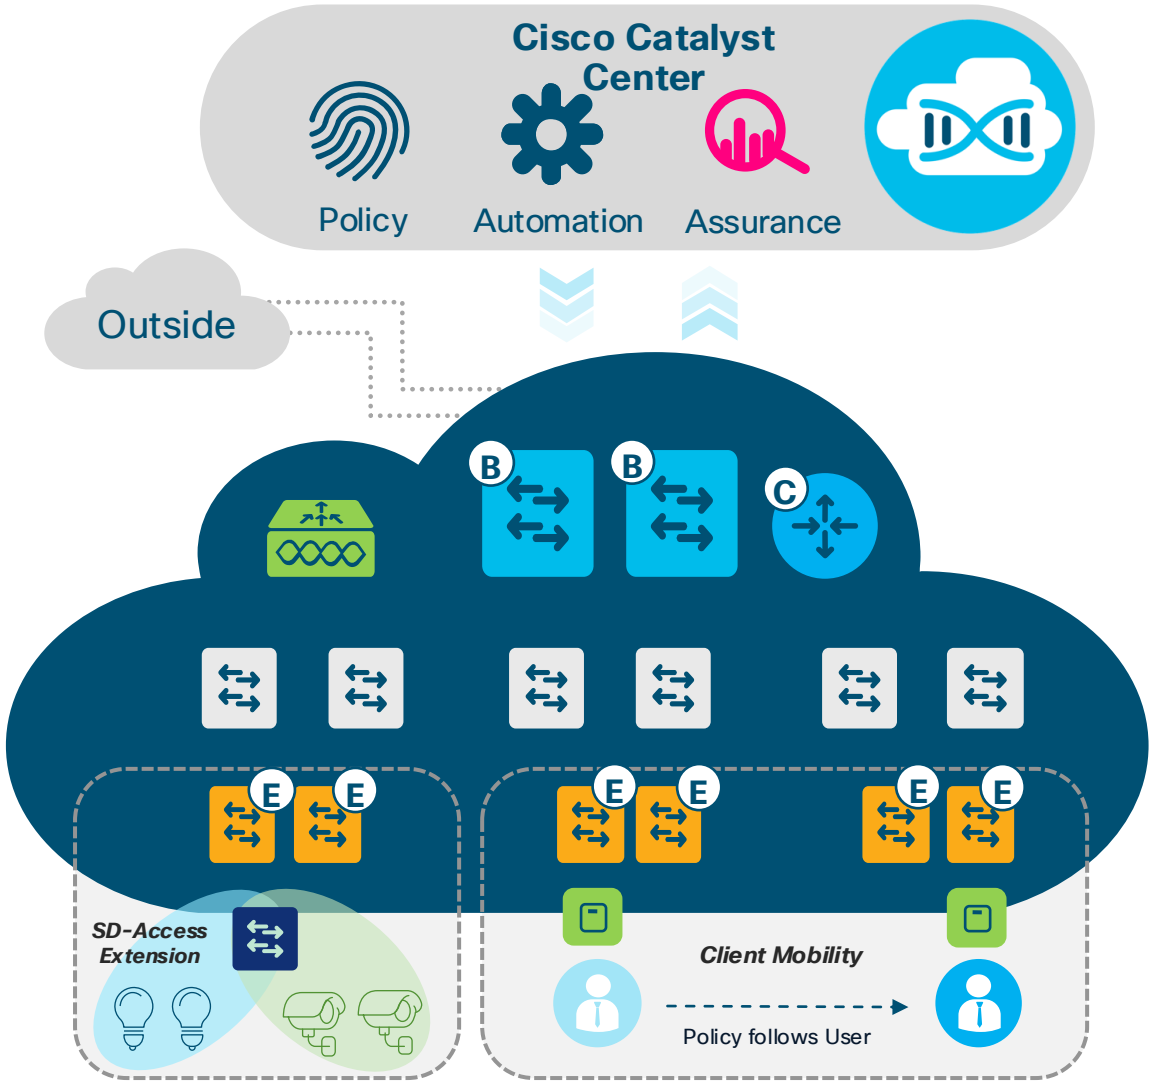

One Automated Network Fabric

Single fabric for wired and wireless with full automation

Identity-Based Policy and Segmentation

Policy definition decoupled from VLAN and IP address

AI-Driven Insights and Telemetry

Analytics and visibility into user and application experience

Modern, Open and Scalable Fabrics

IETF Standard based Protocols

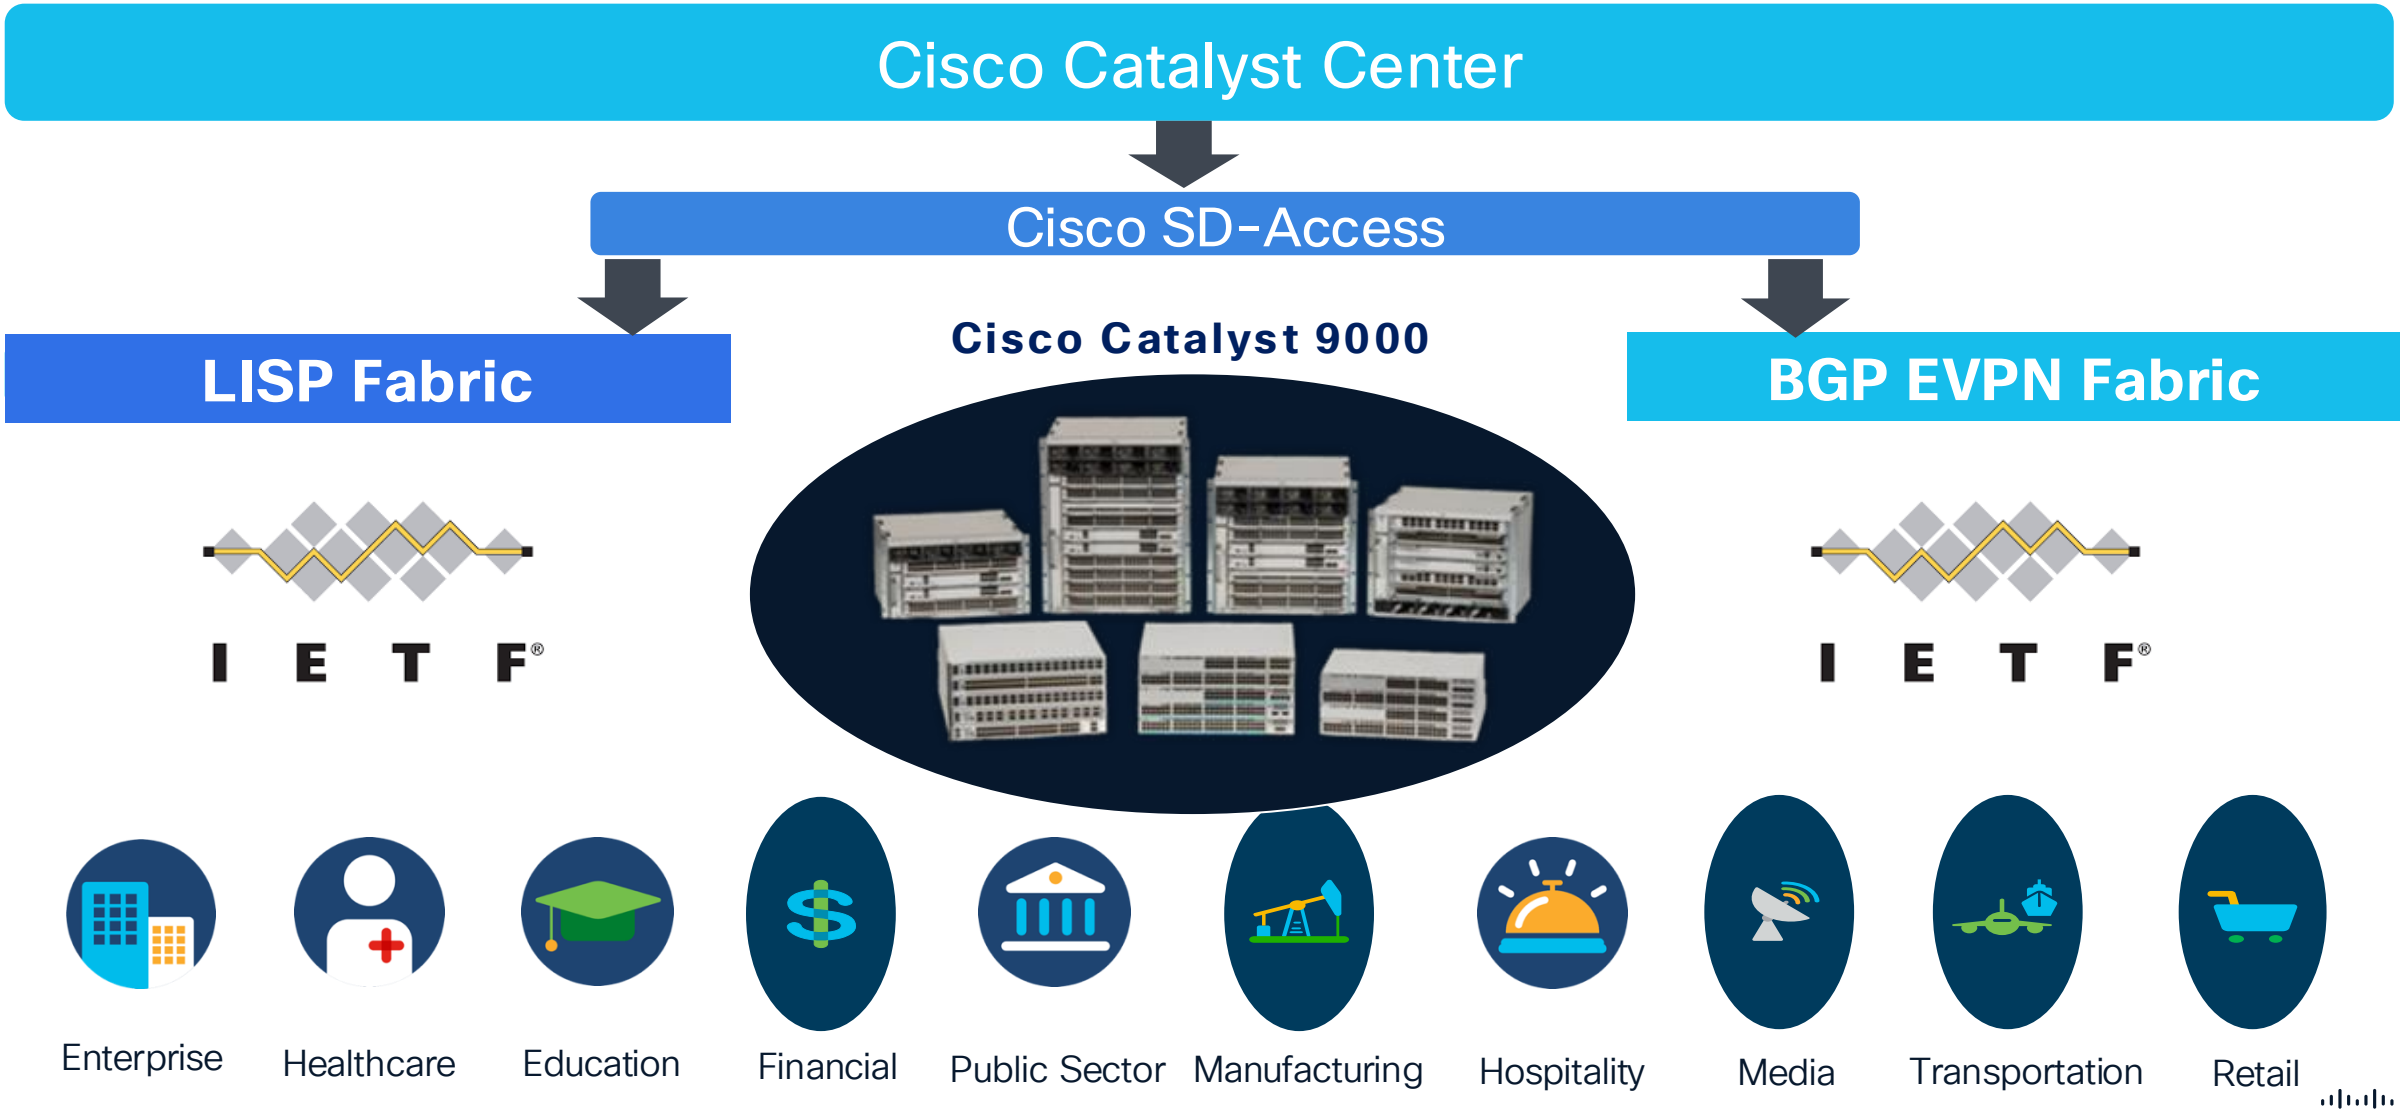

Flexible Fabric Options Tailored to *Customer Outcomes*!

Cisco SD-Access with LISP Control Plane VXLAN Data Plane

Network Simplification

Lightweight, extensible, massive scale with rapid convergence. Single overlay for wired and wireless

Mobility First Requirement

Fabric Integrated Wireless, L2 Mobility, enhanced wireless performance

Segmentation

Zero-Trust Architecture with Unified Wired and Wireless Policy

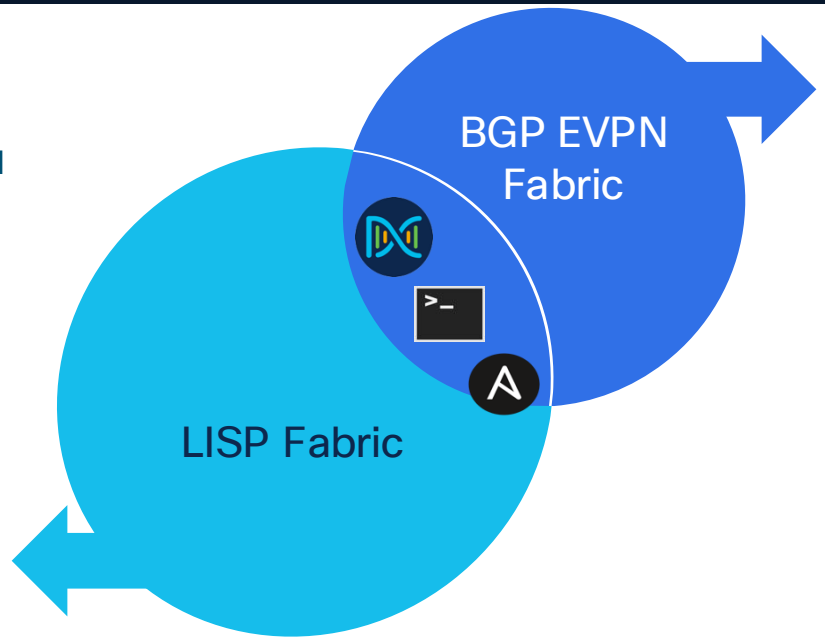

BGP EVPN Control Plane VXLAN Data Plane

One Fabric Architecture (Campus and DC)

Operational ease with a single familiar protocol

Multi-vendor interoperability

Vendor-agnostic solution with unique Cisco differentiators

One Infrastructure | Single Data plane | Consistent Zero-Trust Experience

Roles and Terminology

1. Concepts
2. SD-Access Roles
3. Fabric Constructs

What is a Network Fabric?

- Transports data from source to destination.
- Mesh of connections between network devices.
- Usually refers to a virtualized, automated lattice of overlay connections.

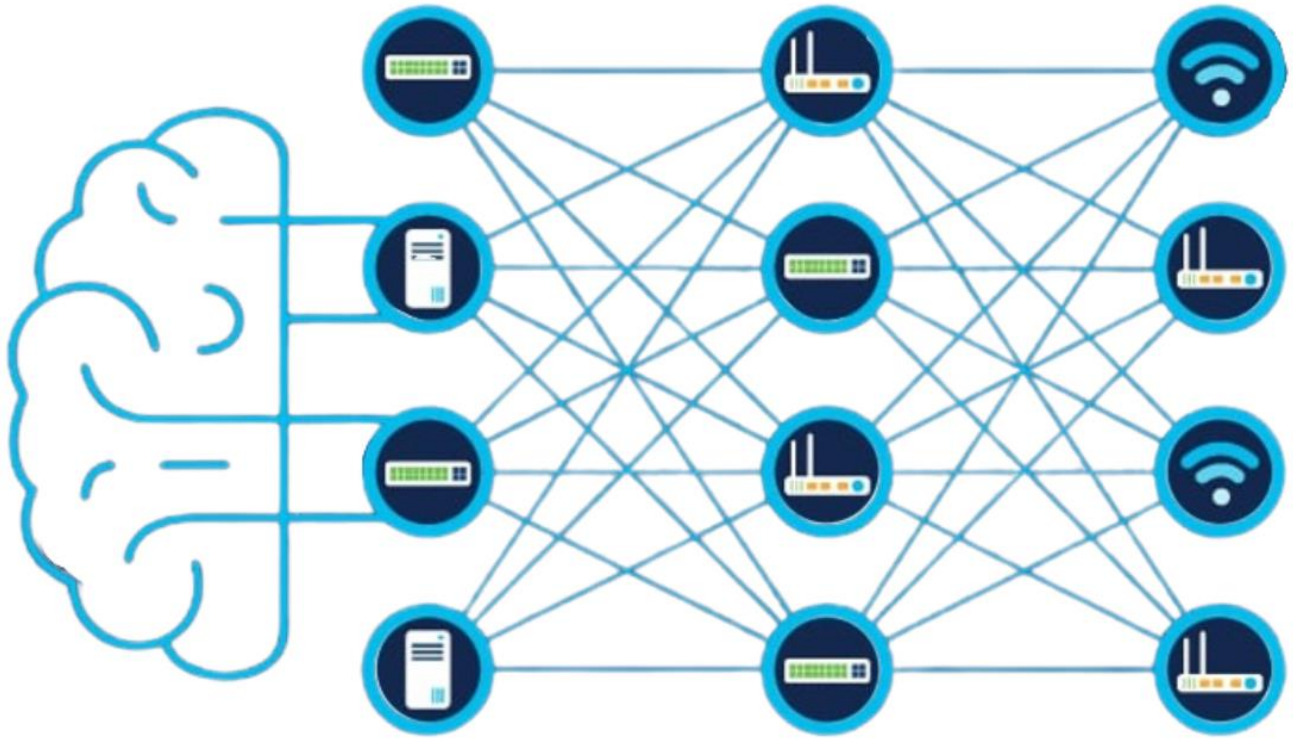

Underlay and Overlay

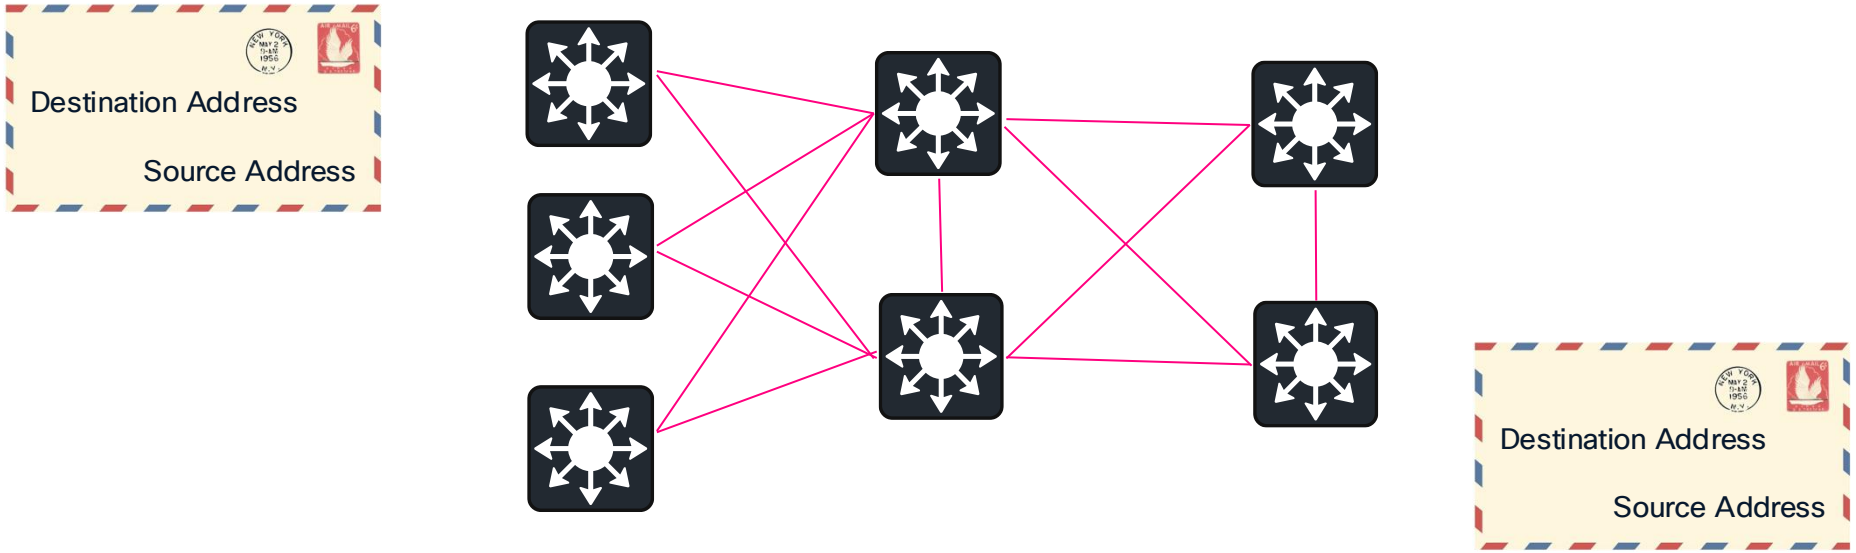

Underlay Network = Physical Infrastructure to provide IP reachability with redundancy and resiliency.

Underlay and Overlay

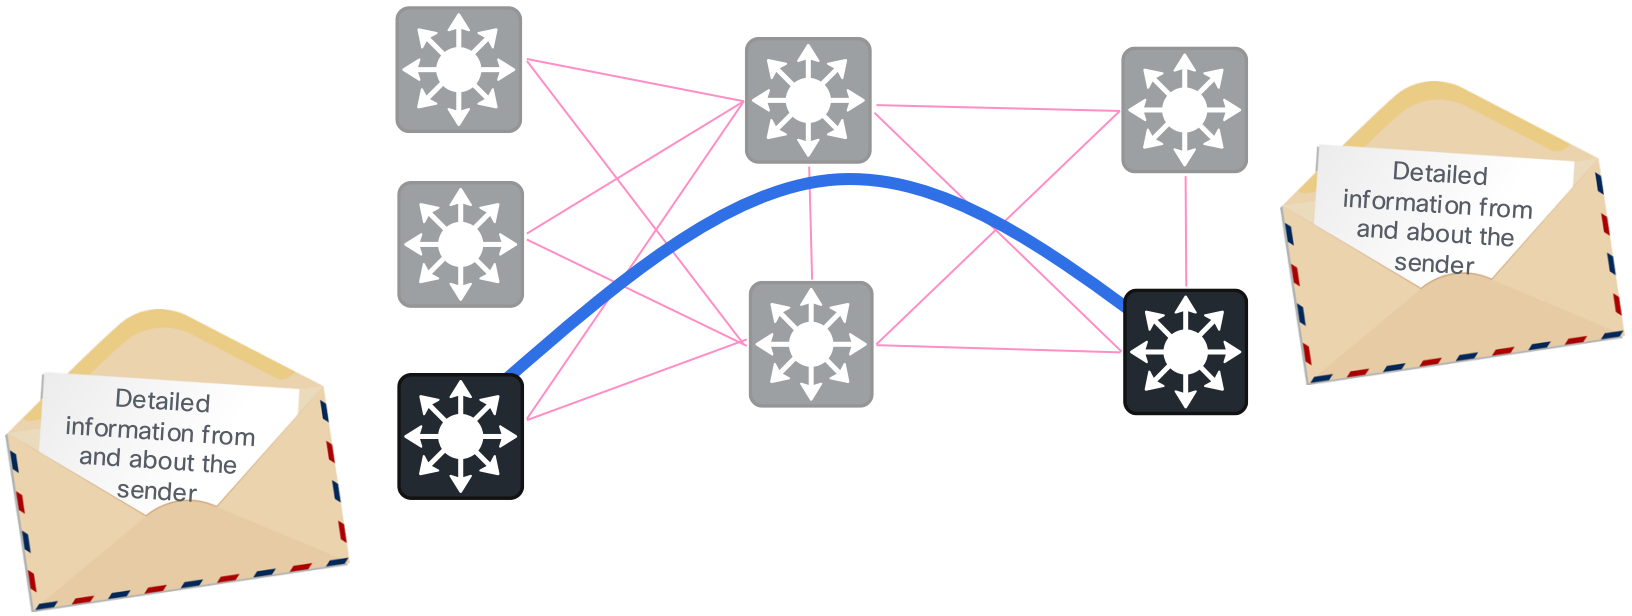

Overlay Network = Logical topology used to virtually connect devices to provide additional services, not delivered by the Underlay.

Underlay and Overlay

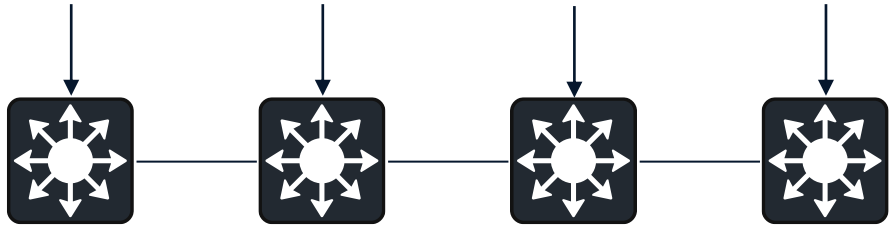

Would you configure network segmentation hop-by-hop?

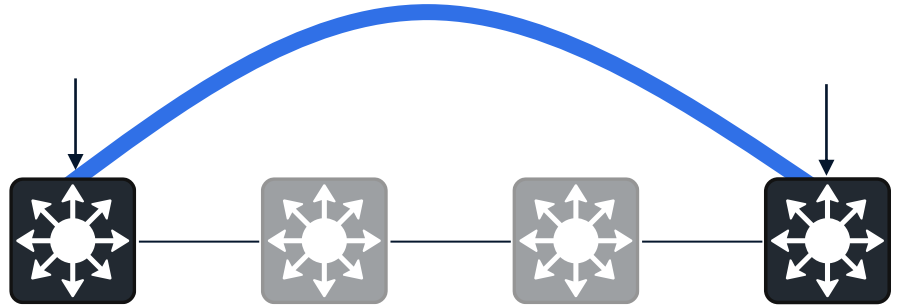

Or simply carry the segmentation tags in the overlay?

Underlay and Overlay

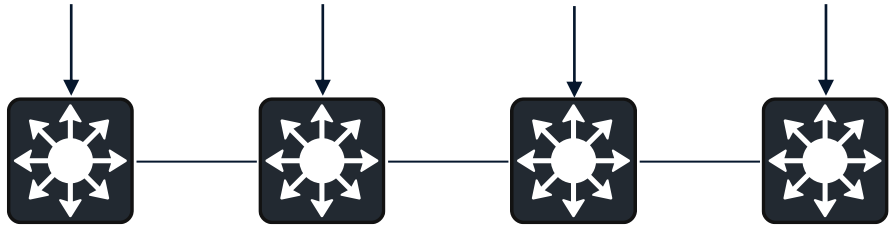

Would you configure network segmentation hop-by-hop?

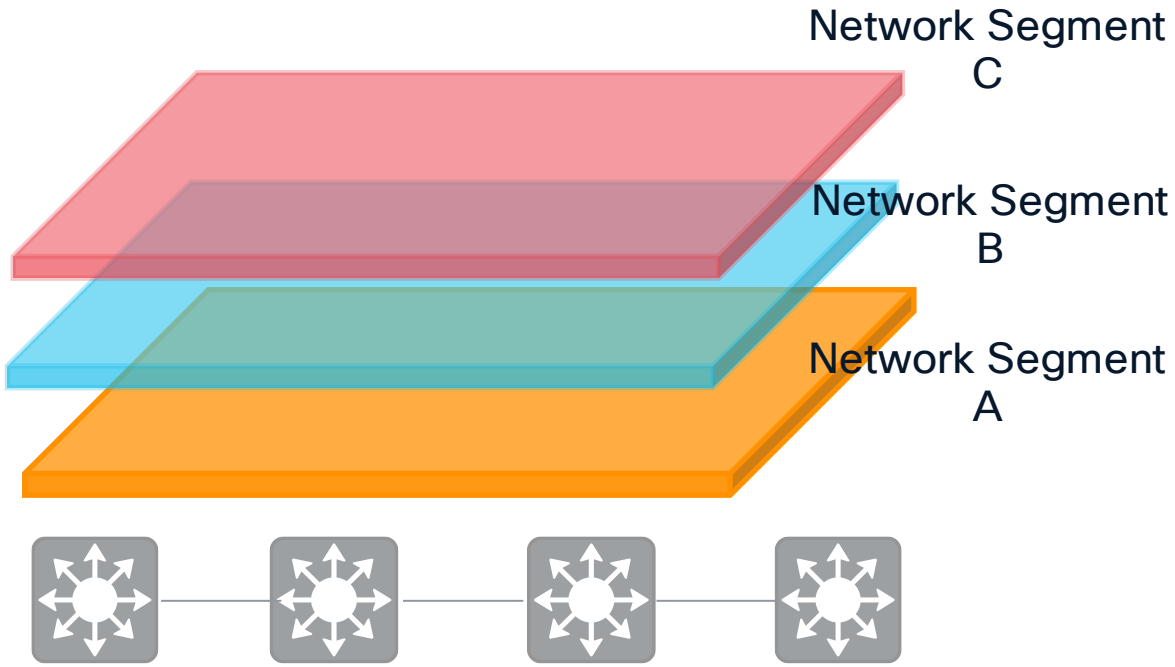

Multiple segments in the overlay that underlay is unaware of!

Underlay and Overlay

In context of SD-Access LISP

What about Underlay?

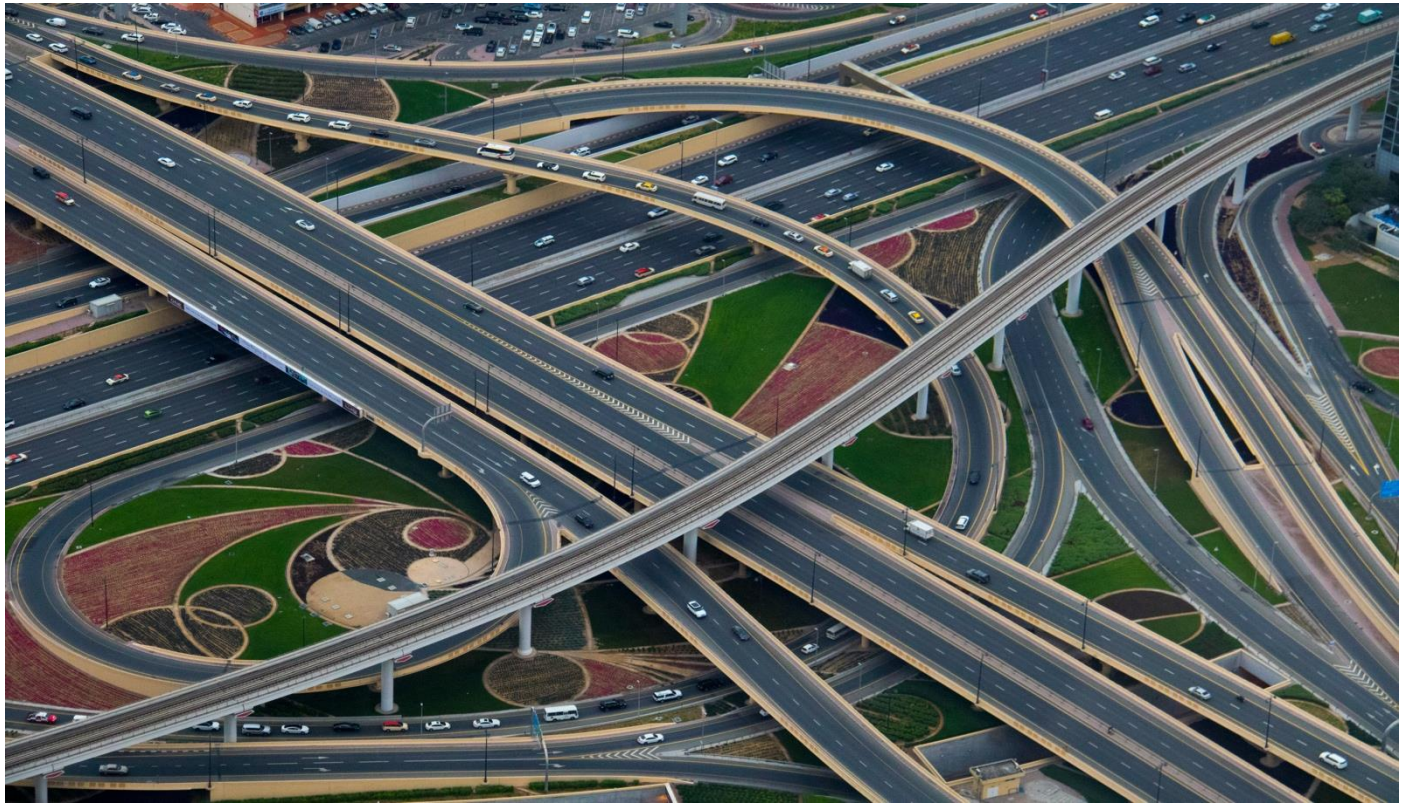

IGP of your choice that gets information from Source RLOC to Destination RLOC, the best way it can.

BGP VRF-Lite for external communications

What is an SD-Access Fabric Site?

SD-Access Fabric Site offers programmable overlays for wired and wireless campus networks, enabled on a single physical infrastructure.

A single fabric site could be demarcated and defined based upon:

- Collection of Edge Nodes, Border/CPs, and optionally Wireless LAN Controllers/Access Points.
- Geographical location.
- Required scale, network devices.
- Failure domain scoping.
- RTT.
- Underlay connectivity attributes.

Multiple fabric sites interconnected by a “Transit”.

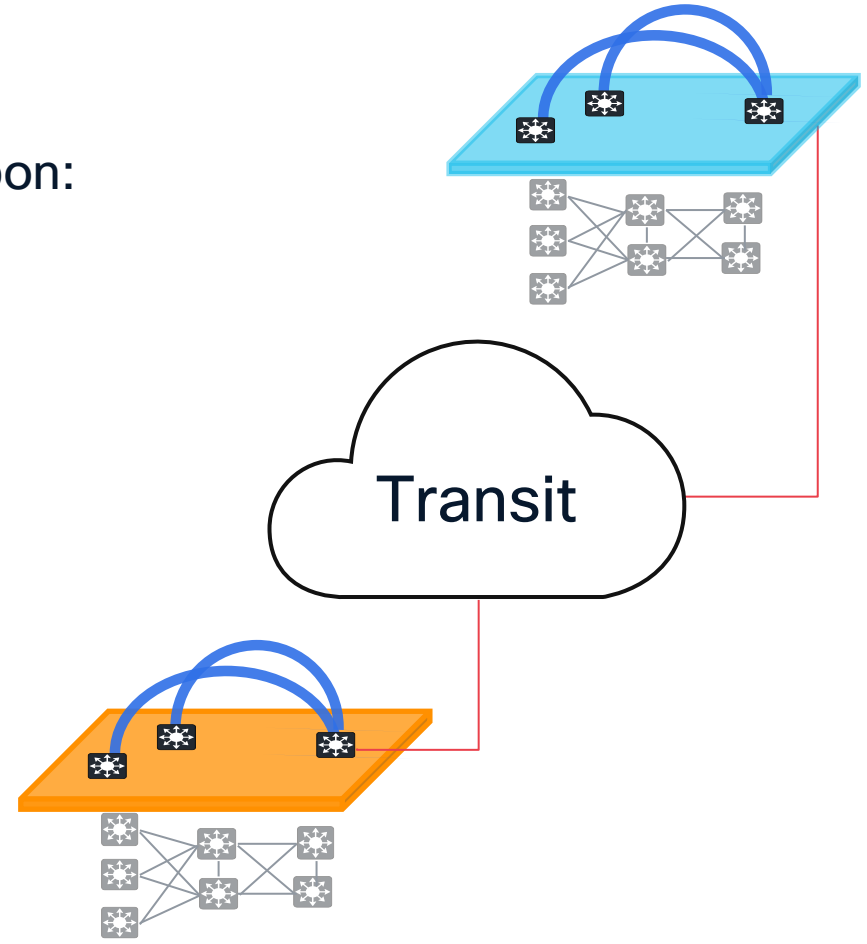

Roles and Terminology

1. Concepts
2. SD-Access Roles
3. Fabric Constructs

Cisco SD-Access Roles

Key Roles for a Complete Wired and Wireless SDA Fabric Experience

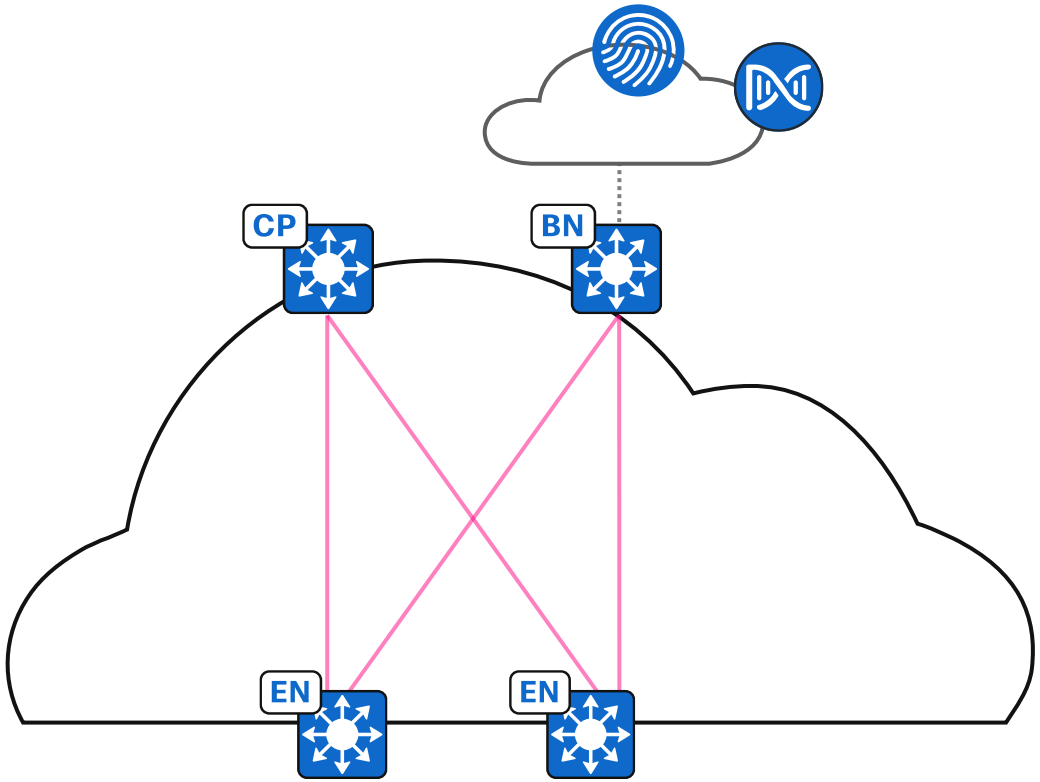

Cisco Catalyst Center

GUI and APIs for intent-based automation of wired and wireless fabric devices.

Identity Service Engine

NAC and ID services for dynamic endpoint to Security Group Tag mapping and policy distribution.

Control Plane Node

Map System that tracks endpoint to fabric node relationships.

Border Nodes

Connects external L3 and L2 networks to the Cisco SD-Access fabric.

Edge Nodes

Connects wired endpoints to the Cisco SD-Access fabric and optionally enforces micro-segmentation policy.

Cisco SD-Access Roles

Key Roles for a Complete Wired and Wireless Campus Experience

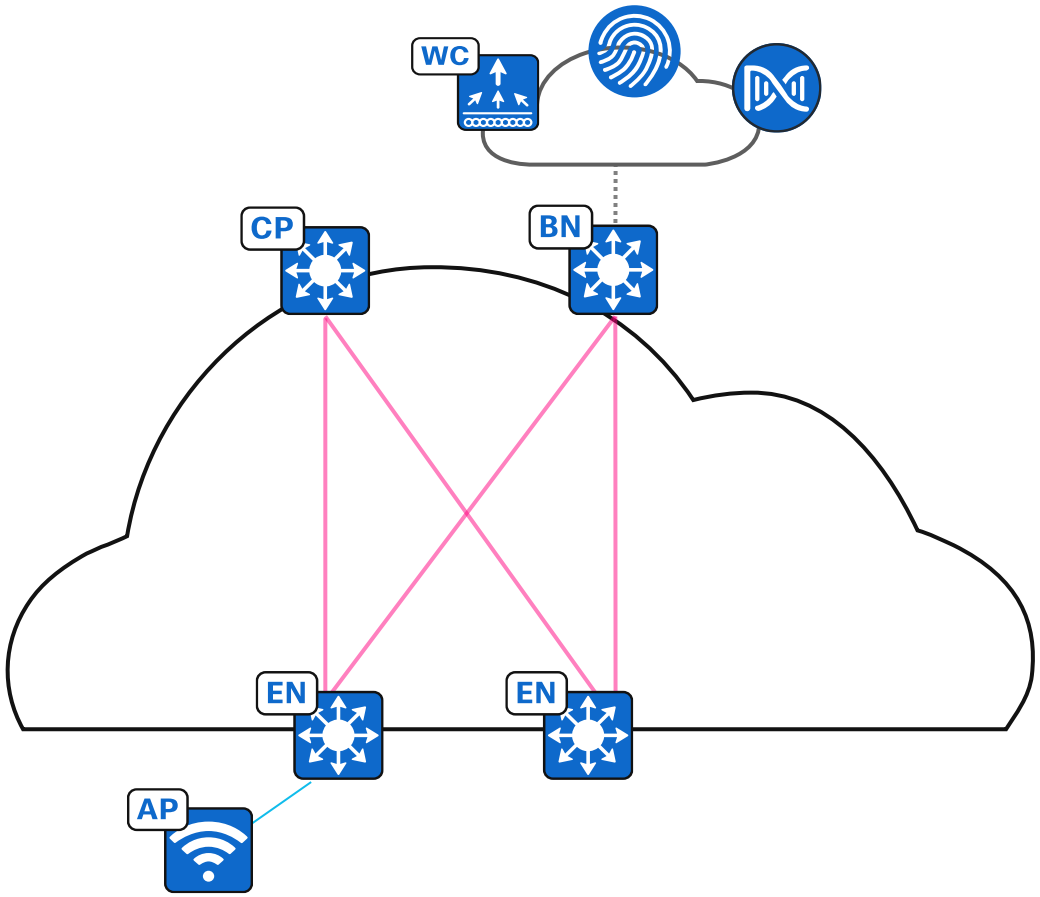

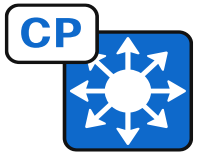

Control Plane Node

Map System that tracks endpoint to fabric node relationships.

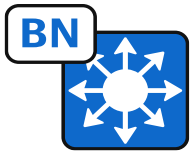

Border Nodes

Connects external L3 and L2 networks to the Cisco SD-Access fabric.

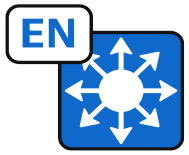

Edge Nodes

Connects wired endpoints and Fabric APs to the Cisco SD-Access fabric and optionally enforces micro-segmentation policy.

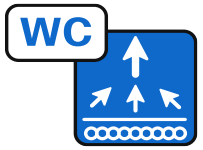

Fabric Wireless Controller

Fabric WLC is integrated into the SD-Access Control Plane (LISP) communication.

Fabric Access Point

Switches endpoint traffic to the adjacent Edge Node.

Cisco SD-Access Roles

Additional Roles for Reference

Additional Information

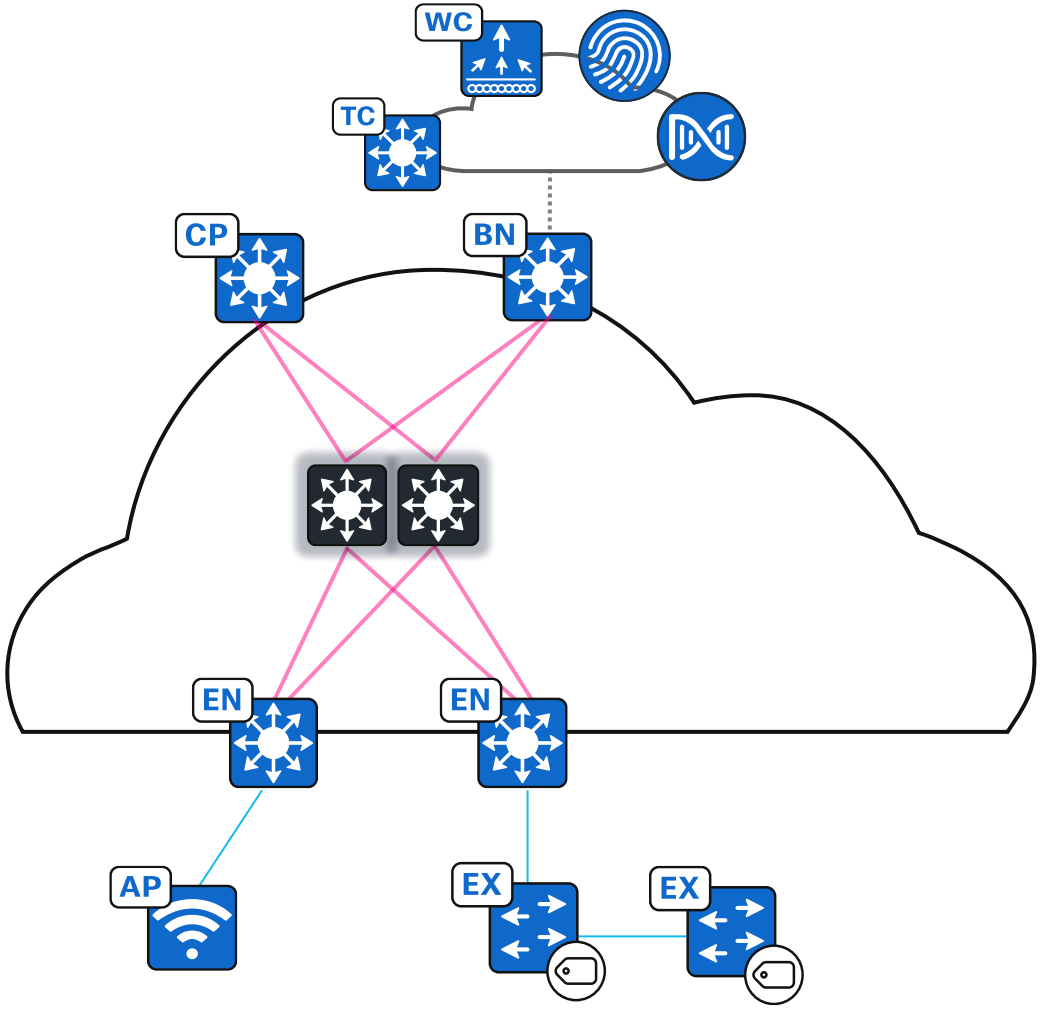

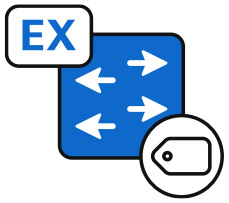

Extended Nodes

A switch operating at Layer 2 that extends fabric connectivity and optionally enforces micro-segmentation policy.

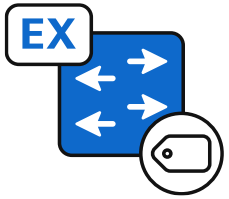

Policy Extended Nodes

Switch able to do Auth, VLAN & SGT assignment and policy enforcement at the edge, without VXLAN tunnelling.

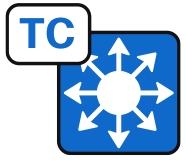

Transit Control Plane Nodes

Facilitates connectivity of multiple SD-Access fabric sites while preserving end to end segmentation.

Intermediate Nodes

Moves data between fabric nodes. Can be one or many hops. Part of the underlay.

Cisco SD-Access Roles

Some of the Supported Colocations

Additional Information

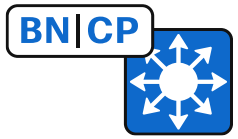

Border Node and Control Plane Node.

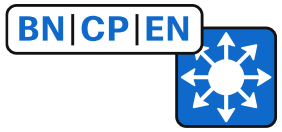

Border Node, Control Plane Node, and Fabric Edge Node.

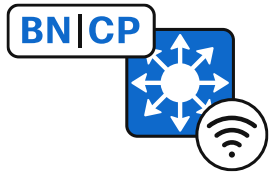

Border Node, Control Plane Node, and Embedded Wireless Controller.

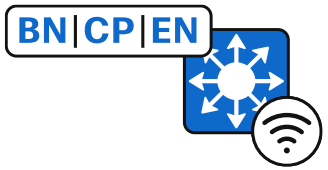

Border Node, Control Plane Node, Fabric Edge Node, and Embedded Wireless Controller.

Cisco SD-Access Fabric

Control Plane Node Maintains a Host and Network Tracking Database

A simple Host Database that maps Endpoint IDs to locations, along with other attributes.

Host Database supports multiple types of Endpoint ID lookup types (IPv4, IPv6 or MAC).

Receives Endpoint ID map registrations from Edge Nodes, Border Nodes and Fabric Wireless LAN Controllers.

Publishes registrations to Subscribers (Border Nodes).

Resolves lookup requests from Edge Nodes and Border Nodes, to locate destination Endpoint IDs.

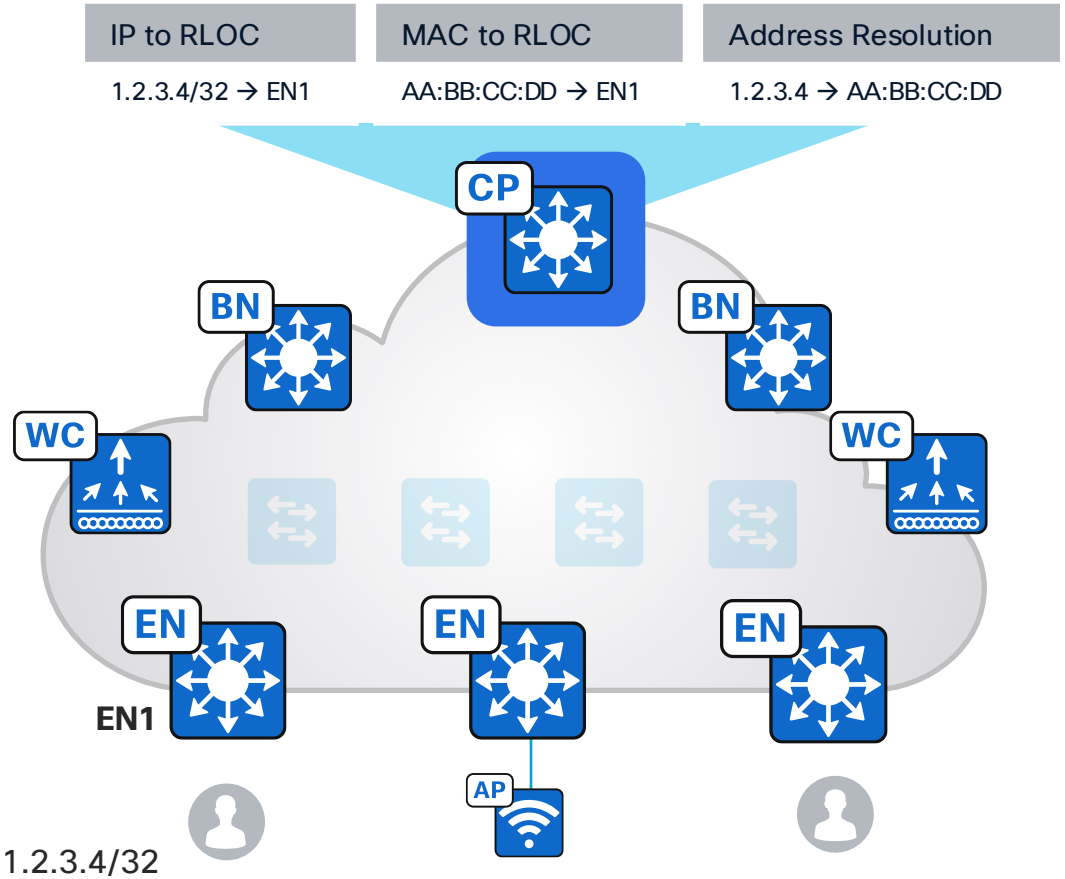

Cisco SD-Access Fabric

Edge Node Provides First Hop Services for Endpoints

Responsible for Authenticating and Authorizing wired endpoints (802.1x, MAB, static) in concert with ISE.

Register Endpoint IDs (IPv4, IPv6, MAC) with the Control Plane Nodes.

Provide an Anycast Gateway for the connected wired and wireless endpoints.

Performs VXLAN encapsulation and decapsulation of traffic to and from all connected wired endpoints.

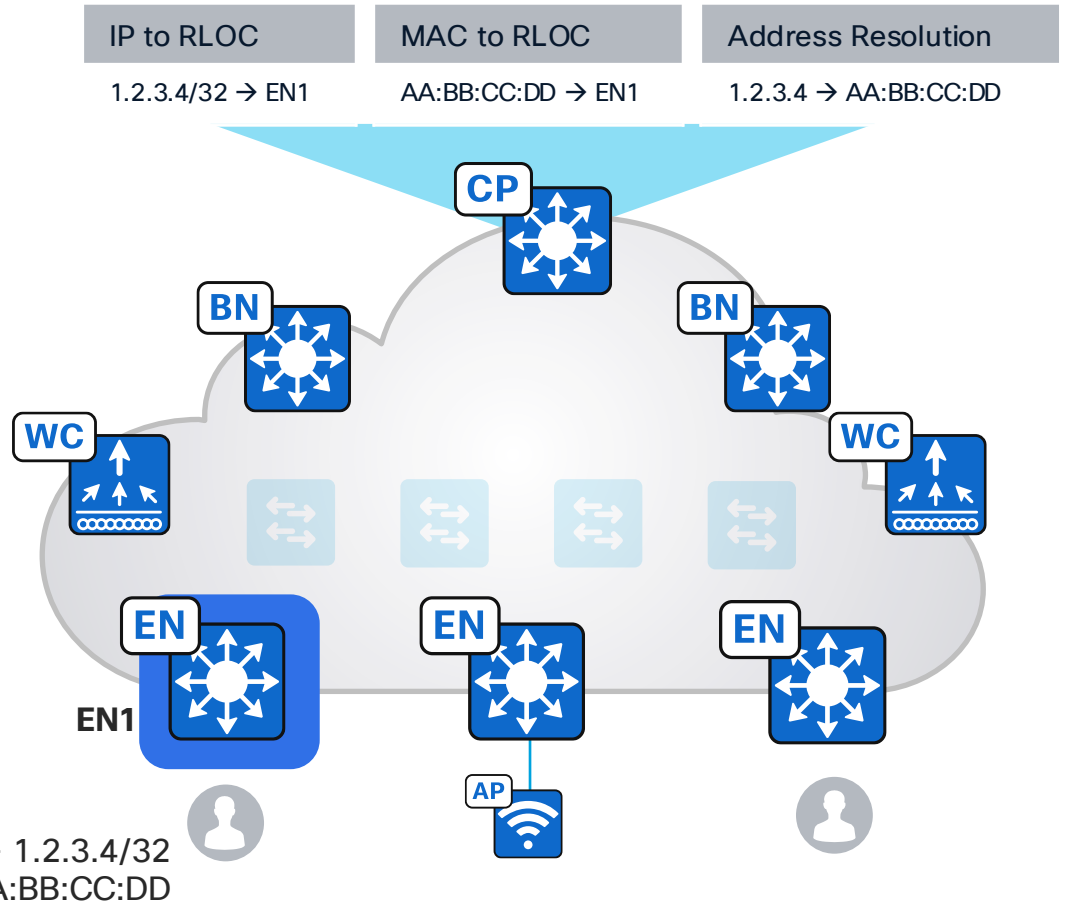

Cisco SD-Access Fabric

Border Node is the Fabric Site Entry and Exit for Network Traffic

Subscribes to LISP Control Plane Node IPv4 and IPv6 Tables.

There are 4 types of Border Node:

- External Border Node.
- Internal Border Node.
- Internal + External Border Node.
- Layer 2 Border Node.

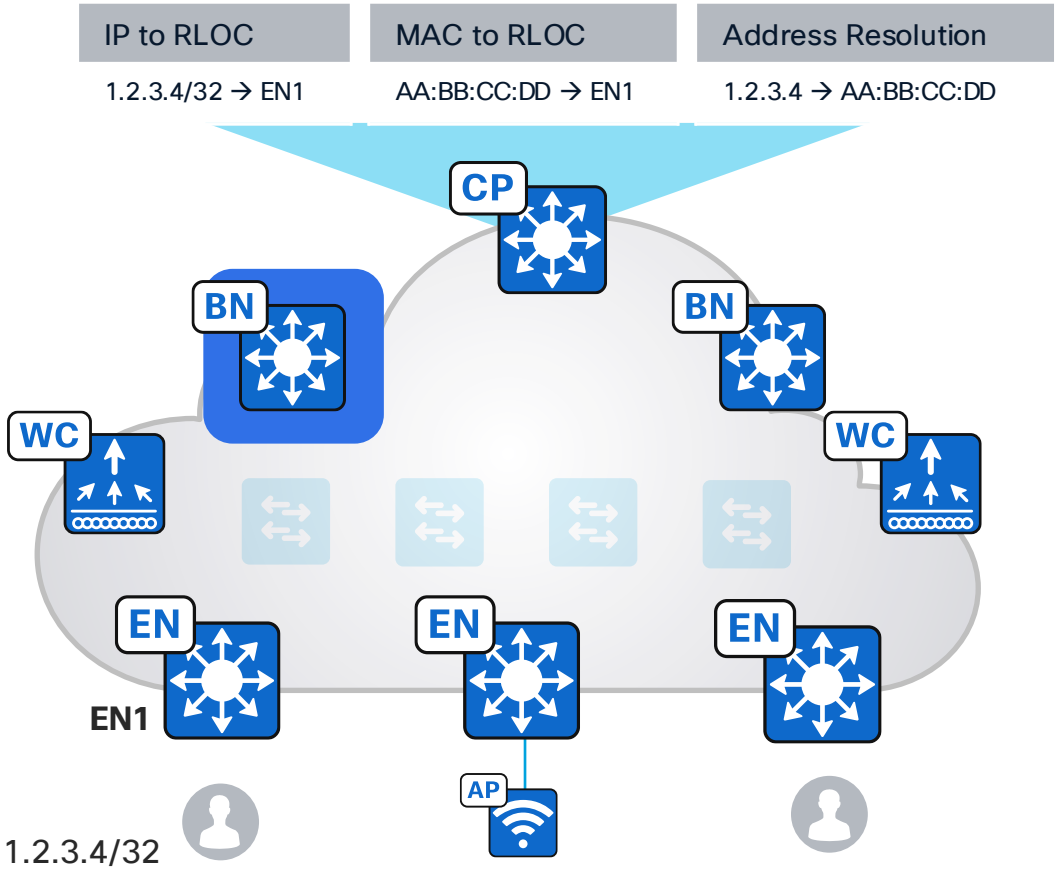

Cisco SD-Access Fabric

Border Node is the Fabric Site Entry and Exit for Network Traffic

External Border Node:

- The most common configuration.
- Exports all fabric subnets to outside the Fabric Site as eBGP summary routes.
- Acts as a gateway of last resort for the Fabric Site.
- Does not register eBGP prefixes from outside the Fabric Site into the fabric Control Plane.

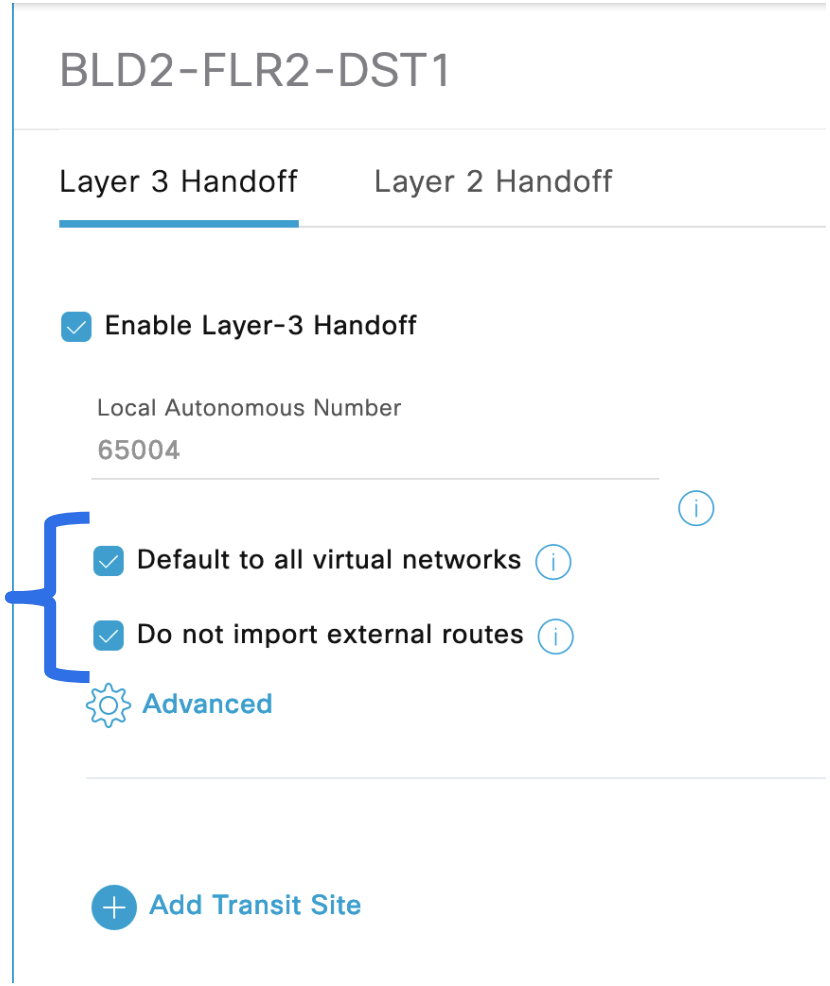

Cisco SD-Access Fabric

Border Node is the Fabric Site Entry and Exit for Network Traffic

Internal Border Node:

Exports all fabric subnets to outside the Fabric Site as eBGP summary routes.

Imports and registers eBGP-learned IPv4/IPv6 prefixes from outside the Fabric Site, into the fabric Control Plane.

Does not act as a gateway of last resort for the Fabric Site.

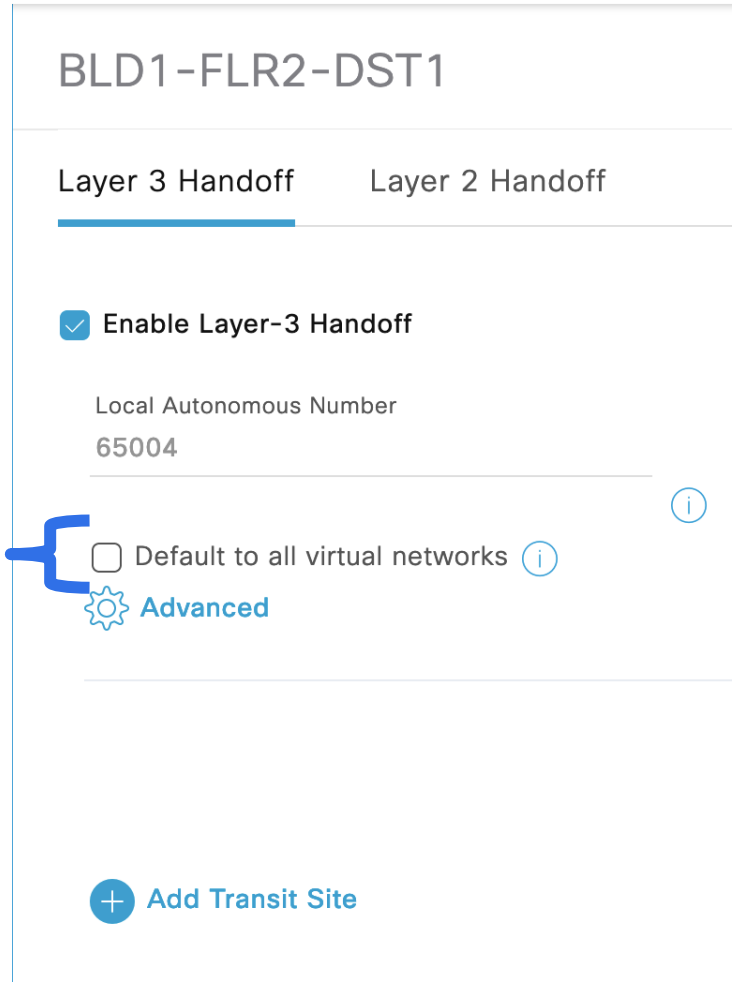

Cisco SD-Access Fabric

Border Node is the Fabric Site Entry and Exit for Network Traffic

Internal + External Border Node:

- Exports all fabric subnets to outside the Fabric Site as eBGP summary routes.
- Imports and registers eBGP-learned IPv4/IPv6 prefixes from outside the Fabric Site, into the fabric Control Plane.
- Acts as a gateway of last resort for the Fabric Site.

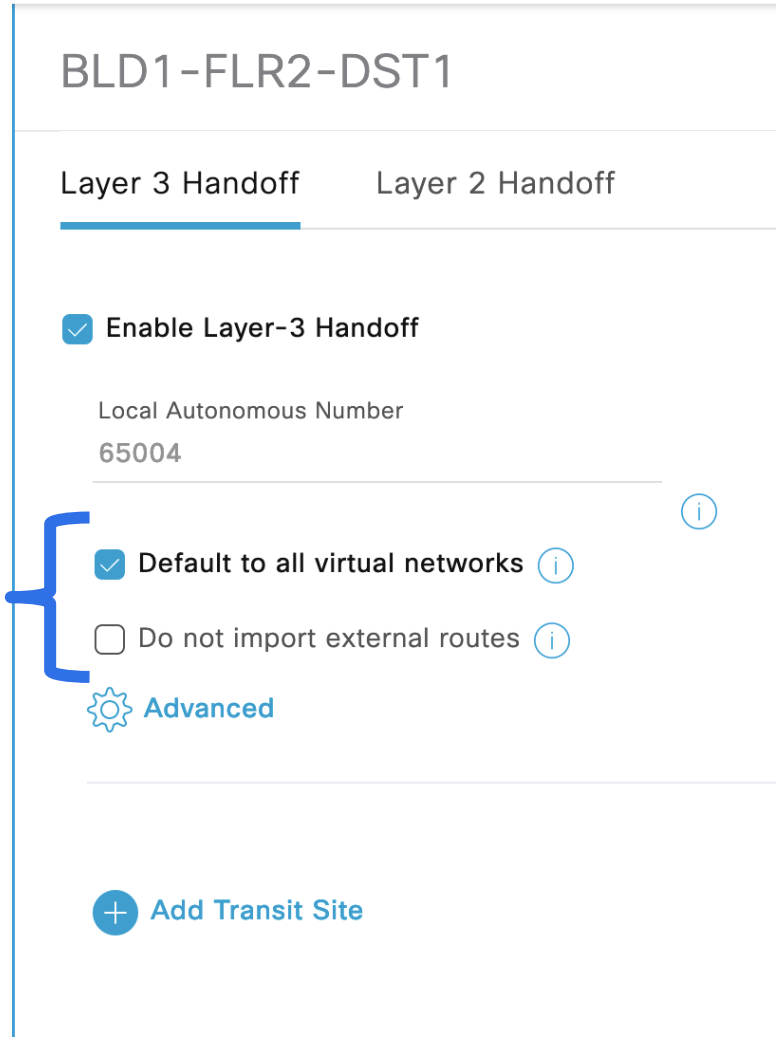

Cisco SD-Access Fabric

Border Node is the Fabric Site Entry and Exit for Network Traffic

Layer 2 Border Node:

- Acts as Layer 2 handoff for pure Layer 2 Overlays or Layer 2 + Layer 3 Overlays.
- Allows VLAN translation between SD-Access network segments and non-fabric VLAN IDs.
- Dual homing requires link aggregation; STP it not tunneled within the SD-Access Fabric.
- Ideally should be separate device from the Layer 3 Border Node.

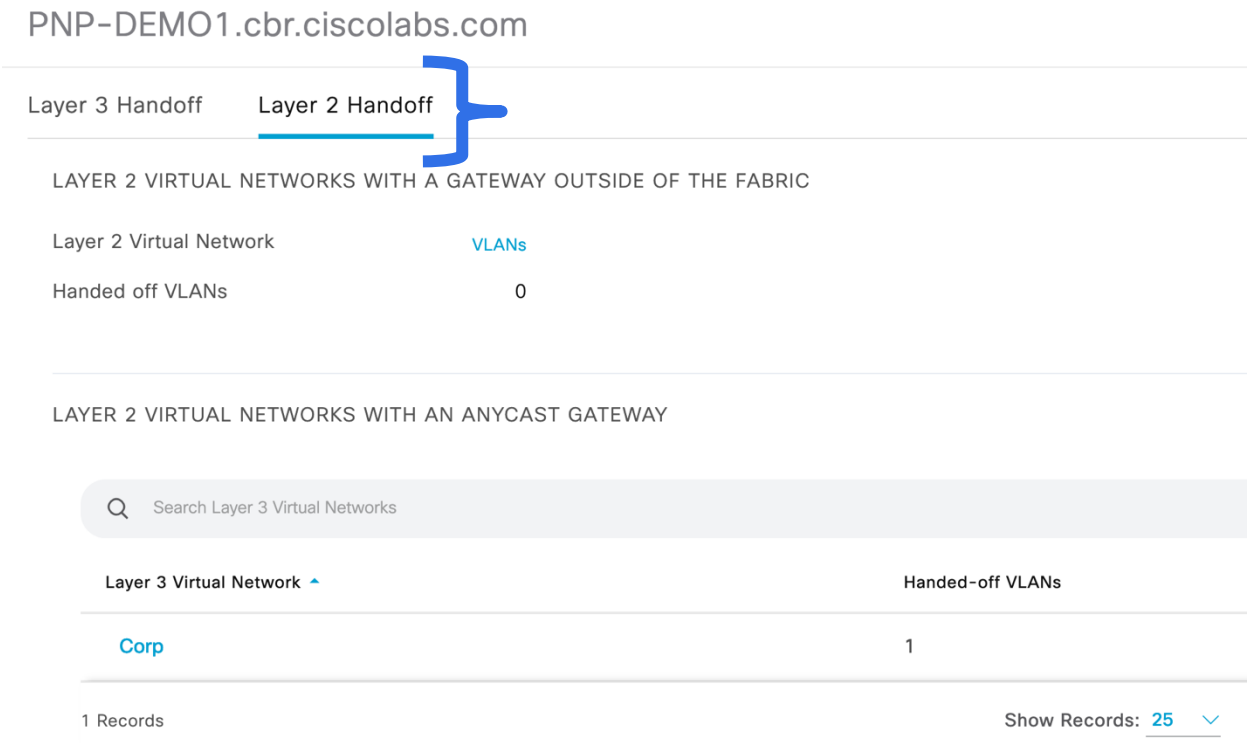

Cisco SD-Access Fabric

Fabric Enabled Wireless for Unified Management, Policy and Data Planes

Fabric WLC accessible though a Fabric Border Node (Underlay). Can be several hops away.

Fabric Enabled APs reside in a dedicated IP range and communicate with the Fabric WLC (CAPWAP Control).

Fabric WLC registers endpoints with the Control Plane Node.

Fabric APs switch endpoint traffic to the adjacent Edge Node.

Wireless endpoints use same data plane and policy plane as wired endpoints.

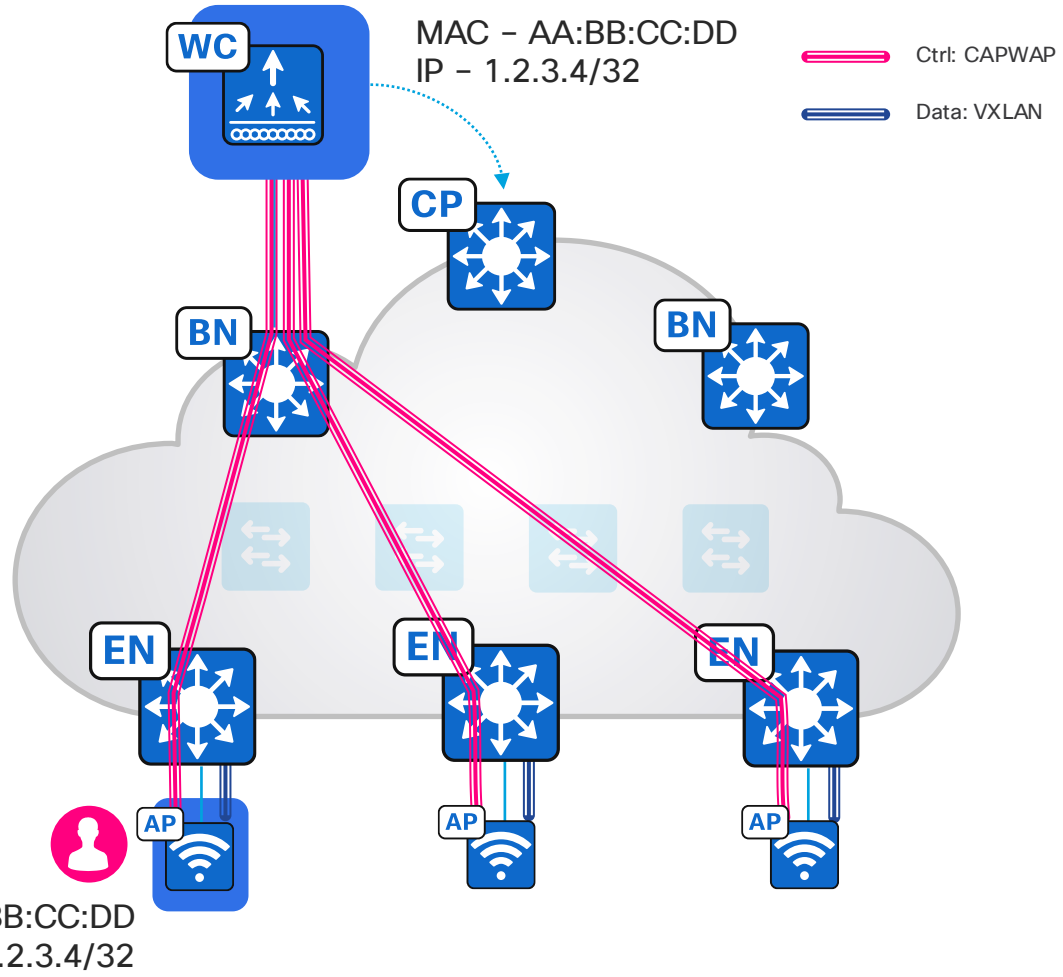

Roles and Terminology

1. Concepts
2. SD-Access Roles
3. **Fabric Constructs**

Cisco SD-Access Fabric

Virtual Networks

Additional Information

- Layer 3 Virtual Networks use VRFs and LISP Instance IDs to maintain separate routing topologies.
  - Endpoint IDs (IPv4/IPv6 addresses) are routed within an L3VN.
- Layer 2 Virtual Networks use LISP Instance IDs and VLANs to maintain separate switching topologies.
  - Endpoint IDs (MAC addresses) are switched within an L2VN.
- Edge Nodes, Border Nodes and Fabric APs add a VNID (the LISP IID) to the fabric encapsulation.

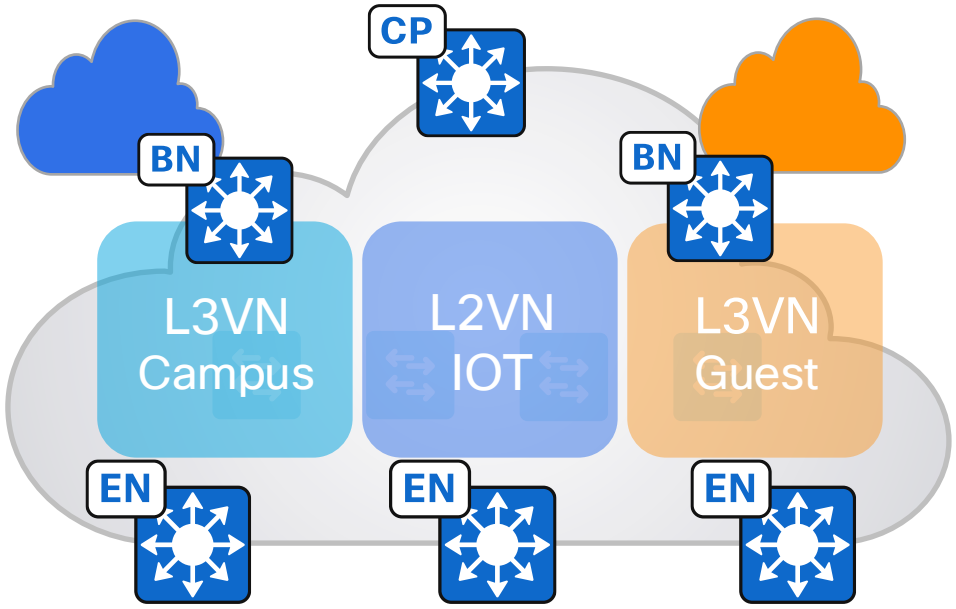

Layer 3 Virtual Networks

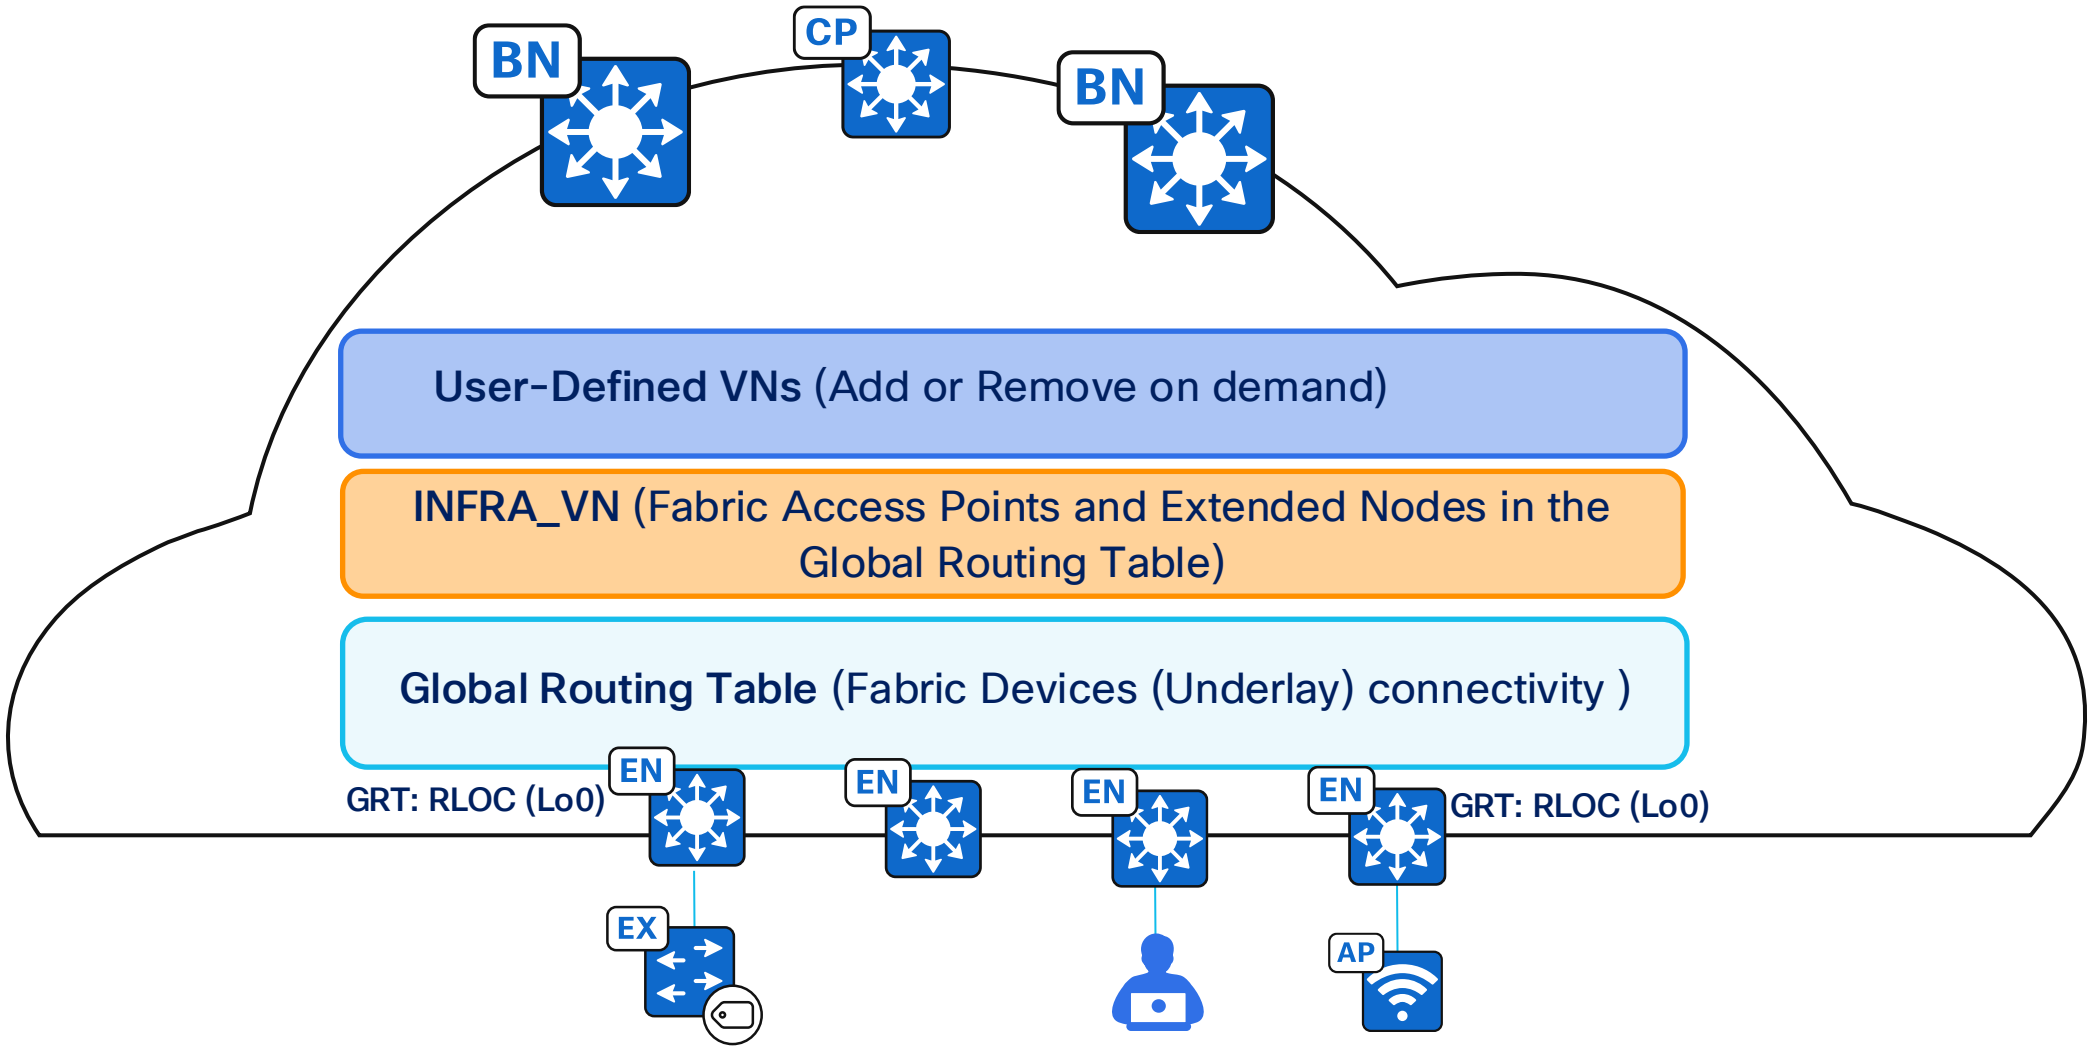

Layer 3 Handoff

Per-Layer-3-Virtual-Network Layer 3 Handoff using Peer Device

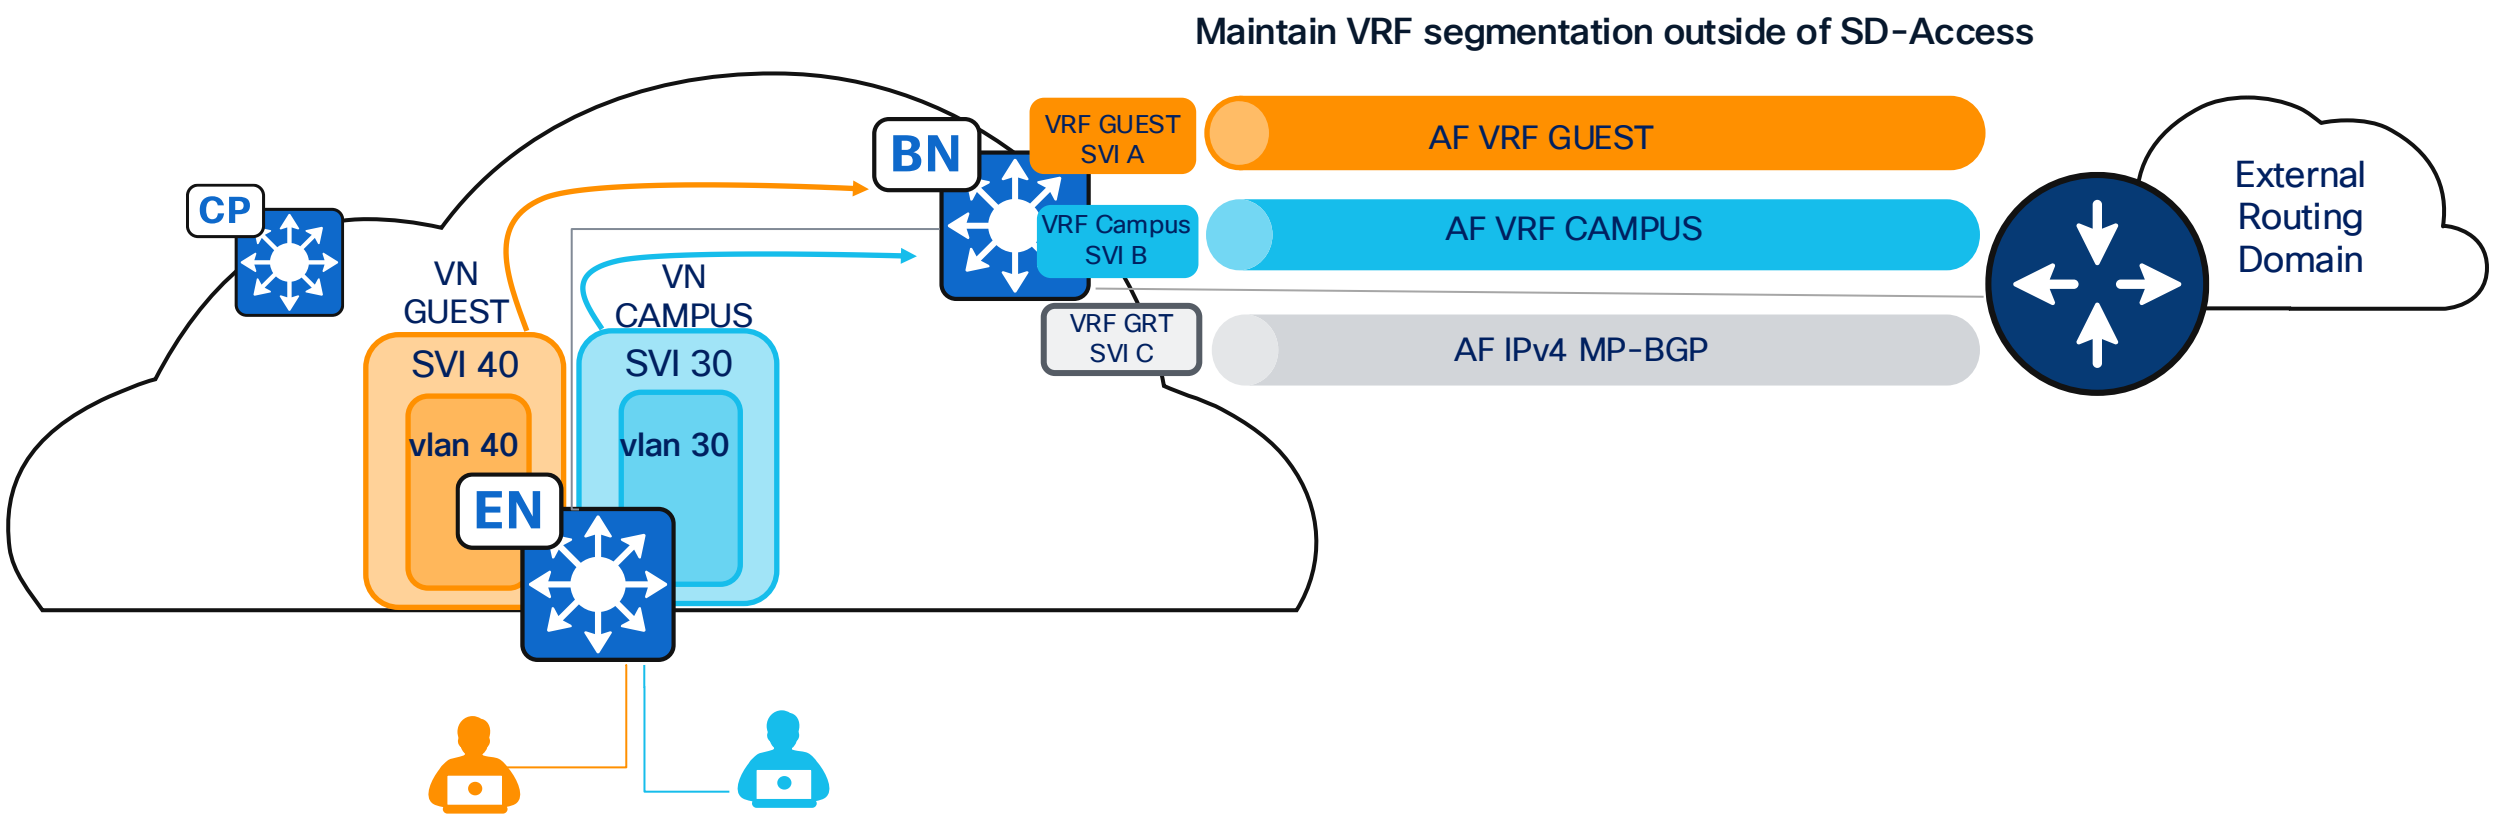

Extranet Layer 3 Handoff

Helps achieve route-leaking natively in LISP SD-Access Fabric

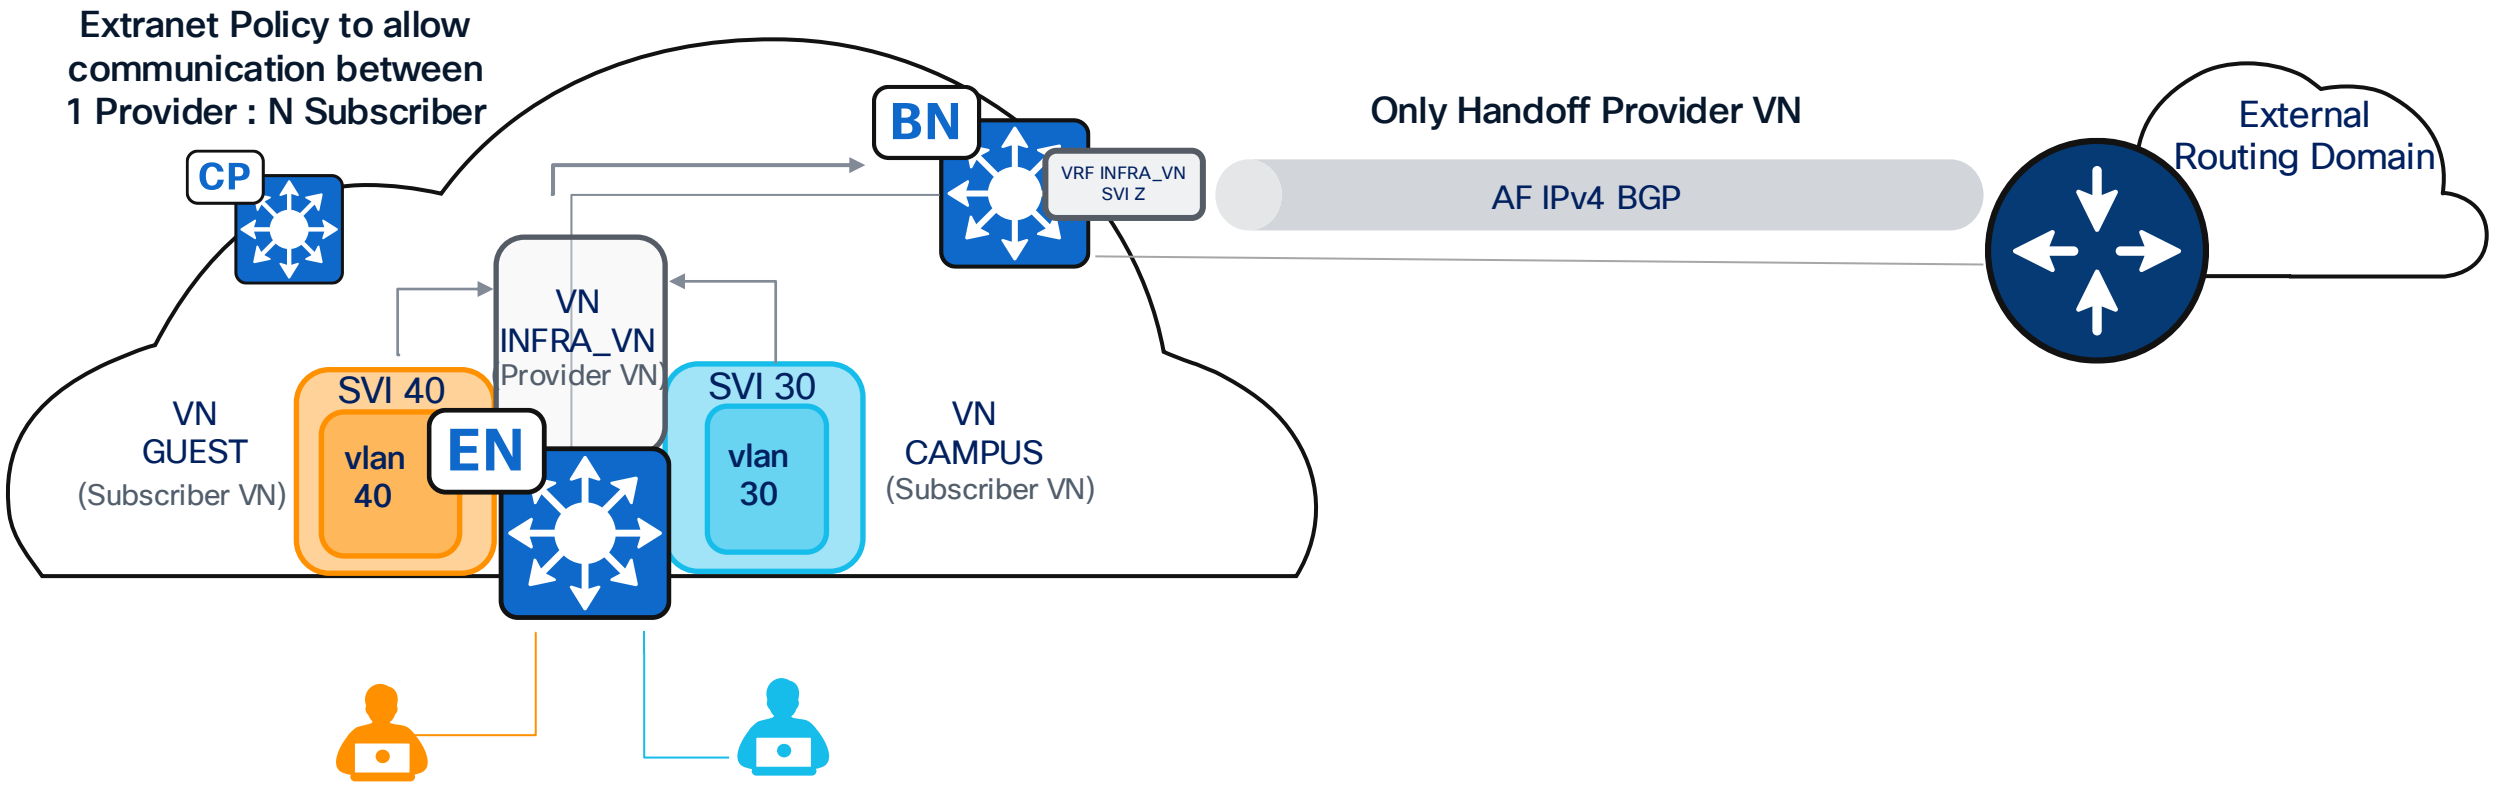

Layer 2 Handoff

Layer 2 Virtual Networks handoff through a user-defined VLAN

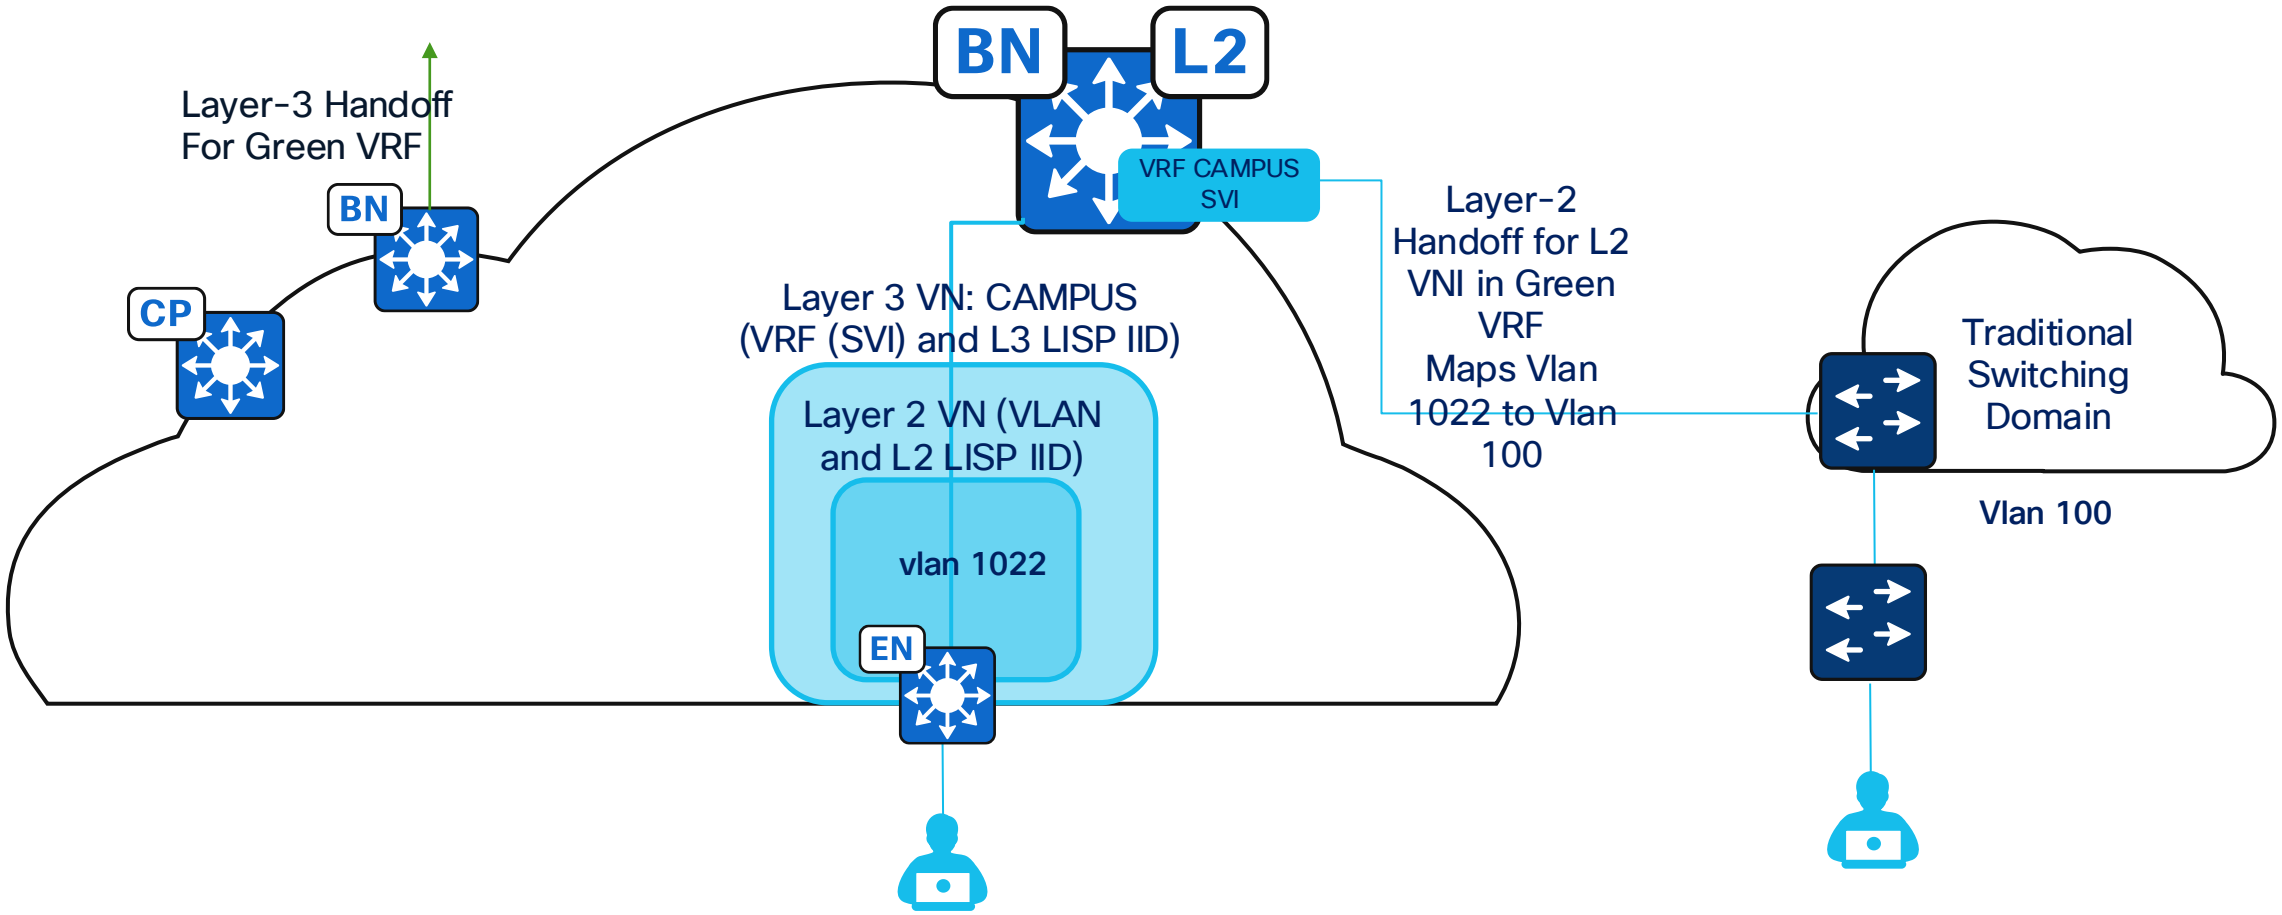

Security Group Tag

A Security Group Tag Assigns a “Group” to Each Endpoint

- Edge Nodes and Fabric APs assign a unique Scalable Group Tag (SGT) to each end endpoint in concert with ISE.
- Edge Nodes and Fabric APs add an SGT to the fabric encapsulation.
- SGTs are used to implement IP-address-independent access policies.
- SGTs can be extended to numerous other networking technologies e.g., Cisco Secure Firewall, Cisco SD-WAN, some third-party devices, etc.

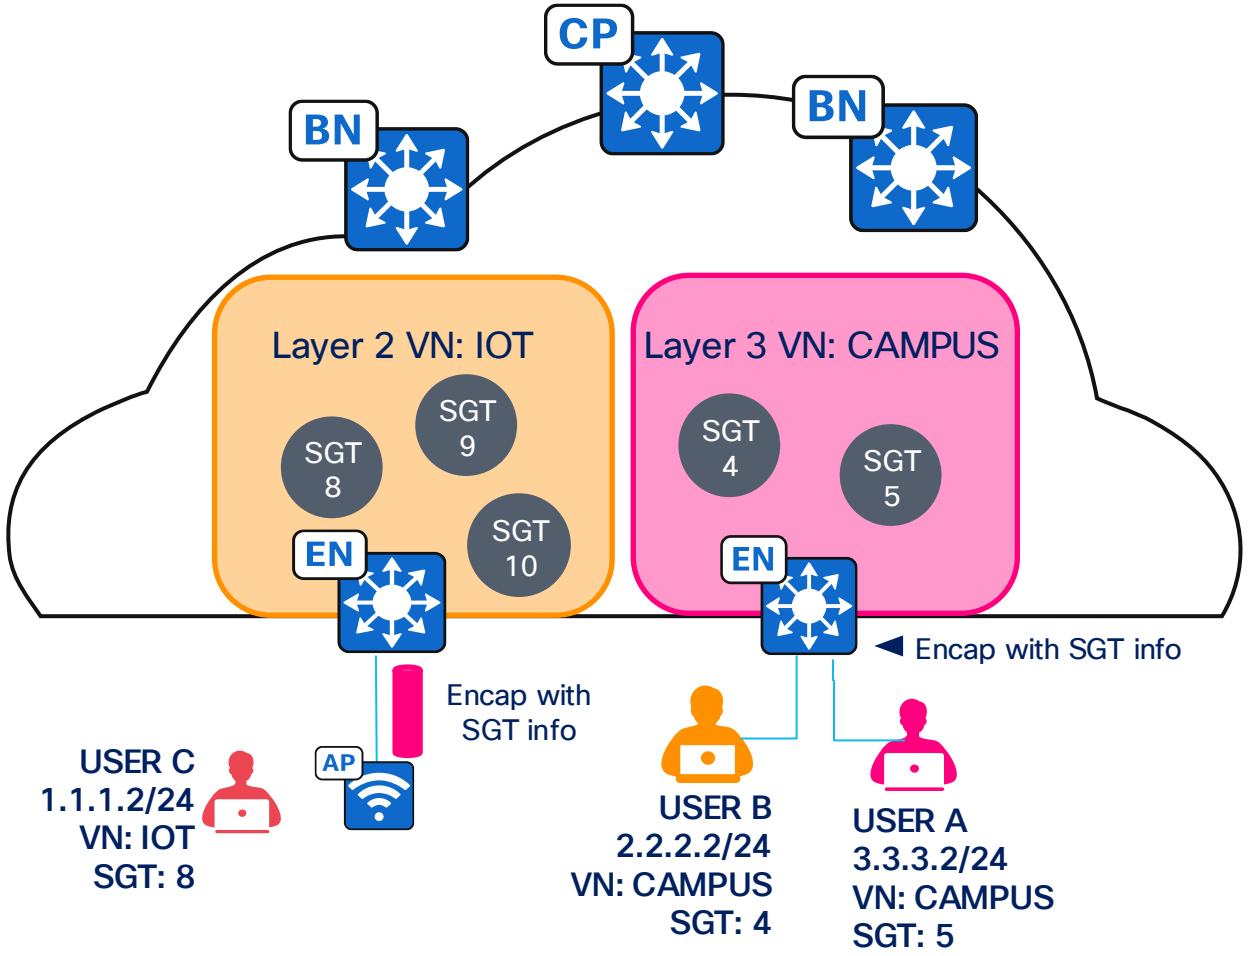

Cisco SD-Access Fabric

Host Pools Provide a Default Gateway and Basic IP Services for Endpoints

- Edge Nodes instantiate an access VLAN and a Switched Virtual Interface (SVI) with user-defined IPv4/IPv6 addresses per Host Pool.
- Host Pools assigned to endpoints dynamically by AAA or statically per port.
- Edge Nodes and Fabric WLCs register endpoint IDs (/32, /128 or MAC) with the Control Plane, enabling IP mobility; any IP address anywhere.

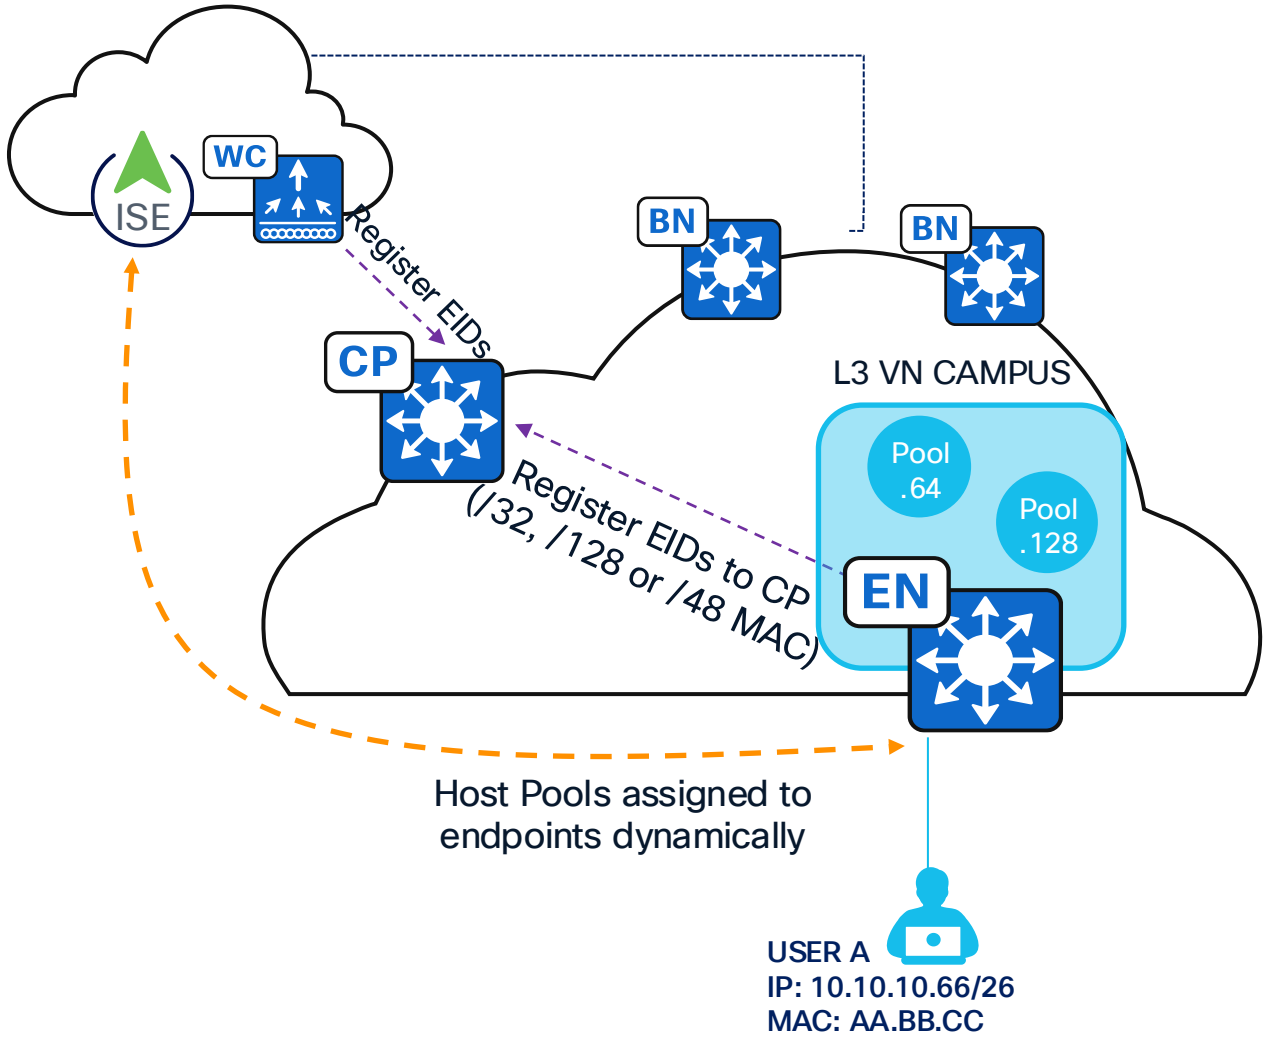

Cisco SD-Access Fabric

Anycast Gateway Provides a Default Gateway for IP-Capable Endpoints

- Similar principle and behavior to FHRP with a shared virtual IPv4/IPv6 addresses and MAC address.
- The same Switch Virtual Interface (SVI) is present on all Edge Nodes with the same virtual IP and MAC.
- The wired or wireless endpoint can connect to any switch or AP in the fabric and communicate with the same Anycast Gateway.

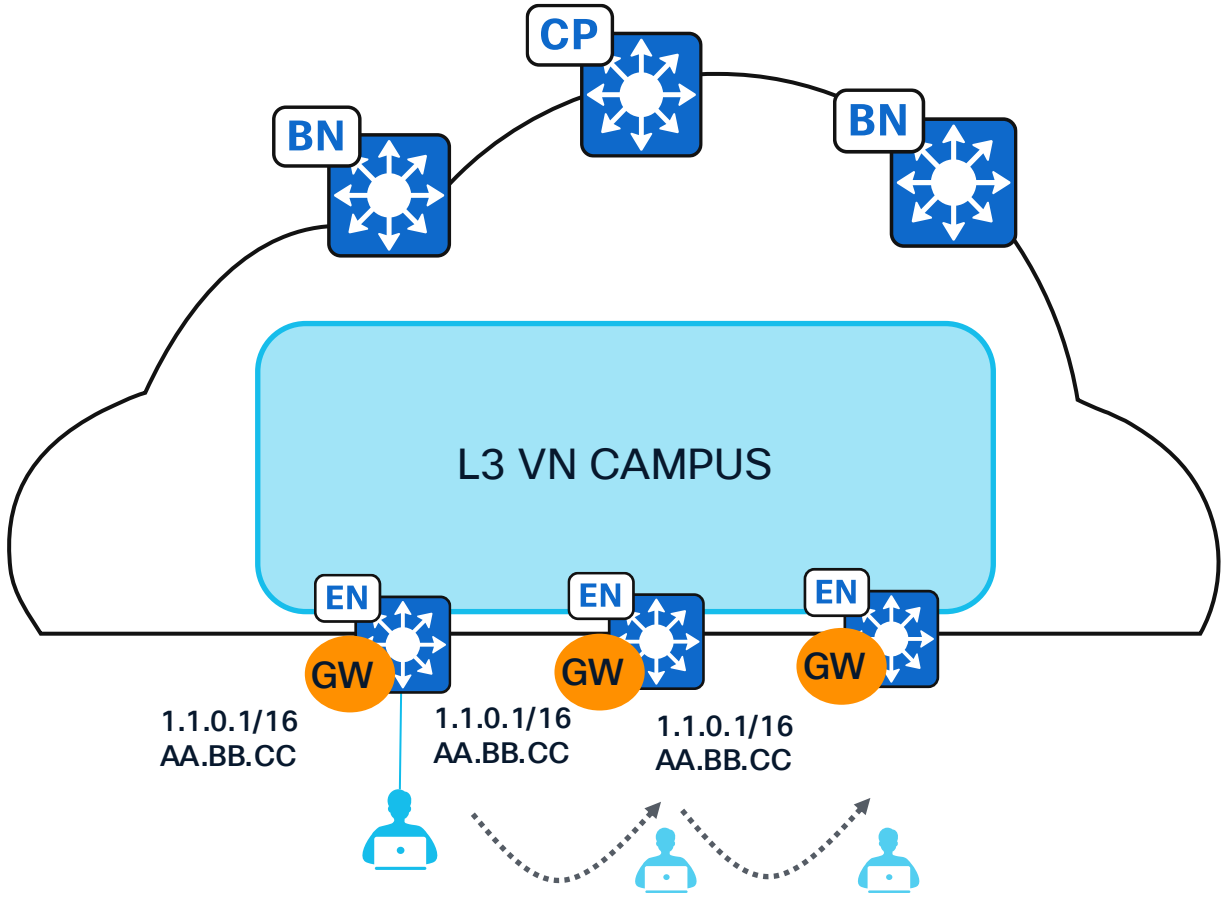

Cisco SD-Access Fabric

Host Pools are “stretched” via the Overlay

- Endpoint IPv4/IPv6 traffic arrives on an Edge Node and is then routed or switched by the Edge Node.
- Fabric Dynamic EID mapping allows endpoint-specific (/32, /128, MAC) advertisement and mobility.
- VLANs to connect endpoints across Edge Nodes, this happens in the Overlay without broadcast flooding.

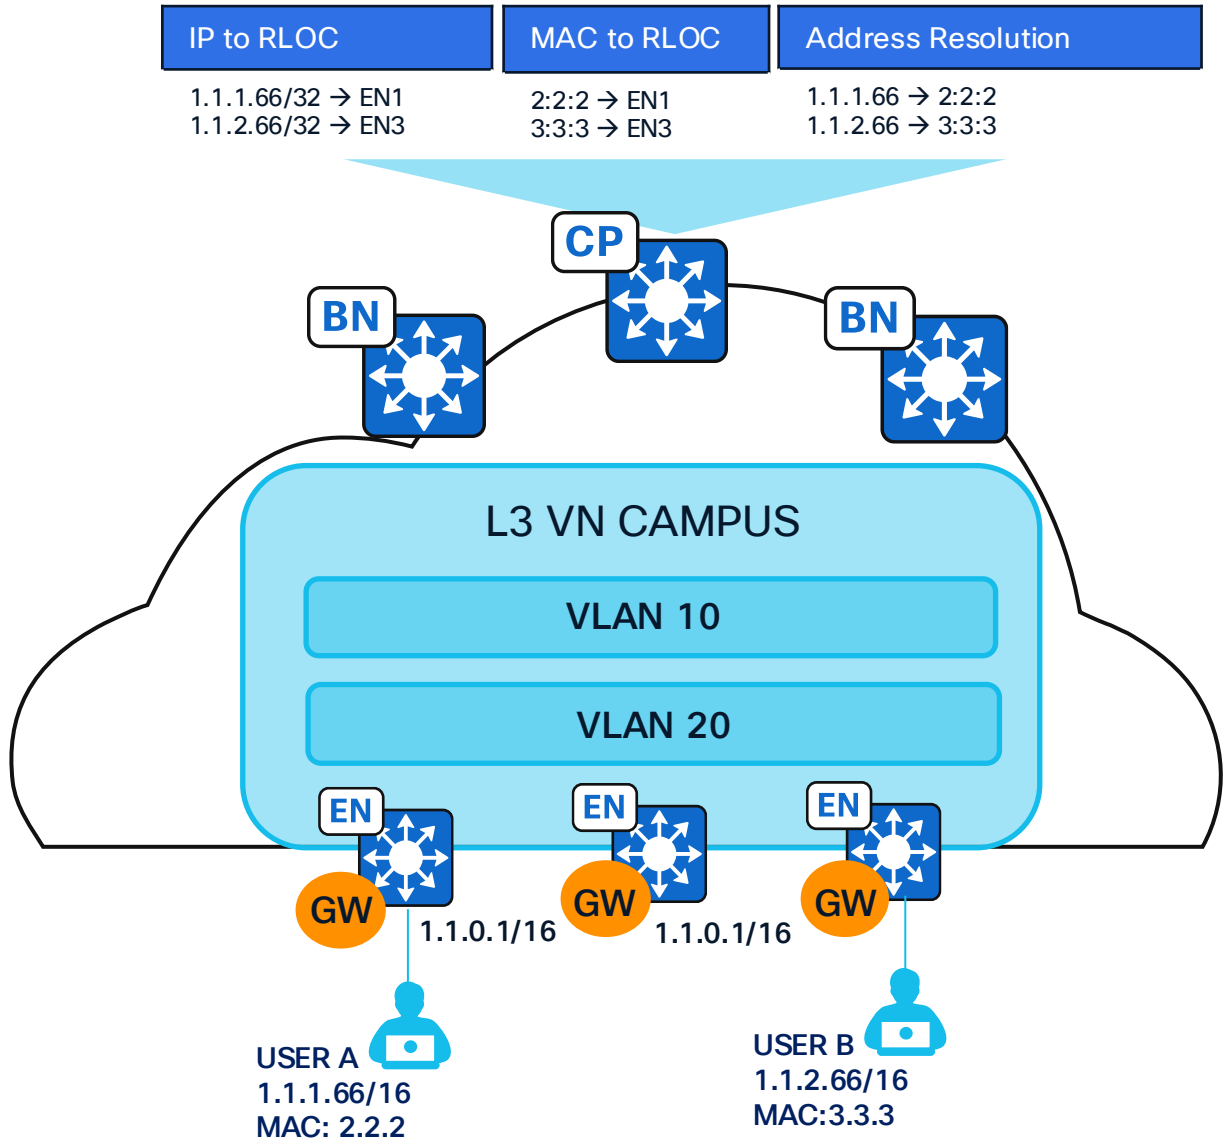

Layer 2 Virtual Networks

Fabric is just a transport for Layer-2 adjacencies

- By default, an L2VN is deployed with each Anycast Gateway and Layer 2 Flooding is disabled. Layer 2 Flooding can be enabled, if necessary, to service niche applications.
- L2VN can be deployed without an Anycast Gateway, and Layer 2 Flooding cannot be disabled.
- Sometimes referred to as “Gateway Outside the Fabric”.
- If Layer 2 Flooding is enabled, a Multicast Underlay P2MP tunnel needs to be configured between all Fabric Nodes.

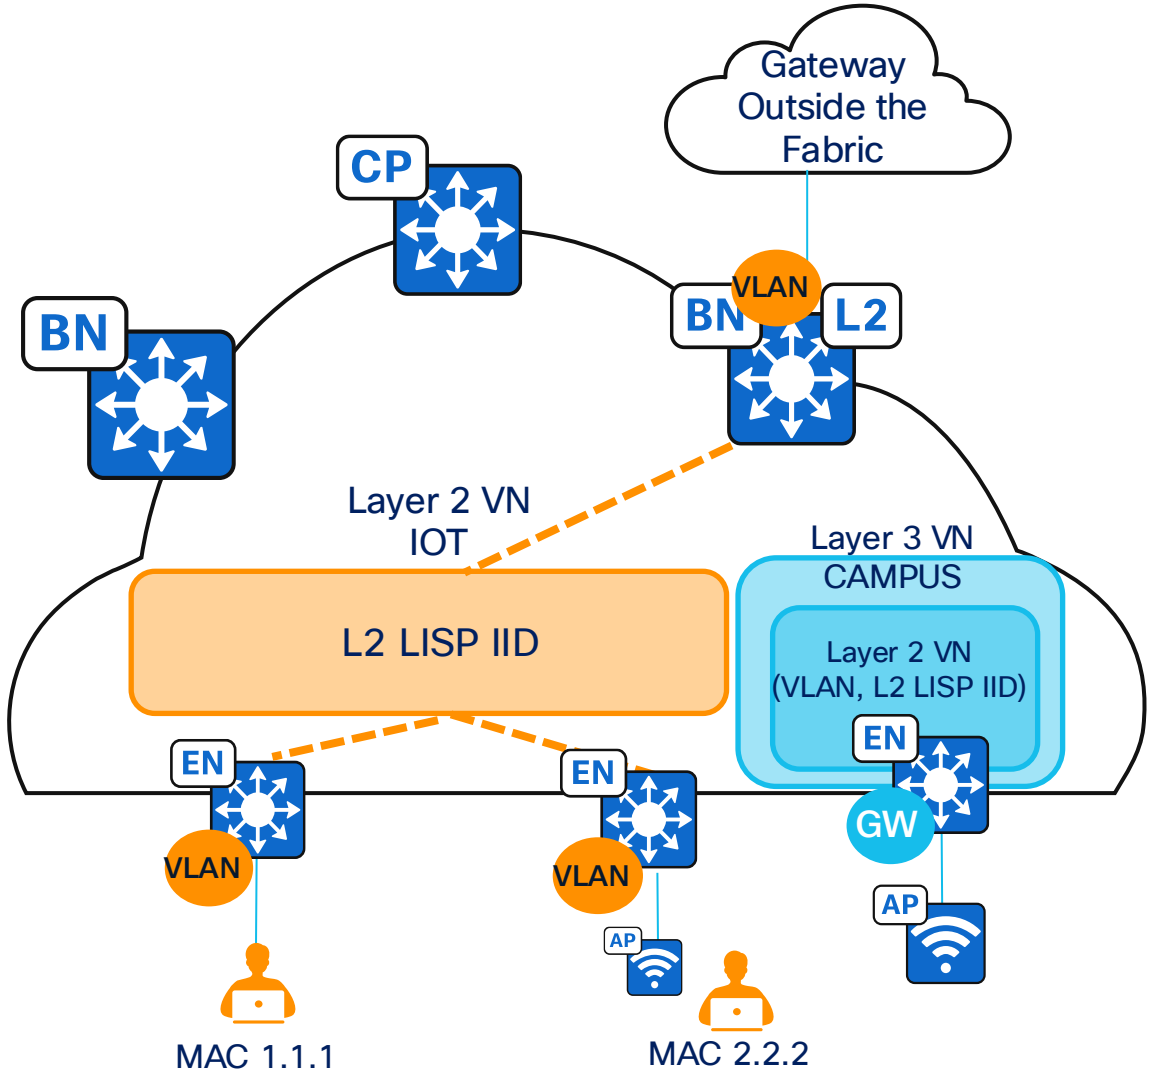

Fabric Fundamentals

Control Plane

Cisco SD-Access Fabric

Control Plane: Locator/ID Separation Protocol (LISP)

Where you are in a network can change, but ***who*** you are in the network remains the same.

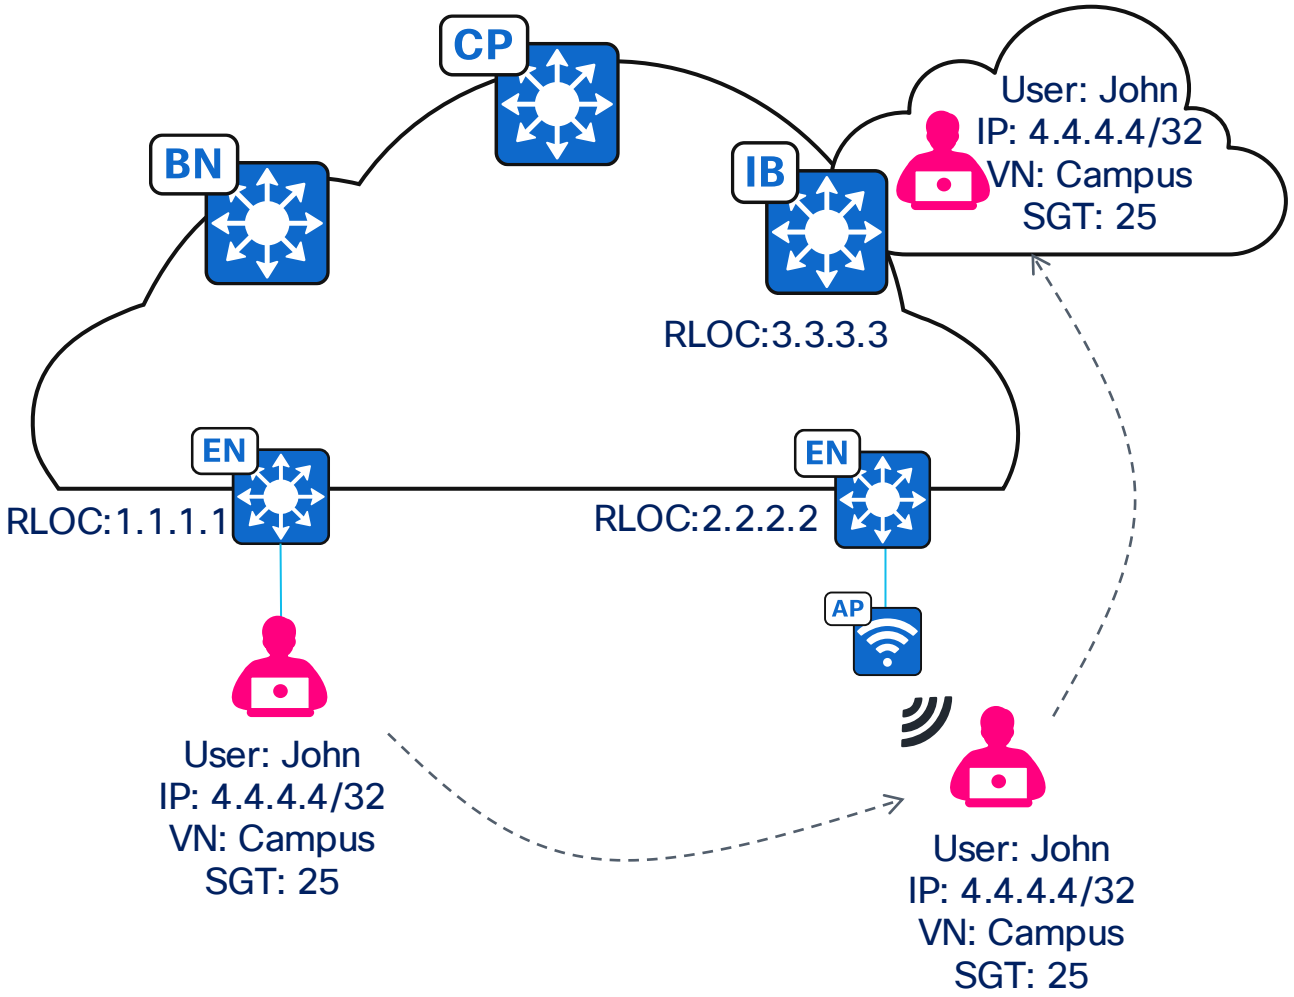

(**IETF Standards** Track RFC9300-RFC9306 and Informational RFC9299)

LISP Operations

LISP :: Mapping Resolution “Level of Indirection”

- LISP “Level of Indirection” is analogous to a DNS lookup
  - DNS resolves **IP addresses** for **URL** **Answering the “WHO IS” question**

DNS
Name-to-IP
URL Resolution

  - LISP resolves **locators** for queried **identities** **Answering the “WHERE IS” question**

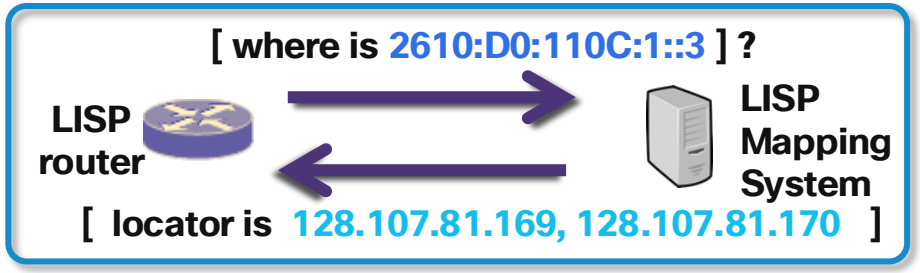

LISP
Identity-to-locator
Mapping Resolution

Fundamental Design Principle in LISP

A key basic design objective:
Distribute routing/mapping information **only where it is required**

A basic working principle:
*Use traffic signals to **pull** routes* **when required**

LISP in Cisco SD-Access

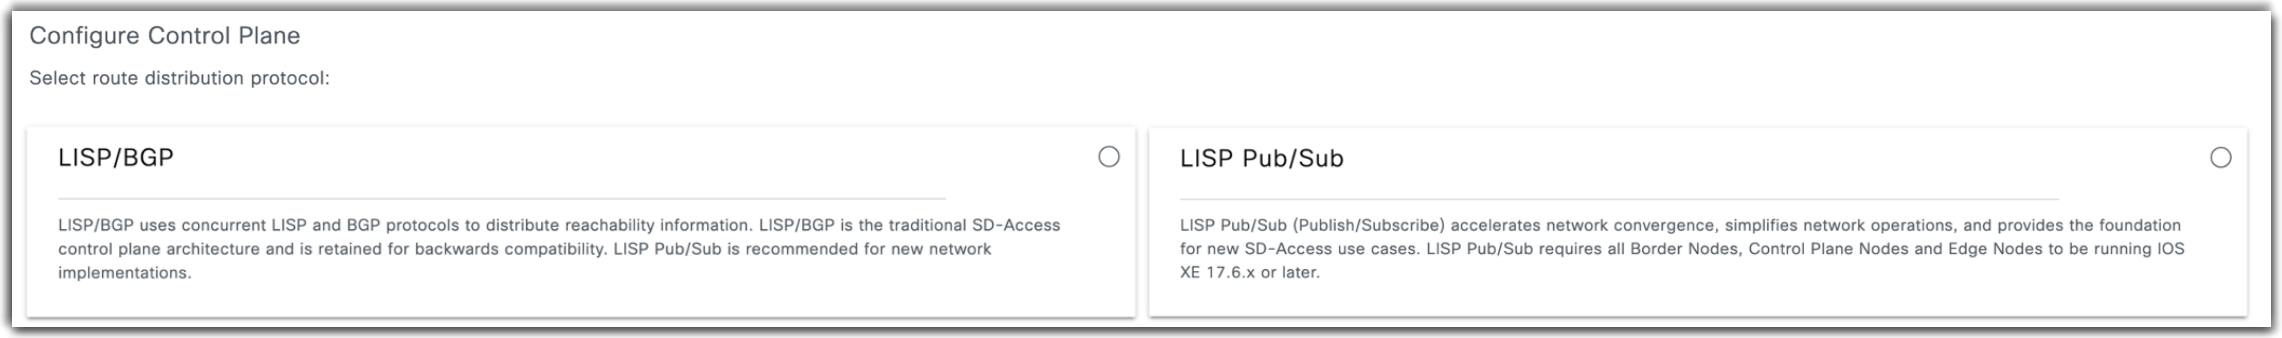

LISP/BGP

- Released circa 2017.
- Reliable and stable.
- BGP transport.

LISP Pub/Sub

- Released in 2022 with Cisco DNA Center* 2.2.3.x.
- Reliable and stable.
- Native LISP transport.
- Less Control Plane load.
- Faster convergence.
- Highly extensible.

*Rebranded to Catalyst Center in late 2023

LISP Control Plane

Host Registration

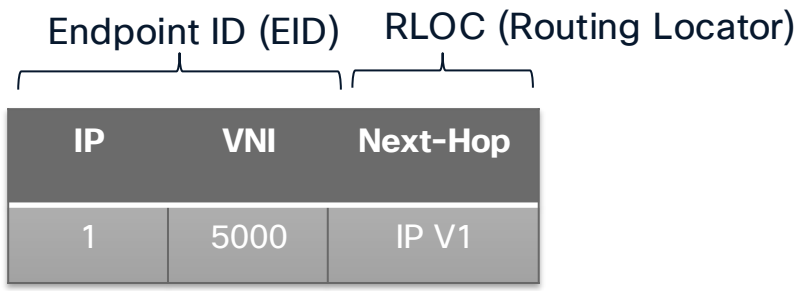

Map Register
EID = IP1, VNI 5000
RLOC = xTR IP V1

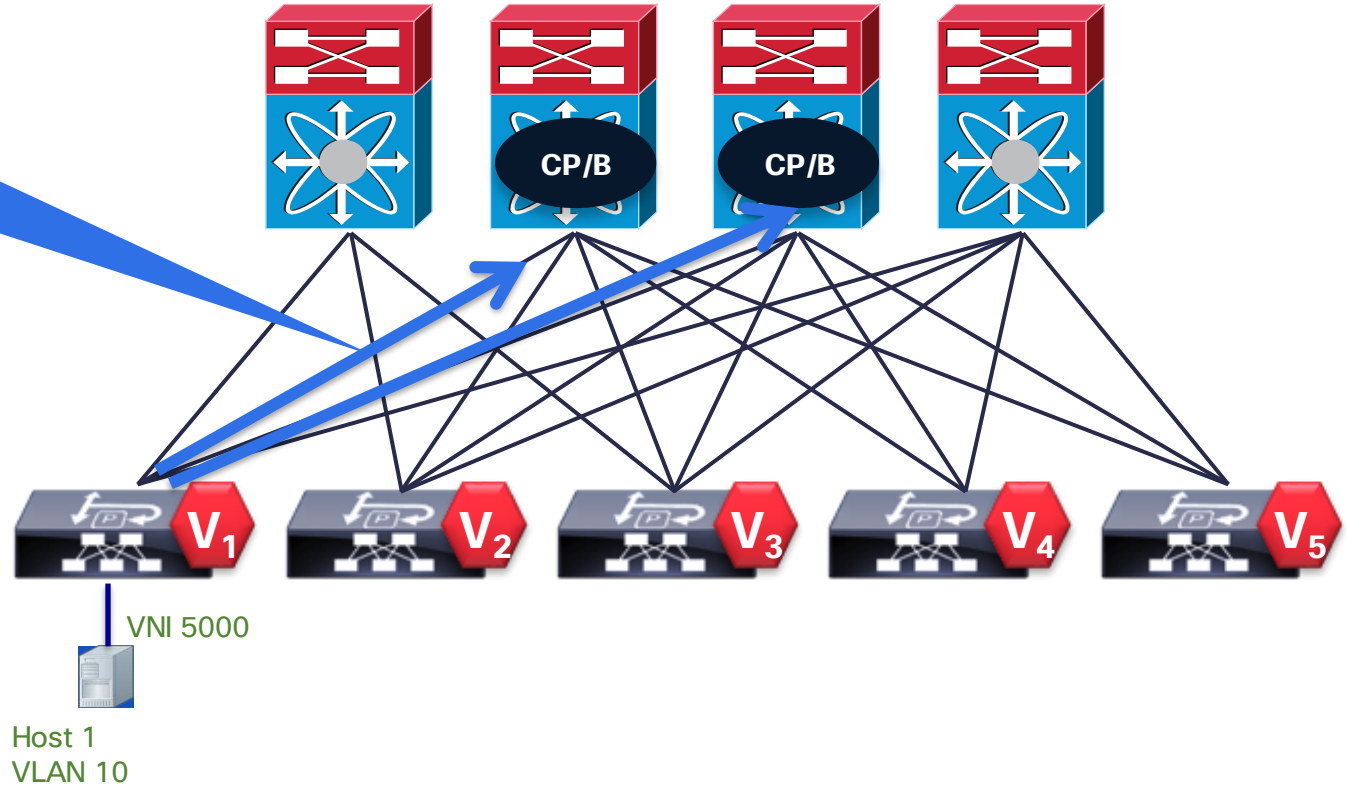

V1 LISP Tunnel Router (xTR) & VTEP*

CP/B CP + Border

* VTEP = VXLAN Tunnel End-Point

1. Attachment xTR registers host's IP (+MAC) in LISP
2. Scoped signaling between fabric nodes – fast convergence, scales and uses hardware resources efficiently

LISP Control Plane

Host Registration

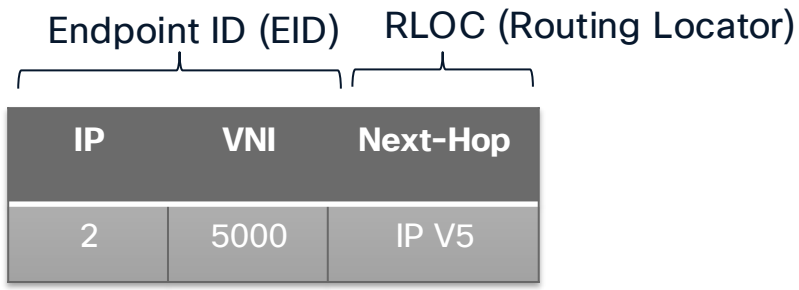

Map Register
EID = IP2, VNI 5000
RLOC = xTR IP V5

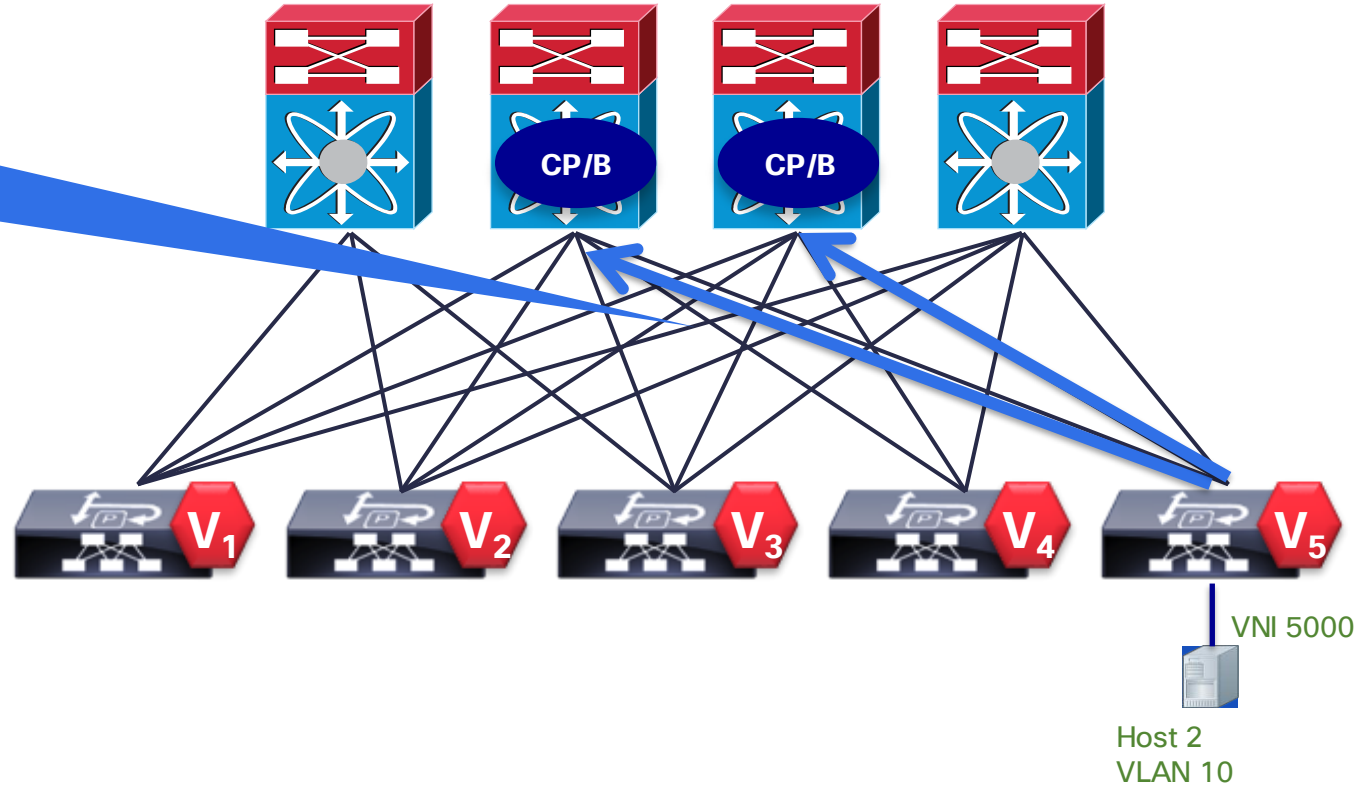

V1 LISP Tunnel Router (xTR) & VTEP*

CP/B CP + Border

* VTEP = VXLAN Tunnel End-Point

1. Attachment xTR registers host's IP (+MAC) in LISP
2. Scoped signaling between fabric nodes – fast convergence, scales and uses hardware resources efficiently

LISP Control Plane

Host Resolution

A key basic design objective:
Distribute routing/mapping information
only where it is required

A key basic working principle:
*Use traffic signals to **pull** routes on-demand*

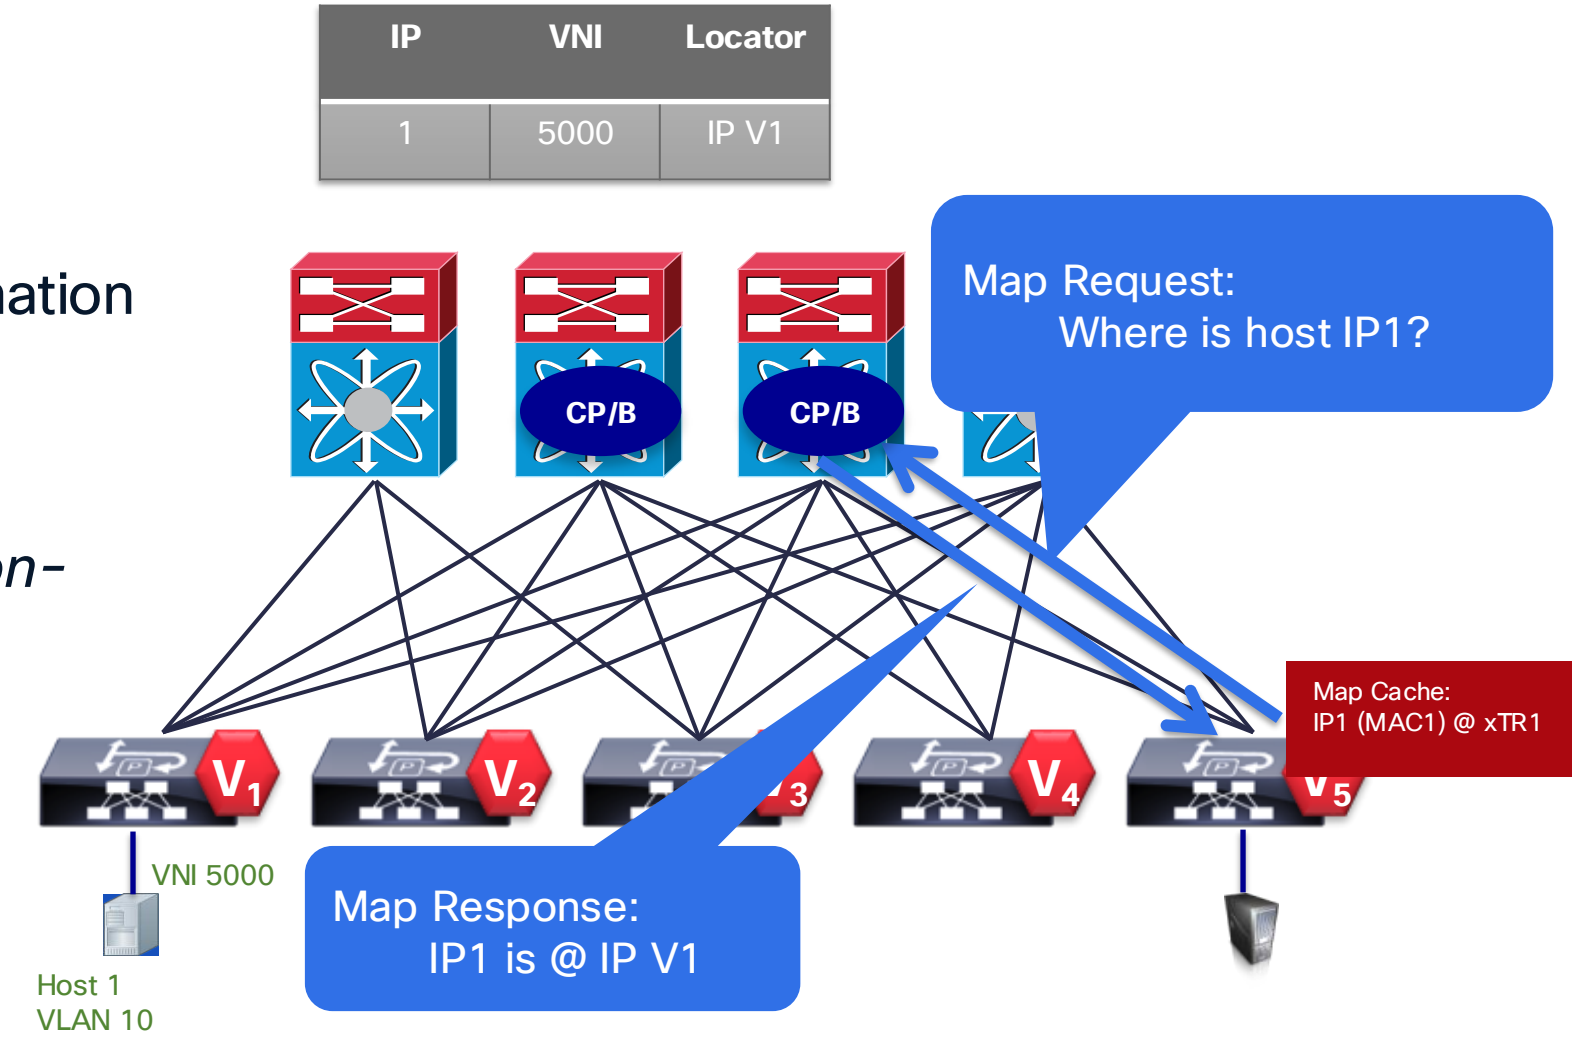

1. Host 2 wants to talk to host 1, the xTR (V5) issues a map-request
2. The Mapping System responds
3. The response is cached at the requesting xTR (V5): LISP map-cache

Cisco SD-Access Fabric

Control Plane: Locator/ID Separation Protocol (LISP)

LISP/BGP

Reliable and stable.
BGP transport.

2017

2022

LISP Pub/Sub

Released with Catalyst Center 2.2.3.x.
Reliable and stable.
Native LISP transport.
Less Control Plane load.
Faster convergence.
Highly extensible.

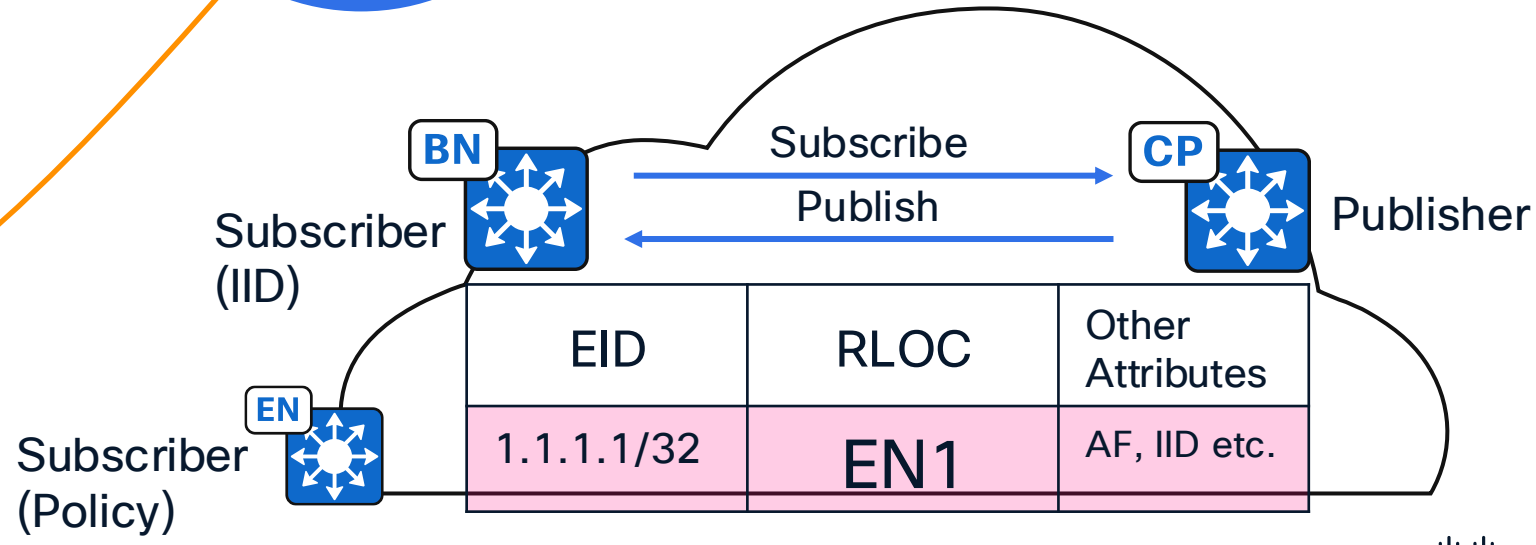

Fabric Fundamentals

Data Plane

Cisco SD-Access Fabric LISP Data Plane

Virtual Extensible Local Area Network (VXLAN)

VXLAN extends Layer 2 and Layer 3 overlay networks over a Layer 3 underlay network

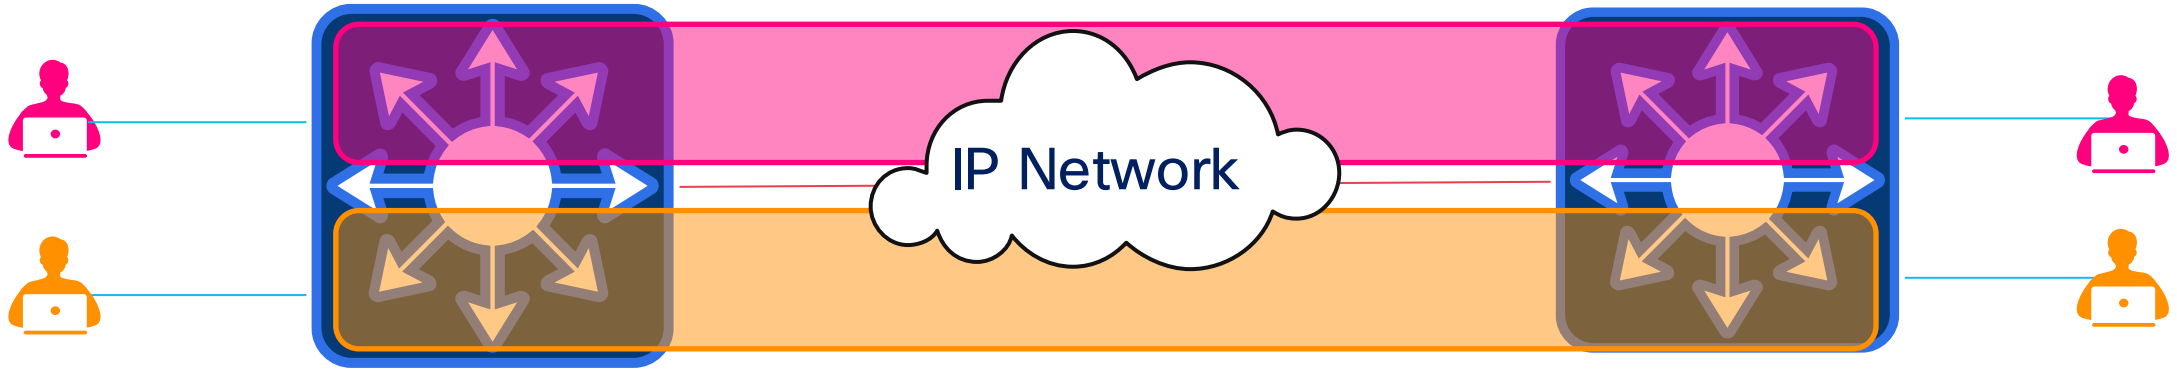

- ✓ Scalability: 16 million unique identifiers.
- ✓ Runs on top of L3, avoids need for STP.
- ✓ L2 traffic tunnelled over an L3 infrastructure.
- ✓ Handles broadcast, multicast, and unknown unicast traffic using multicast instead of flooding.
- ✓ Carries segmentation information.

Cisco SD-Access Fabric LISP Data Plane

VXLAN Packet Format

1. **Control Plane: LISP**
2. **Data Plane: VXLAN**

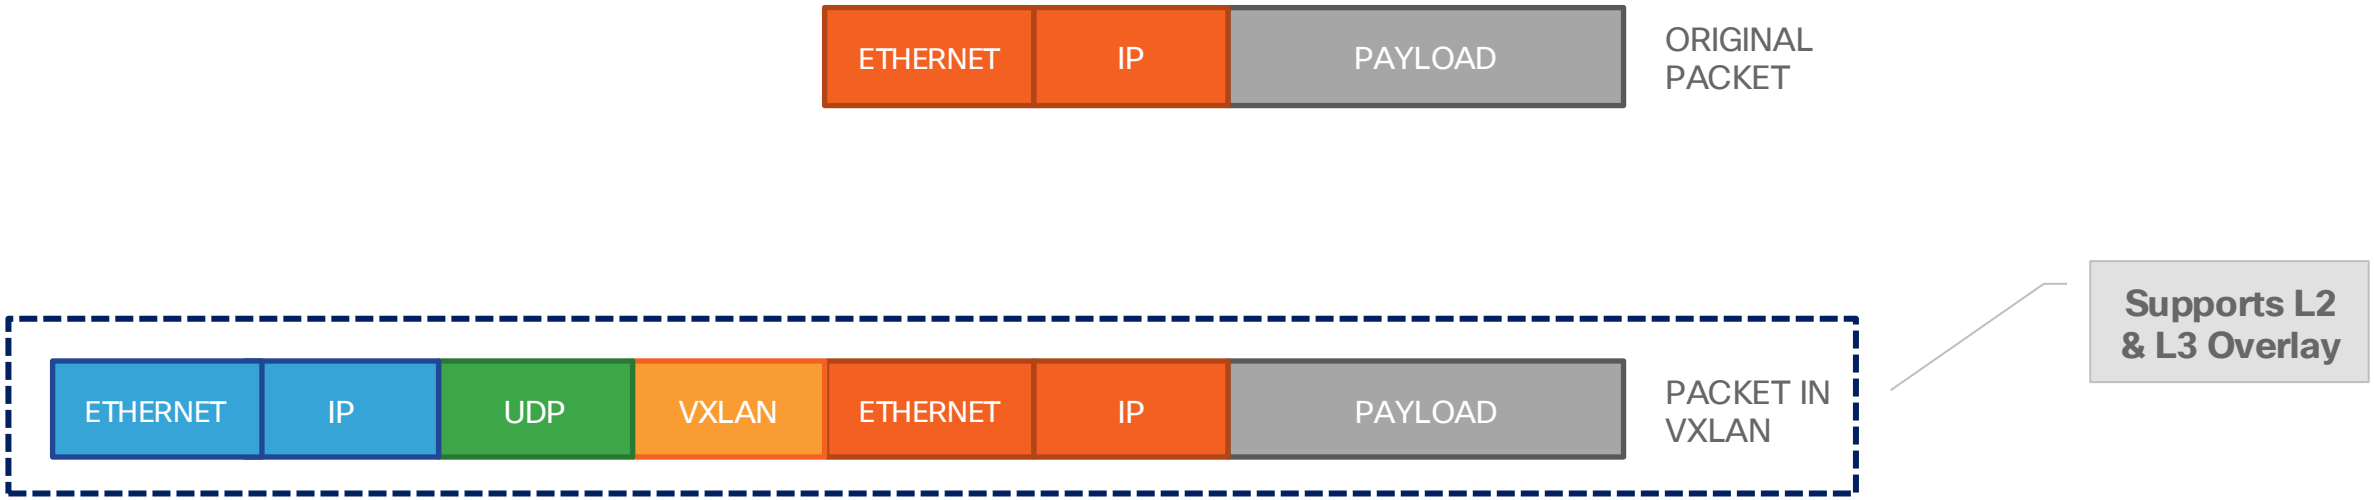

VXLAN-GPO Header

MAC-in-IP with VN ID and SGT ID

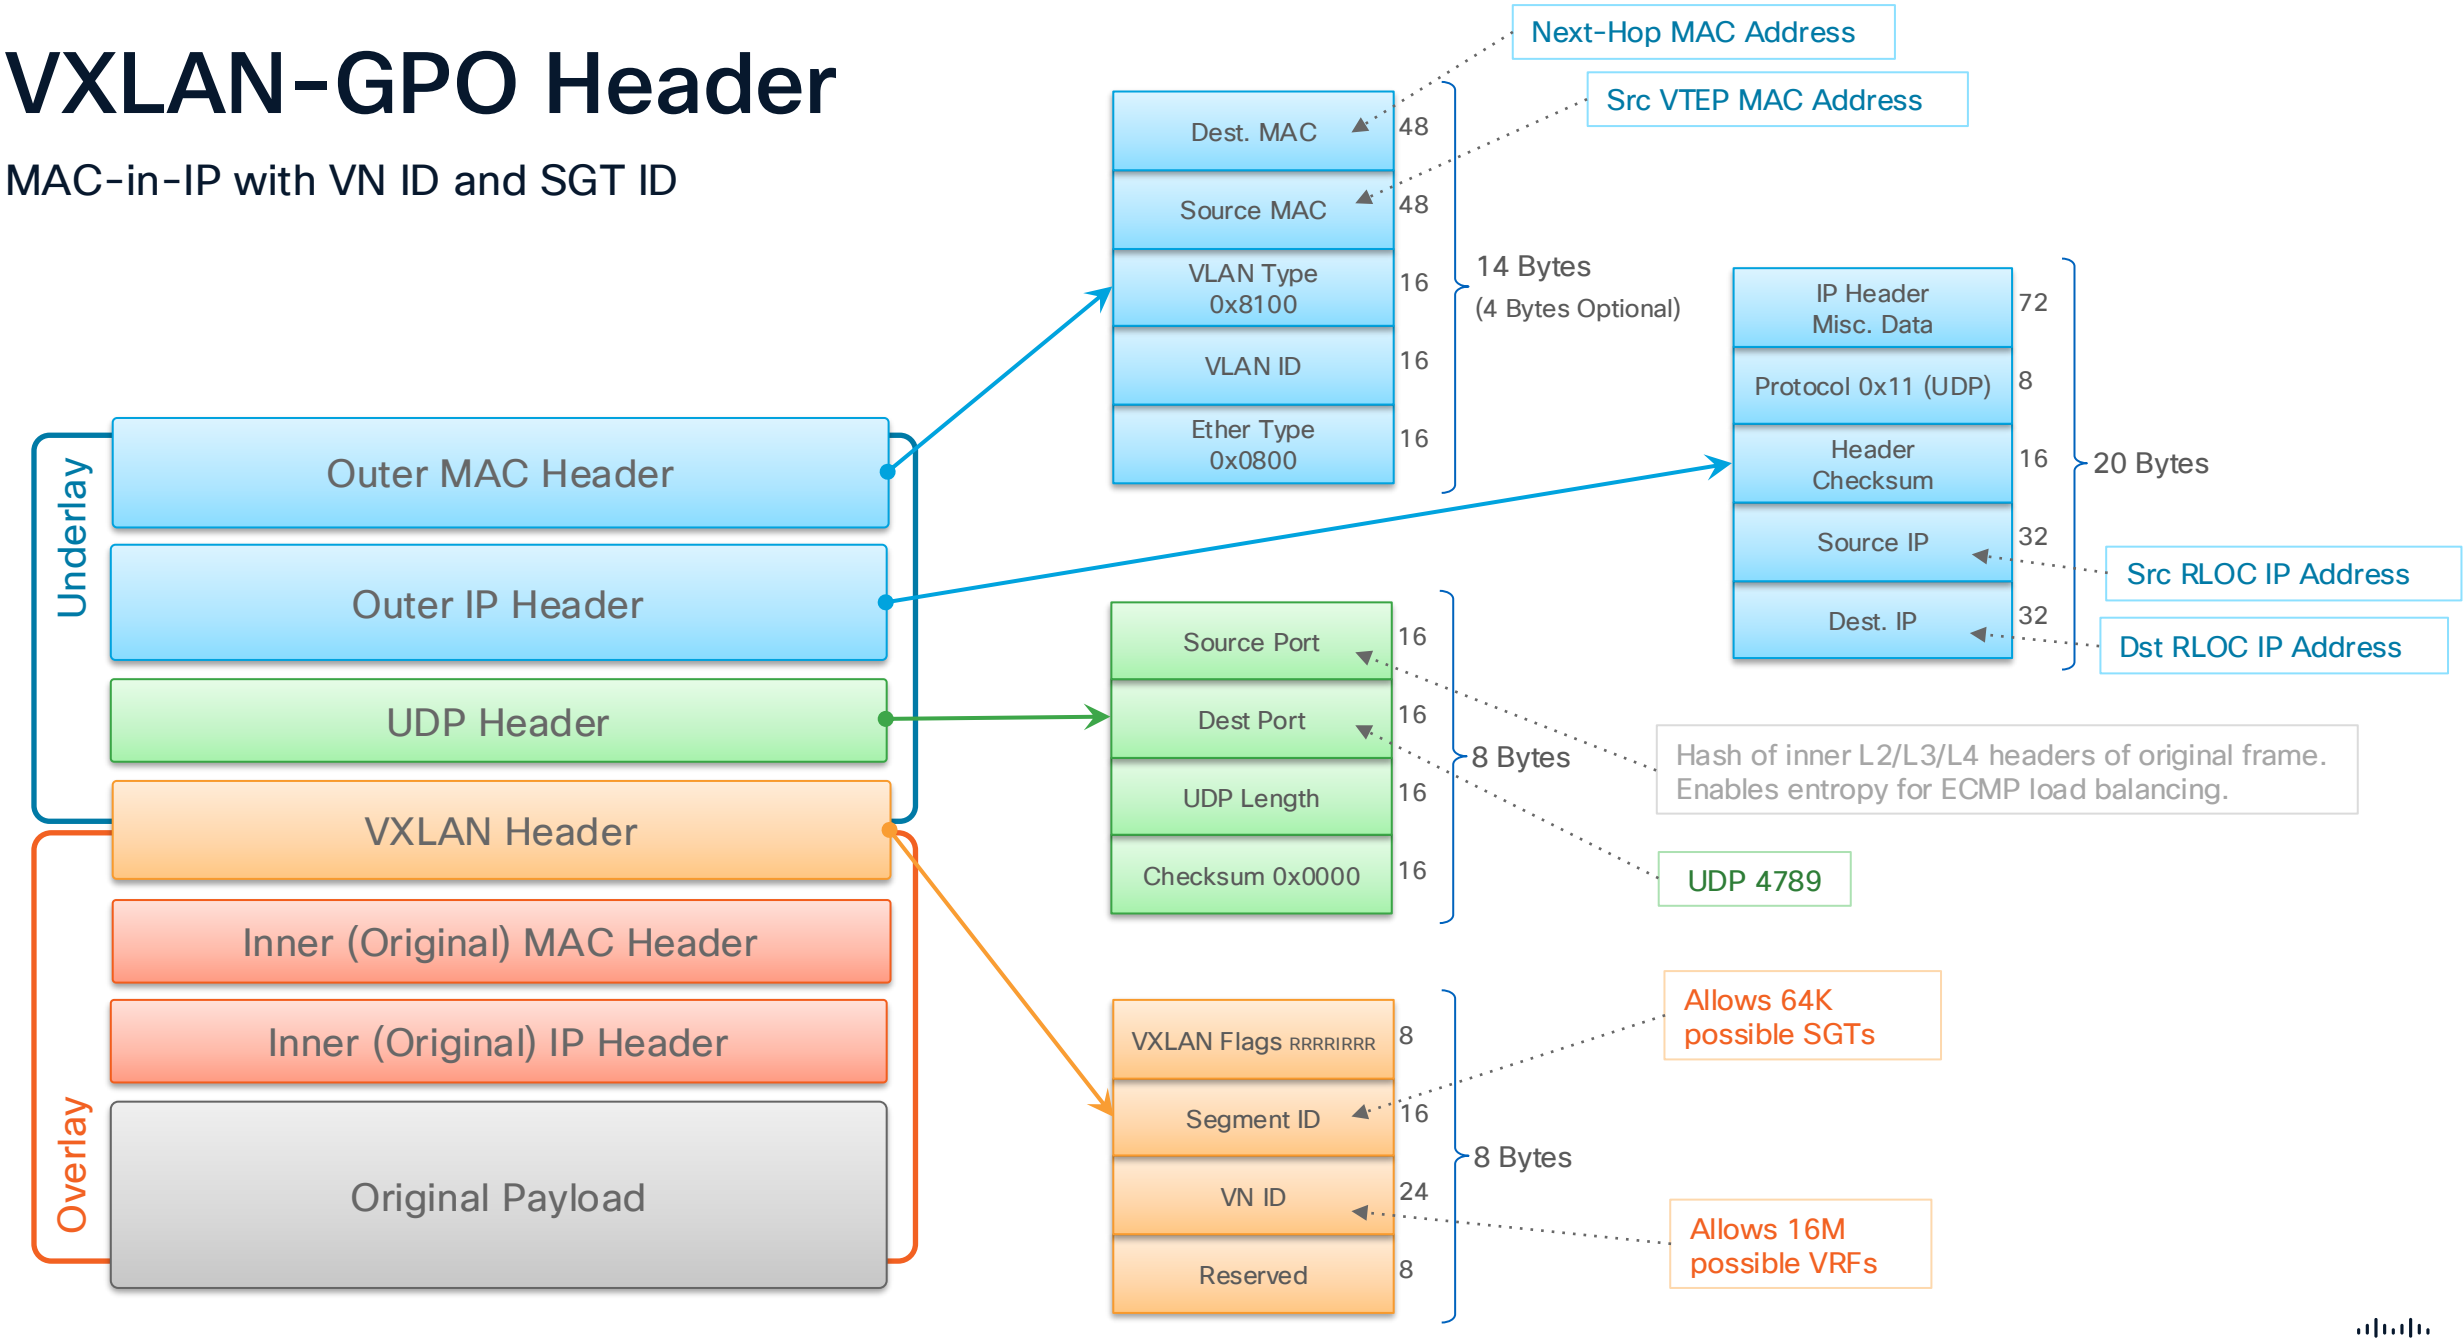

Fabric Fundamentals

Policy Plane

Cisco SD-Access Fabric LISP Data Plane

Policy Plane

1. **Control Plane: LISP**
2. **Data Plane: VXLAN**
3. **Policy Plane: Group-Based Policy**

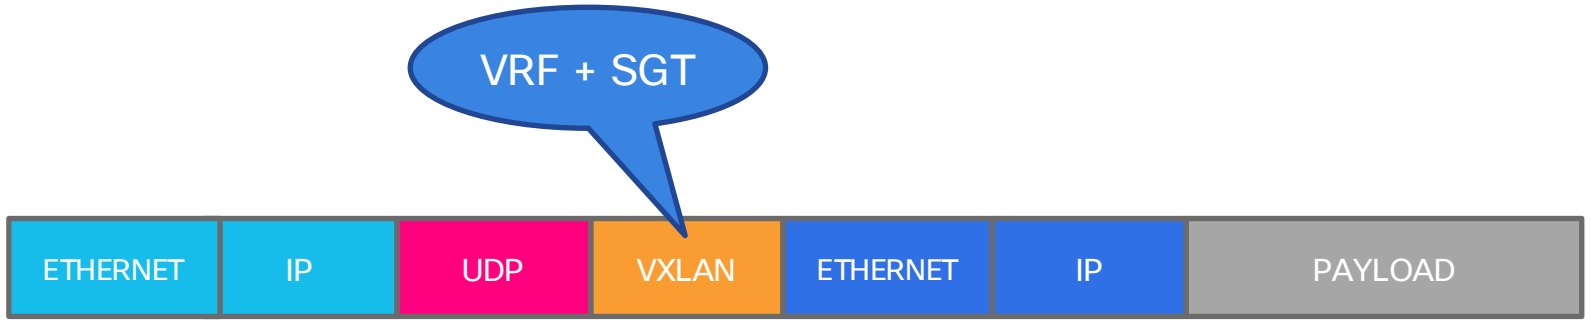

What is Security Group Tag and Group-Based Policy?

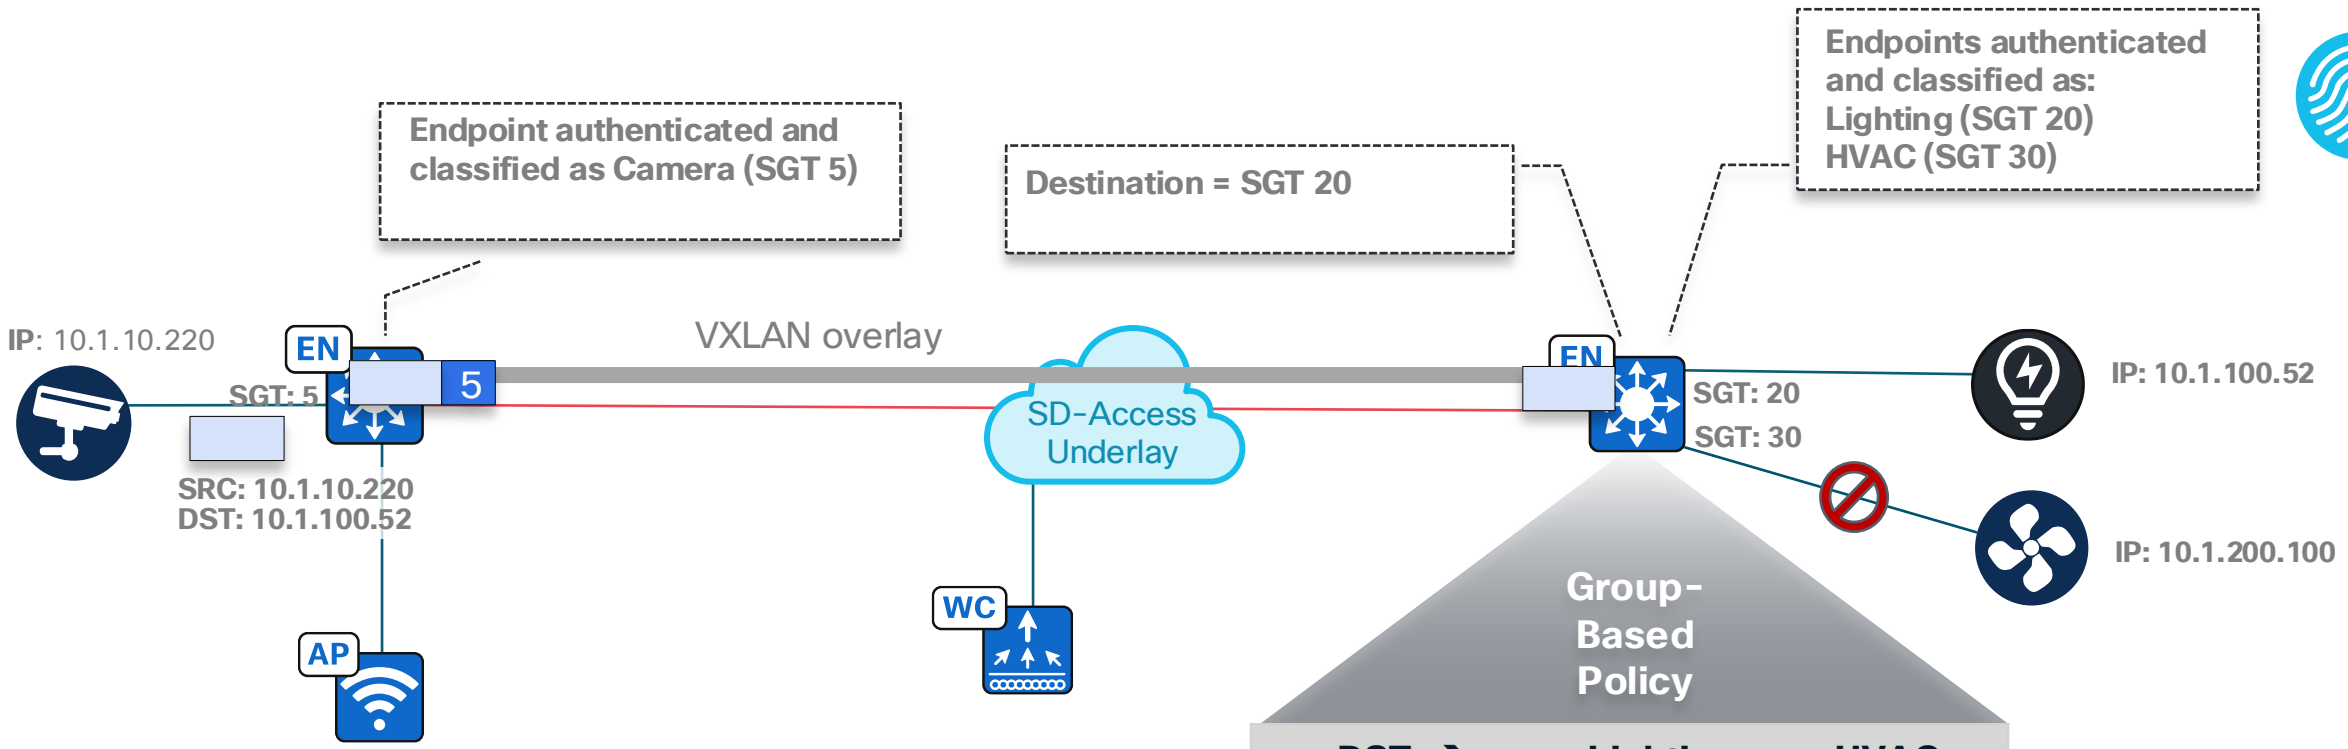

DST → / ↓ SRC	Lighting (20)	HVAC (30)
Camera (5)	Permit	Deny
BYOD (7)	Deny	Permit

SD-Access Policy

Macro-Segmentation and Micro-segmentation

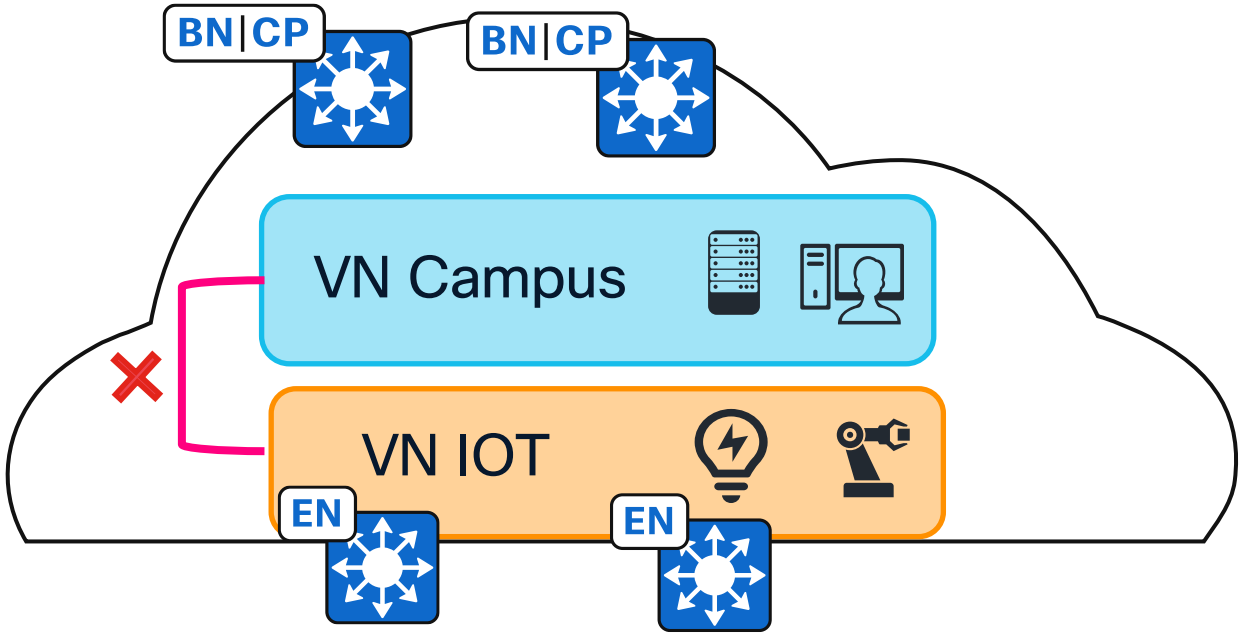

Virtual Network (VN)

First-level segmentation ensures **zero communication** between forwarding domains. Ability to consolidate multiple networks into one management plane.

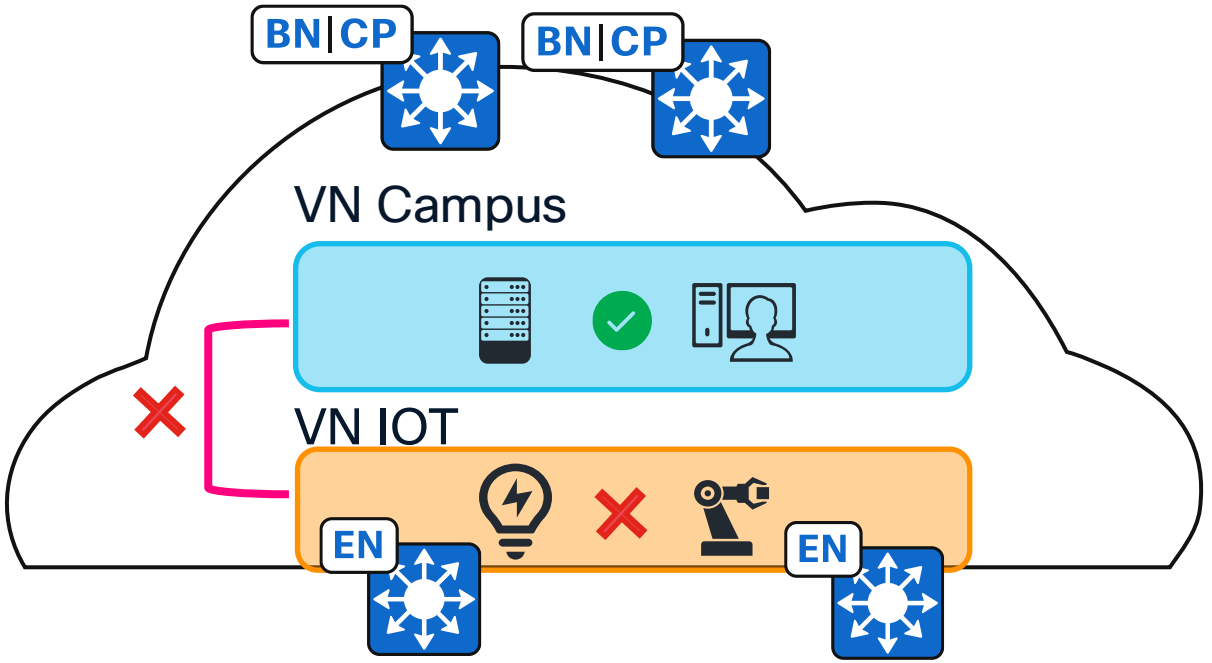

Security Group Tag (SGT)

Second-level segmentation ensures role-based access control between groups in a VN. Ability to segment the network into lines of business or functional blocks.

SD-Access Policy

Access Control Policies

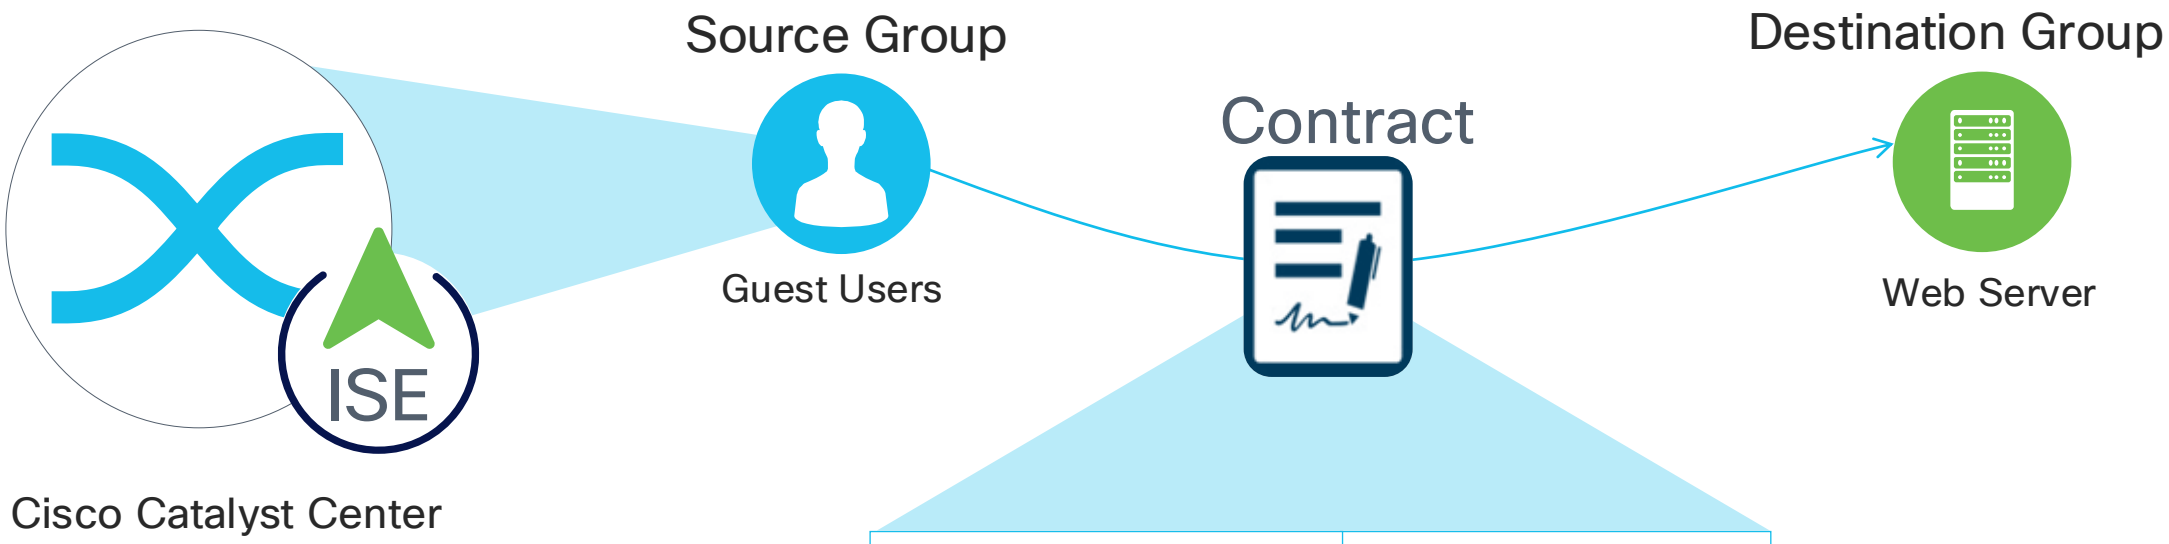

Classifier Type	Action Type
Port Number	Permit
Protocol Name	Deny
Application Type	Copy

Create and edit access contracts without knowing syntax for underlying SGACLs.

Multiple Fabric Sites

Transits for VN and SGT Preservation

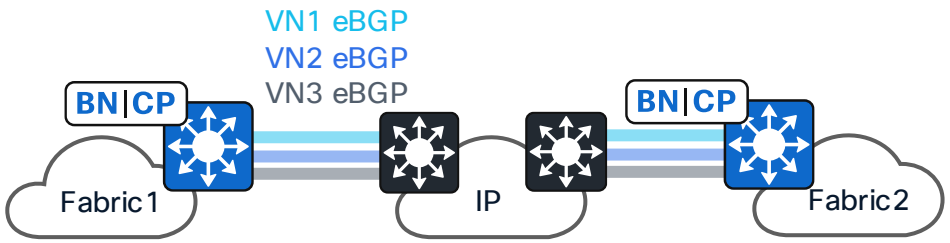

IP-Based Transit

- Per-Layer-3-Virtual-Network eBGP peering to external routing domain, or LISP Extranet Provider VN eBGP peering to external routing domain.
- SGT propagation outside of fabric requires suitable hardware and software.

Transits for VN and SGT Preservation

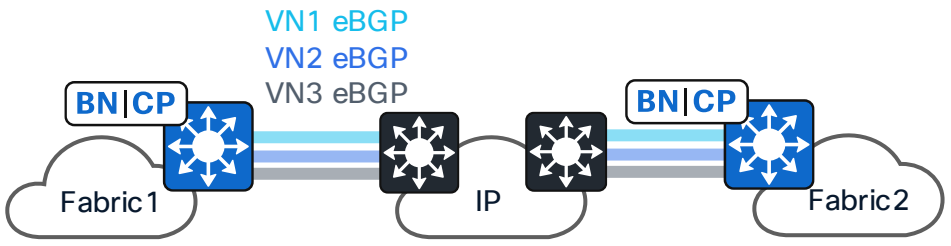

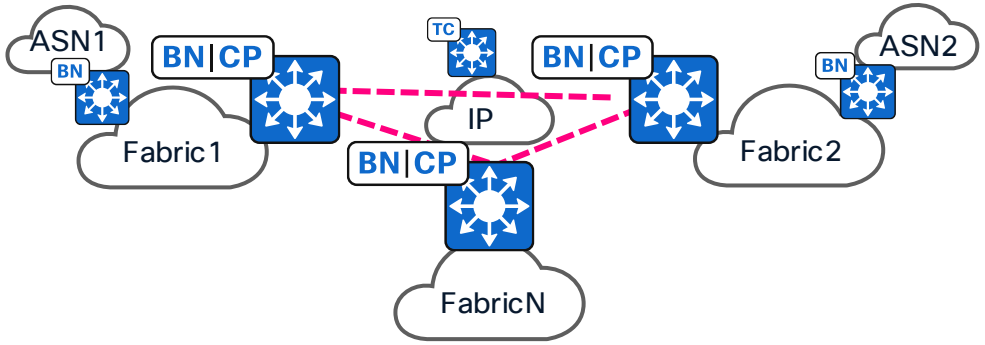

IP-Based Transit

- Handoff from the Border is automated with Cisco Catalyst Center
- Per-Layer-3-Virtual-Network eBGP peering to external routing domain, or LISP Extranet Provider VN eBGP peering to external routing domain.
- SGT propagation outside of fabric requires suitable hardware and software.

SD-Access Transit

- Automated from Cisco Catalyst Center
- Used only for inter-fabric-site traffic
- Uses VXLAN data plane between Fabric Sites.
- Preserves Layer 3 Virtual Networks and SGT.
- Fabric as a transit between external routing domains.

Watch BRKENS-2816 for SD-Access Transit deep dive

Transits for VN and SGT Preservation

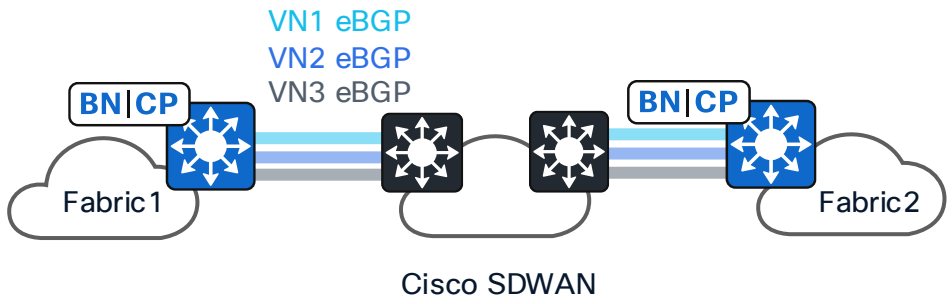

SD-WAN Transit

- Cisco / Meraki SD-WAN between Fabric Sites.
- Capable of preserving Layer 3 Virtual Networks and SGT's
- Dedicated SD-WAN Edge for design flexibility, Border Node port densities and port speeds

Design

External Dependencies

Before you spin up your first SD-Access fabric site, you will need:

- Catalyst Center – automation engine for SD-Access.
- DHCP / DNS – if you intend to provide these services to users connecting to SD-Access network.
- Cisco ISE – if you want to authenticate and authorize users or devices.
- Cisco WLC – if you want to provide wireless access. WLC can enable fabric-enabled wireless for a single site only.
- Fusion device (typically a firewall) to implement VRF route-leaking and enforce security policy at the leaking point.

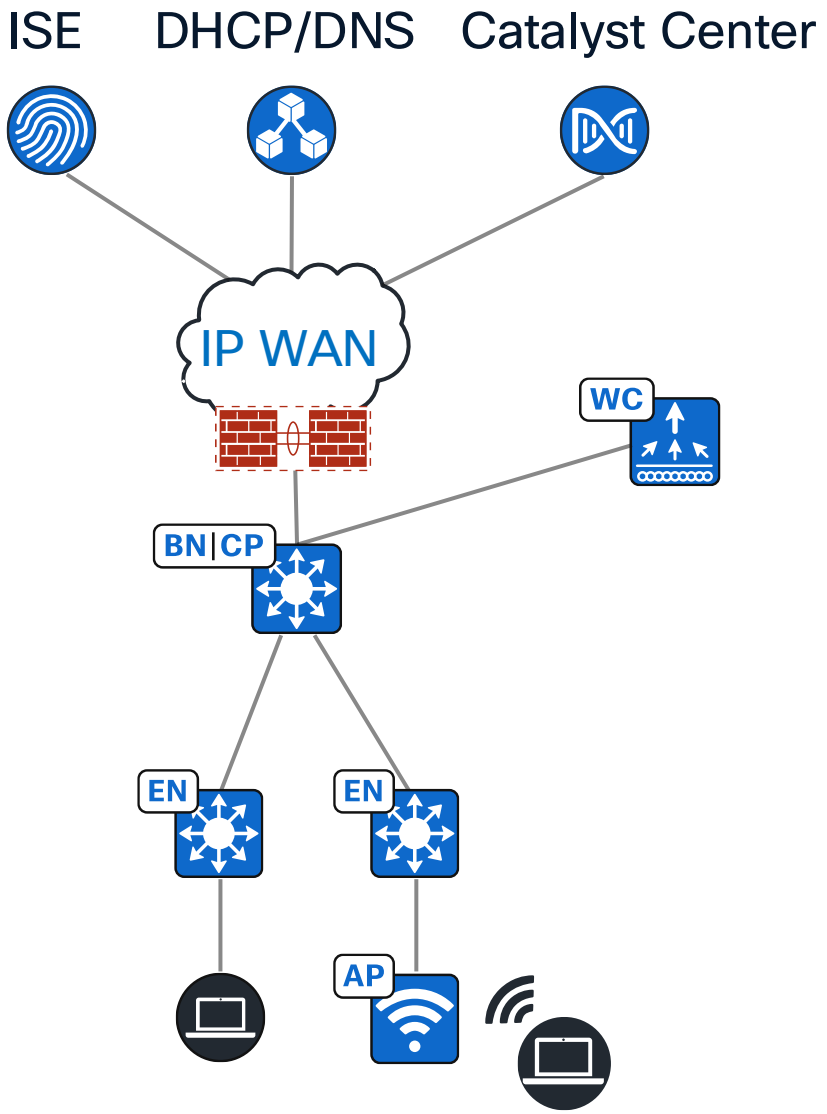

External Dependencies

All external dependencies reside outside the fabric site and just need IP (Layer 3) connectivity to fabric devices. Latency requirements:

- Catalyst Center to fabric devices – 200ms RTT.
- ISE to fabric devices – 100ms RTT.
- Fabric WLC to fabric APs – 20ms RTT (put it onsite).

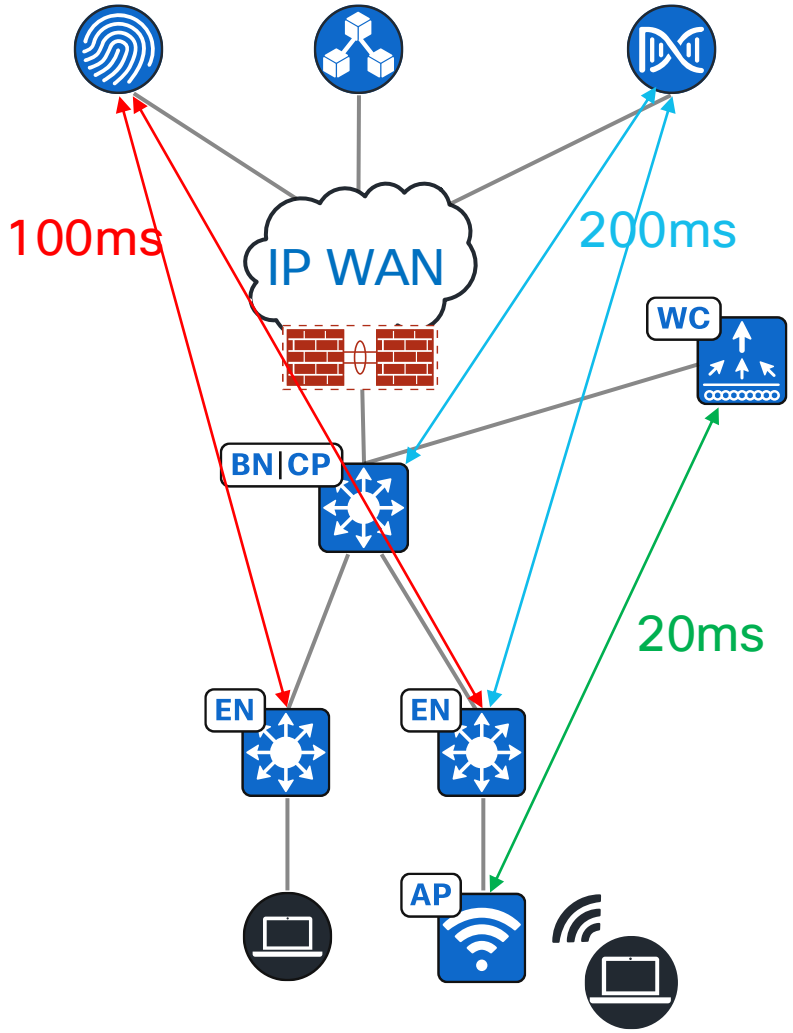

How Would Your Carve Your Fabric Sites?

- Fabric site is an instance of an SD-Access Fabric.
- A collection of Edge Node switches using the same set of CP/BN switches.
- Typically defined by disparate geographical locations, but not always.
- Can also be defined by:
  - Endpoint scale.
  - Failure domain scoping.
  - Underlay connectivity attributes (MTU, multicast).
- Typically interconnected by a “Transit”.

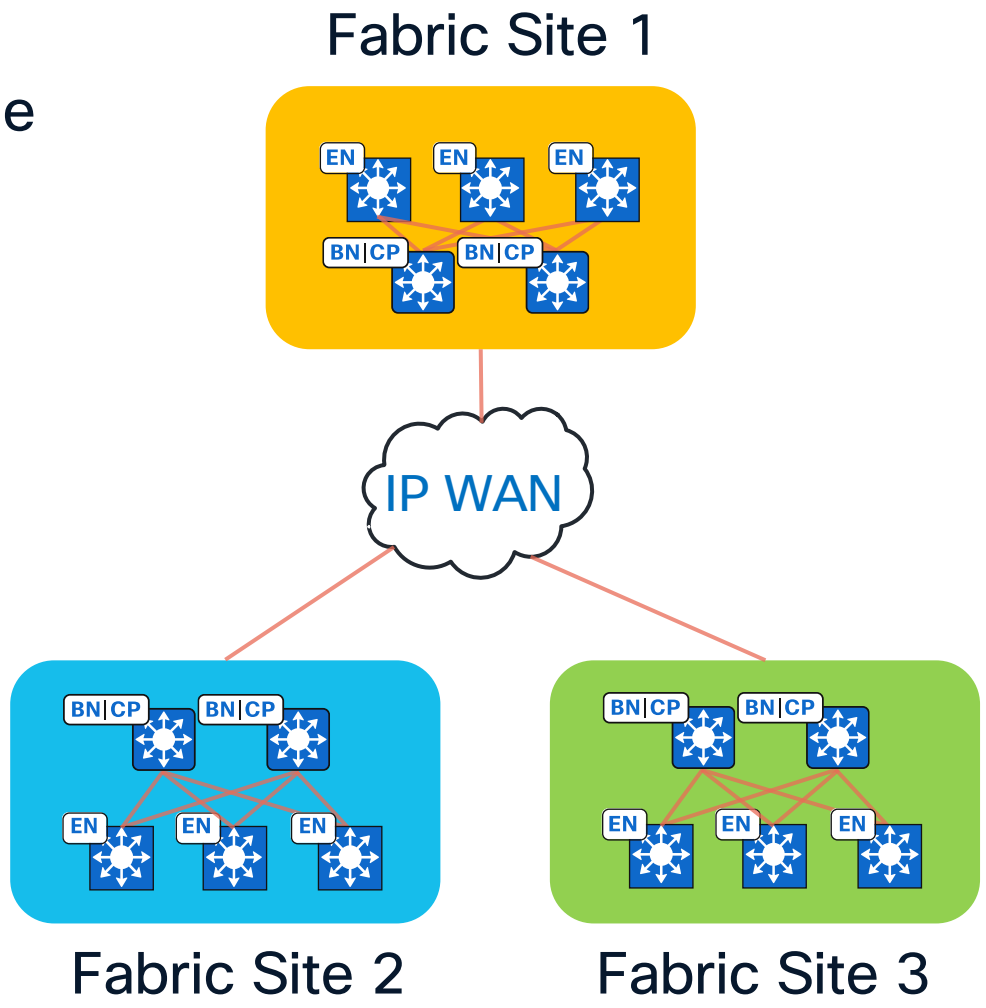

Site Limits – Endpoint Scale

Control Plane Node keeps information about all site endpoints in **RAM** and uses **CPU** to process it (including wireless roaming events).

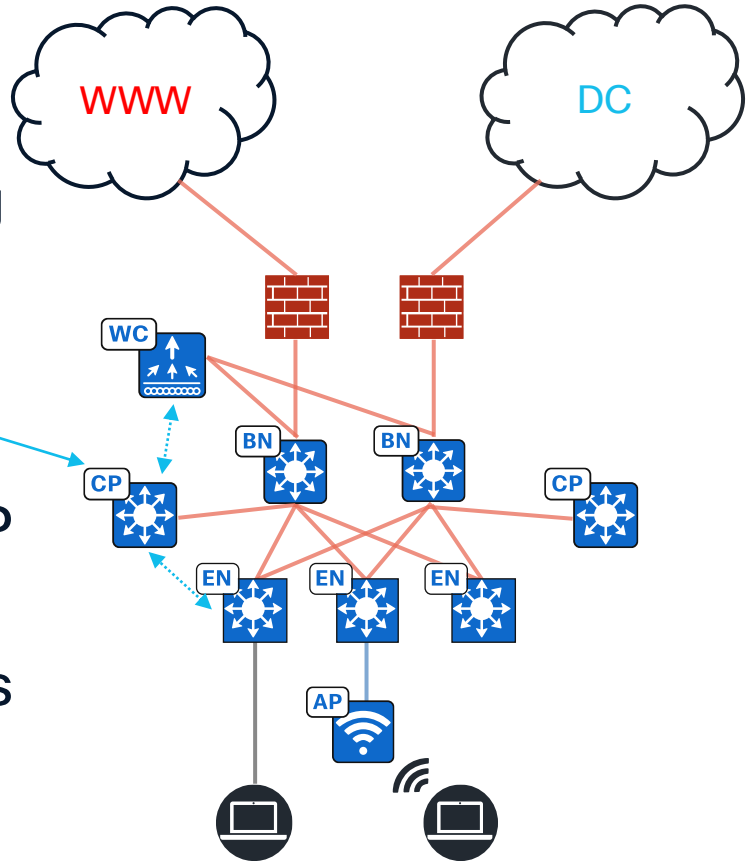

- C9300\L switches can support up to 16,000 EPs as CP node.
- C9500-32C / C9500-48Y4C / C9500-24Y4C switches can support up to 80,000 EPs as CP node.
- Other C9K switches are possible in CP role, sizing values are documented in Catalyst Center Data Sheet.

Site Limits – Endpoint Scale

Border Node keeps all EP information in **TCAM** as host routes. If EP has multiple IP addresses (v4 + multiple v6), each address is counted as individual entry.

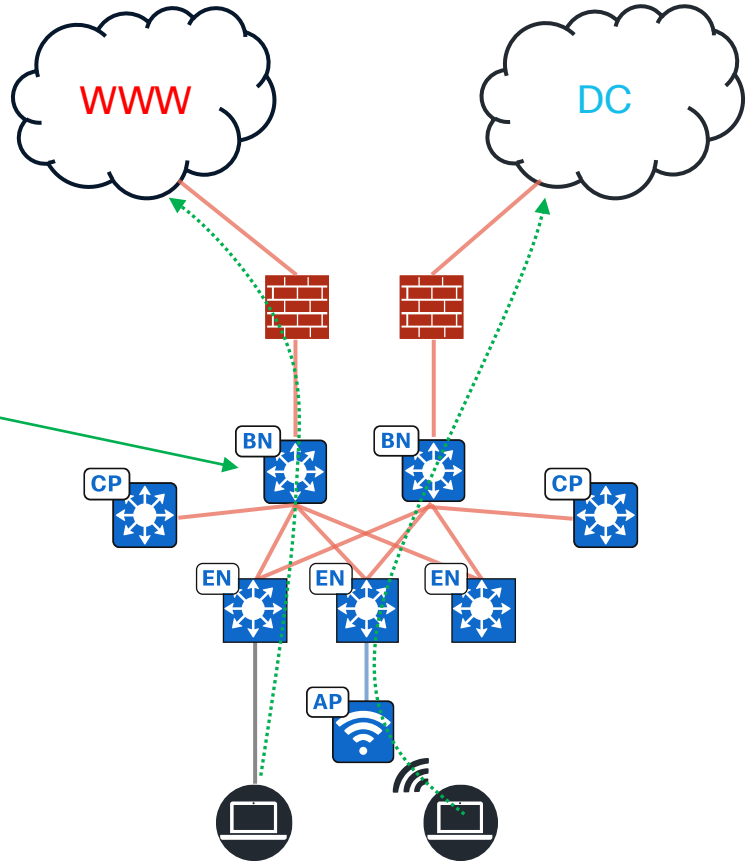

- C9300\L switches can support up to 16,000 IP host routes (/32 or /128) as Border Node.
- C9500-32C / C9500-48Y4C / C9500-24Y4C switches can support up to 150,000 IP host routes as Border Node.

Full border node sizing values for all SD-Access platforms are documented in Catalyst Center Data Sheet.

Site Limits – Failure Domain Scoping

- All Edge Nodes in the site are sharing the same set of Control Plane and Border Nodes. If all CP or BN nodes fail, the site is failed*. SD-Access site with fabric wireless can have 2 CP nodes max.
- A lot of configuration elements (VRF, VLAN, multicast, wireless, default switchport policy) are applied at the site level, to all** fabric site switches at the same time.
- Fabric site is underpinned by a single instance of underlay routing protocol (IGP) as well as overlay routing protocol (LISP) and is visible as single BGP AS from the outside world.

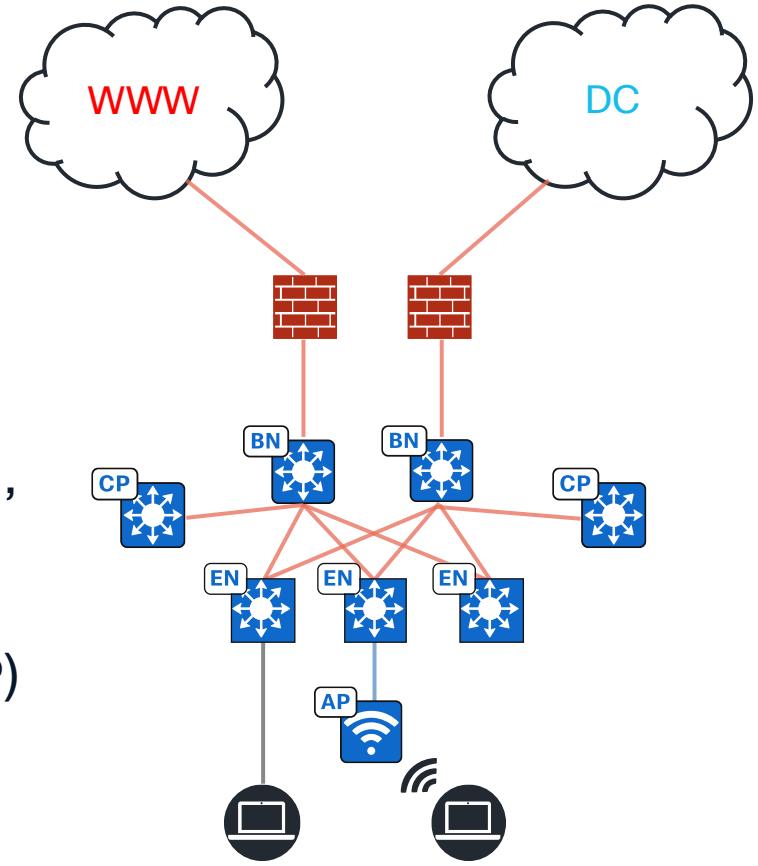

*During a total CP failure, no new endpoints can be onboarded into the fabric and roaming events won't work. Existing traffic flows will be cached for 24 hours.

**Some changes can be scoped to a limited subset of switches via Fabric Zones.

Site Limits – Underlay Connectivity Attributes

- Avoid mixing different underlay connectivity attributes, such as MTU or multicast support because you will end up dropping to the lowest common denominator within a fabric site.

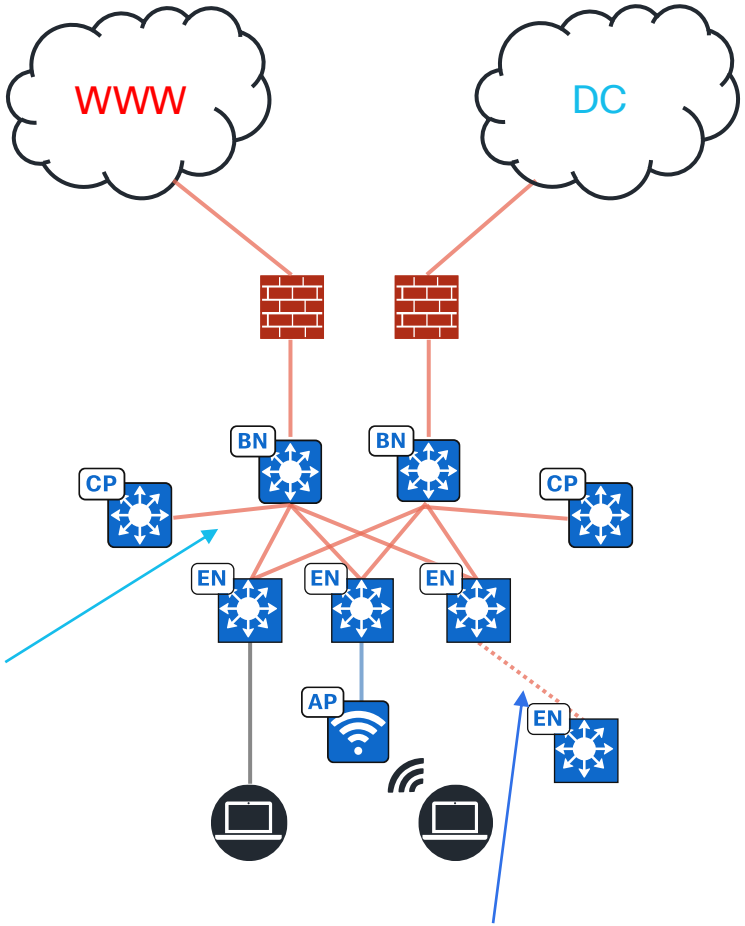

Multiple Fabric Sites vs Single Fabric Site?

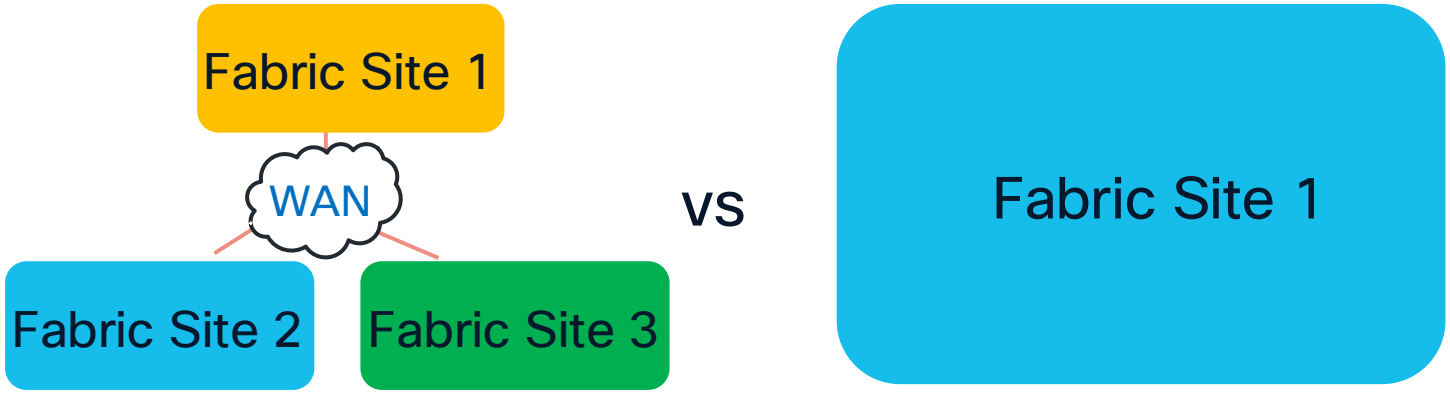

Make large single fabric site within single geographical area

- You hit fabric device (1200 logical switches for –XL Catalyst Center) or endpoint limit (~100,000 EPs).
- Links between parts of your fabric site can support increased MTU (from 1550 to 9000 bytes) and can be multicast-enabled.
- Part of your fabric site needs to be online even if the rest of your site is offline.
- Part of your fabric site needs to provide Direct Internet Access for users in the overlay.

Control Plane – Pub/Sub or Not?

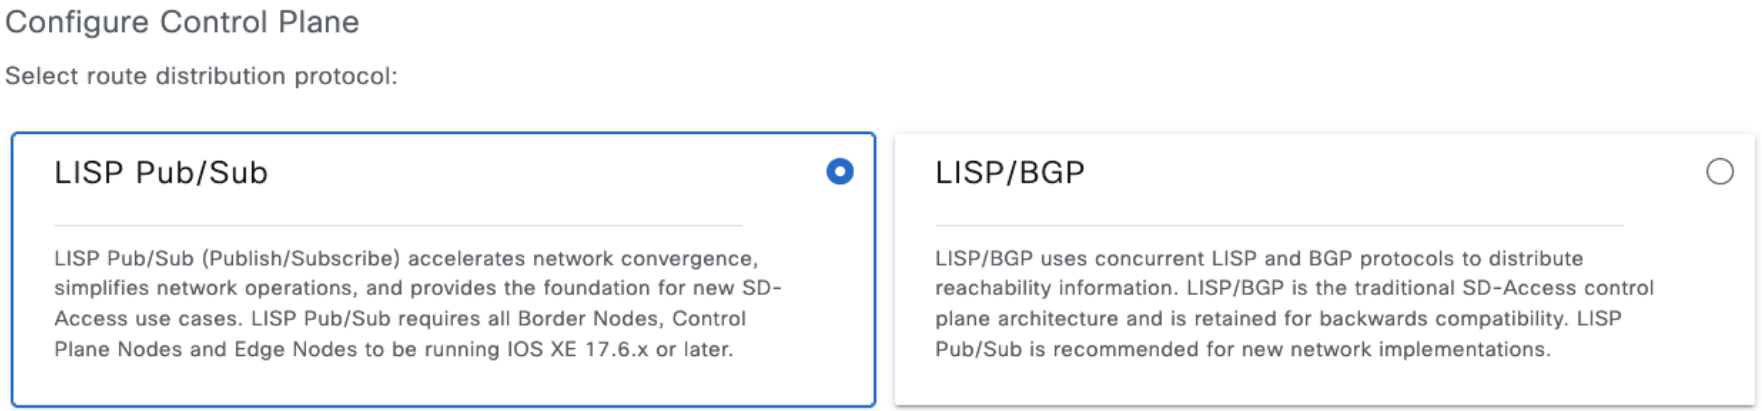

LISP Pub/Sub

- Released in 2022 with Catalyst Center 2.2.3.X and IOS-XE 17.6.X.
- Reliable and stable.
- Less Control Plane load.
- Faster convergence.

→ Greenfield: deploy LISP Pub/Sub.

- Requires default route (0.0.0.0/0) from upstream to work in External Border capacity.
- No longer need per-VN iBGP peering between Border Nodes.
- All sites connected via SDA Transit need to be on the same CP architecture (Pub/Sub or LISP/BGP).

Control Plane – Colocate with Border or Not?

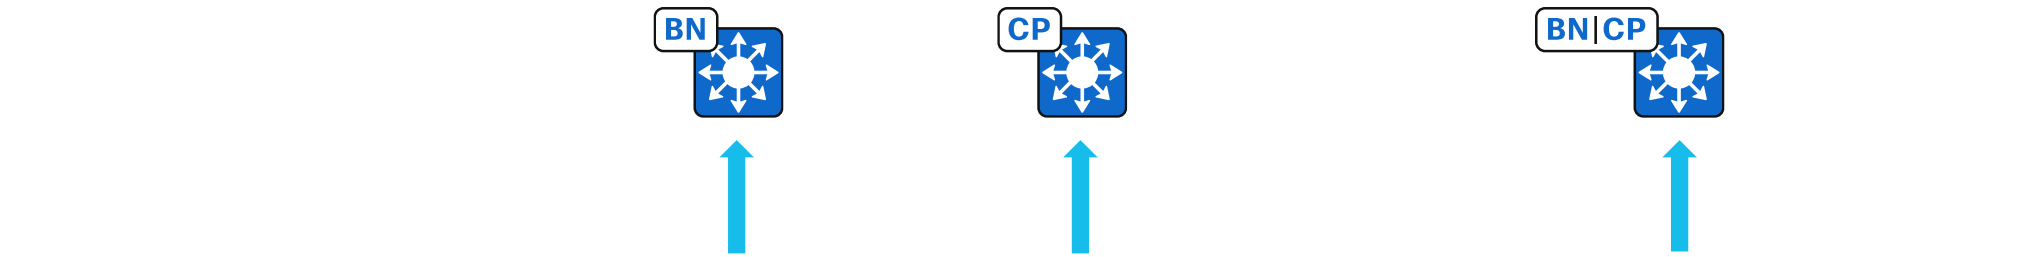

- Border Node downloads all fabric host routes in switch TCAM.
- Control Plane RAM is non-issue from scale perspective.
- Main CPU stress for CP is handling wireless roaming for Fabric Enabled Wireless endpoints.
- It is safe to colocate until 50,000 EPs**, even in wireless-heavy environment.
- Can split BN and CP for architectural reasons (fault isolation, network modularity), rather than technical (scale).
- Avoid using routing platforms (C8K) as Control Plane and/or Border Nodes if possible.

*Number of host (/32 or /128) routes **C9500H or above

SD-Access Design Resources

Cisco Validated Design:
https://cs.co/sda-cvd

SDA Design Tool (use Chrome):
http://cs.co/sda-design-tool

SDA Compatibility Matrix:
http://cs.co/sda-compatibility-matrix

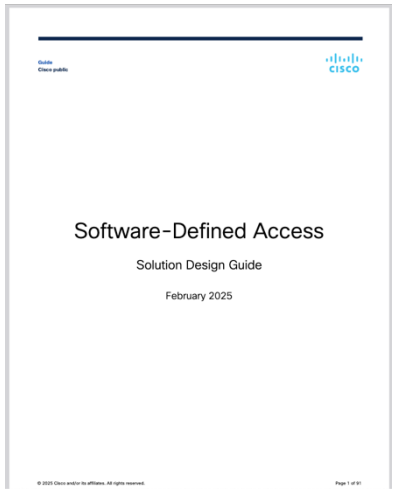

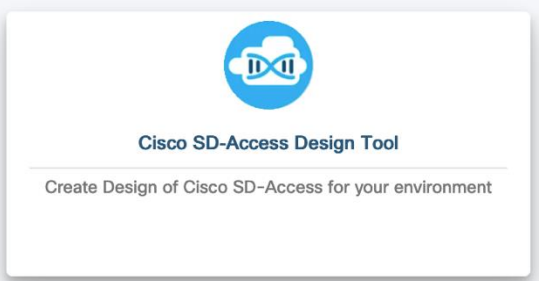

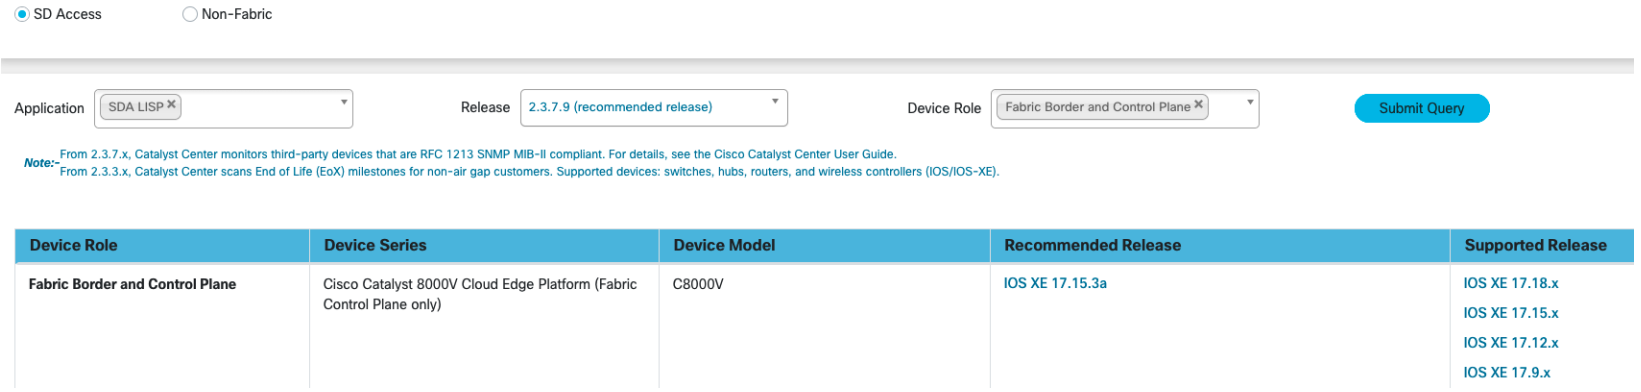

Device Role	Device Series	Device Model	Recommended Release	Supported Release
Fabric Border and Control Plane	Cisco Catalyst 8000V Cloud Edge Platform (Fabric Control Plane only)	C8000V	IOS XE 17.15.3a	IOS XE 17.18.x IOS XE 17.15.x IOS XE 17.12.x IOS XE 17.9.x

Thank you