Advanced Campus Network Design

Multilayer Architectures and Min / Maxing

Shawn Wargo
Principal Tech Marketing Engineer

Jakub Matela
Technical Solutions Architect

BRKENS-2500

Cisco Webex App

Questions?

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

How

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

Webex spaces will be moderated by the speaker until June 13, 2025.

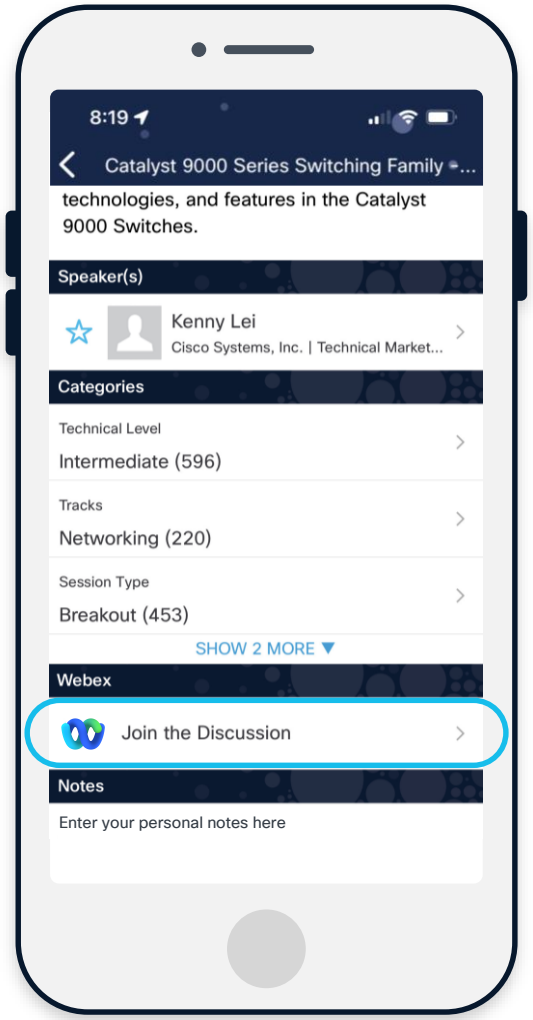

Who are we?

Shawn Wargo
Principal TME

swargo@cisco.com

I'm a **Principal Engineer of Technical Marketing** (PTME) for Cisco Enterprise 'Network Experience' (NX) Product Management team. I've been with Cisco **since 1999**.

I mainly focus on the **Enterprise Switching & Routing** technology areas, with a special emphasis on '**next generation**' **Hardware & Software** products and solutions.

As a Principal TME, I'm currently working on the next generation of **Catalyst Switching, Wireless & Routing** products, and solutions like Software-Defined Access (SDA) & Cisco Catalyst Center.

Jakub Matela
Technical Solutions Architect

jamatela@cisco.com

I'm a **Technical Solutions Architect** (TSA) at Cisco, part of the EMEA Enterprise Networking team. I joined Cisco in 2016.

Since 2021, I have been leading Cisco's **Enterprise Networking Switching, Software-Defined Access, and Catalyst Center** technologies in EMEA Sales.

I am dedicated to enabling the field, partners, and customers in their **transition to intent-based networking, leveraging Software-Defined Access and Cisco Catalyst Center.**

Based in Krakow, Poland, I graduated from AGH University of Science and Technology with a **Master's in Electrical Engineering** and hold a **CCIE in Enterprise Infrastructure**.

Campus Architecture - Series Agenda

Design Fundamentals

1. Campus Design Fundamentals
2. Campus Design Principles
3. Campus Foundational Services

Design Considerations

4. Platform Design Considerations
5. Campus Design Best Practices
6. Campus integration with other PINs

Session Agenda - BRKENS-1500

Design Fundamentals

1 Campus Design Fundamentals

- What is “Campus”?
- Place in Network (PIN)

2 Campus Design Principles

- Multi-Layer Model
  - Hierarchical Design
- Access Layer
- Distribution Layer
- Core Layer

3 Campus Foundational Services

- Layer 1 physical layer & links
- Layer 2 switching protocols
- Layer 3 routing protocols

Design Considerations

4 Platform Design Considerations

- Chassis Considerations (Capacity)
- Cabling Considerations (Speed)
- Feature Considerations (Scale)
  - L2 Features
  - L3 Features
  - Quality of Service (QoS)

5 Campus Design Best Practices

- LAN High Availability
- LAN Security
- Virtual Networking

Session Agenda - BRKENS-2500

Design Fundamentals

1 Campus Design Fundamentals

- What is “Campus”?
- Place in Network (PIN)

2 Campus Design Principles

- Multi-Layer Model
  - Hierarchical Design
  - 1,2,3 & 4+ Tiers
- Access Layer
  - Baseline, Extended Access, Routed Access
- Distribution Layer
  - Baseline, Collapsed Core, Collapsed Distro
- Core Layer
  - Baseline, Interconnect, Edge

3 Campus Foundational Services

- Layer 1 physical layer & links
- Layer 2 switching protocols
- Layer 3 routing protocols
- ECMP, LAG & Load balancing

Design Considerations

4 Platform Design Considerations

- Chassis Considerations (Capacity)
- Cabling Considerations (Speed)
- Feature Considerations (Scale)
  - L2 (Unicast & Multicast)
  - L3 (Unicast & Multicast)
  - Security (AAA & ACL)
  - Quality of Service (QoS)
  - NetFlow (AVC & XDR)

5 Campus Design Best Practices

- LAN High Availability
  - SSO/NSF, Stack/SVL, mLAG, FHRP
- LAN Security
  - NAC, Access Control, FHS, ZTNA
- Virtual Networking
  - MPLS, LISP, EVPN

6 Campus integration with other PINs

- Wireless Integration
- Firewall Integration

Session Agenda

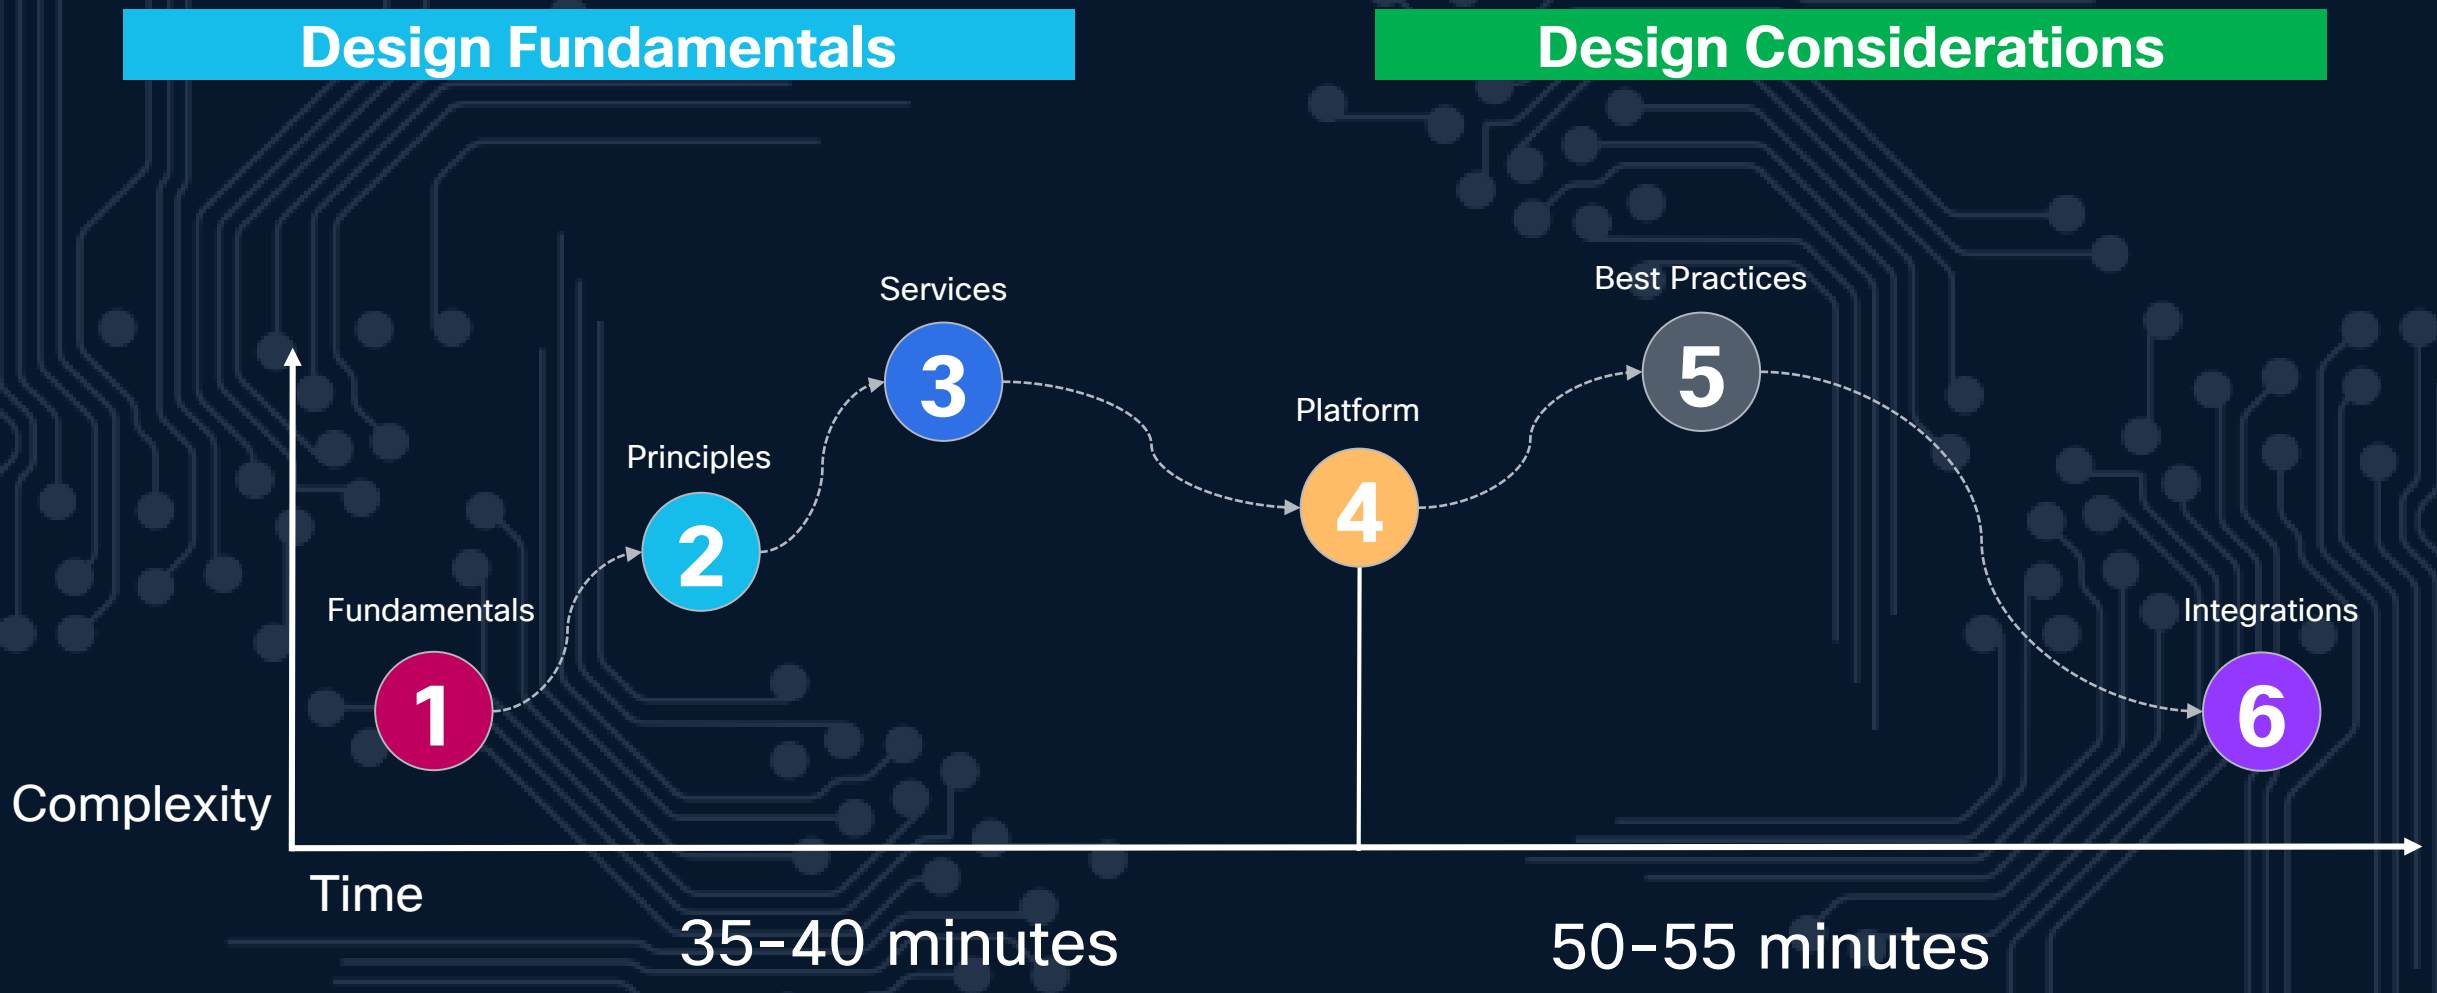

Campus Design Fundamentals

Chapter 1

Jakub Matela
Technical Solutions Architect

Session Agenda

Design Fundamentals | **Design Considerations**

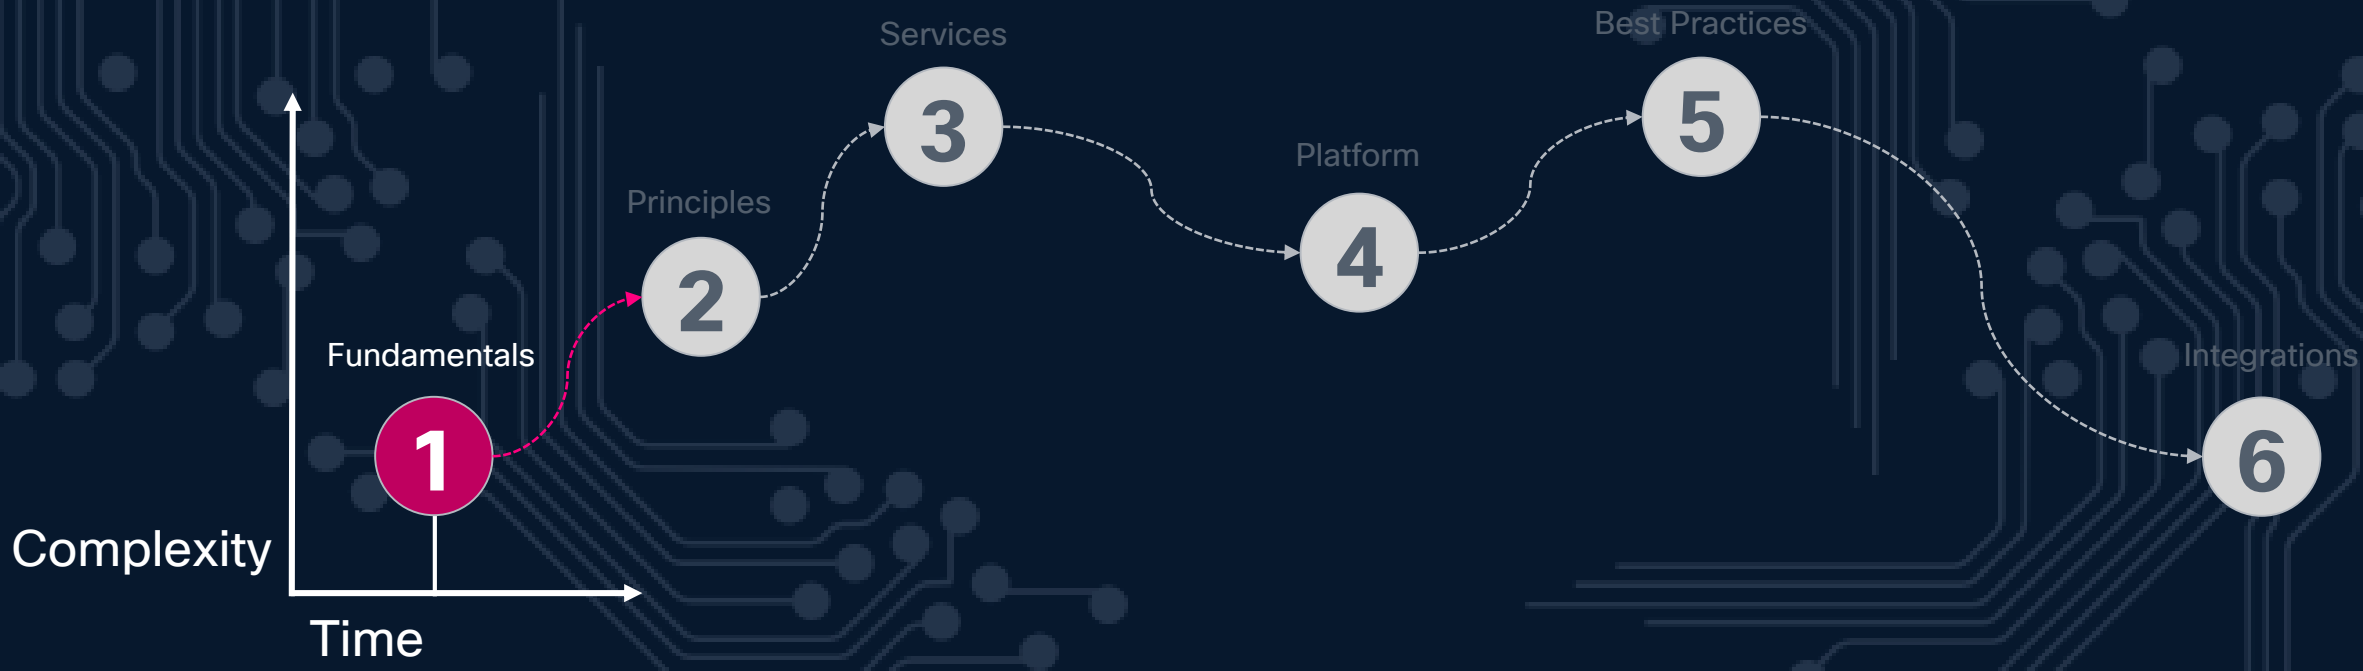

Design Fundamentals

- **What is “Campus”?**
- **Place in Network** (PIN)

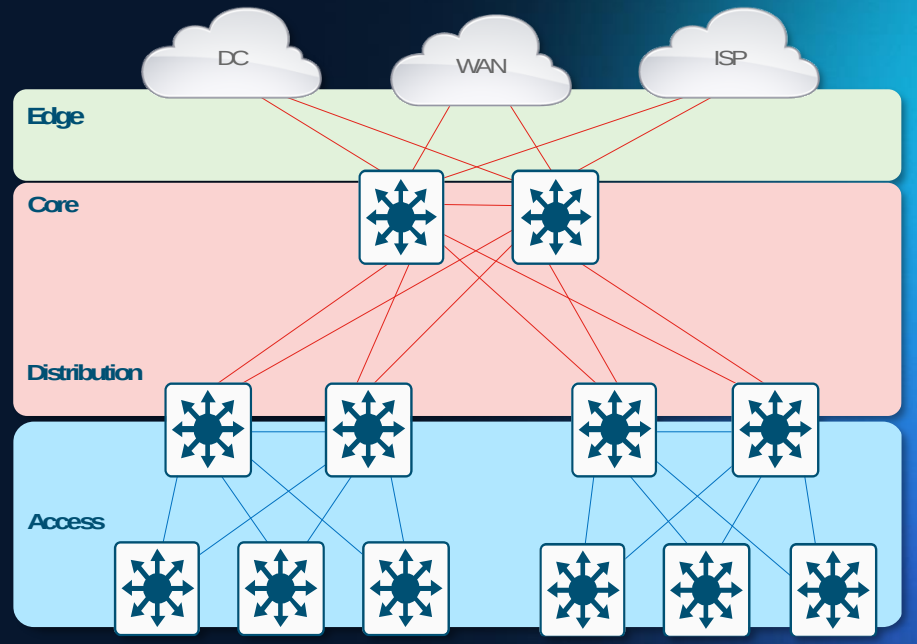

What is a “Campus”?

A basic **Merriam-Webster** definition of a Campus is:

*A group of **one or more buildings**, and surrounding grounds, where **people and their belongings** work together.*

Common examples are **Corporate** & **Government Offices**, **Hospitals**, **Schools**, **Transportation**, **Manufacturing** & more.

Using this – it’s clear a **Campus Network** is focused on:

- **People** (*Users, Vendors, etc.*)
- **People's devices** (*PCs, Phones, Printers, etc.*)
- **Local geographic area** (*LAN, WLAN or MAN, etc.*)
- **Access other domains** (*WAN, ISP, DC & Cloud, etc.*)

This includes **many different network technology areas** (*Wired, Wireless, Security, QoS, Management, etc.*)

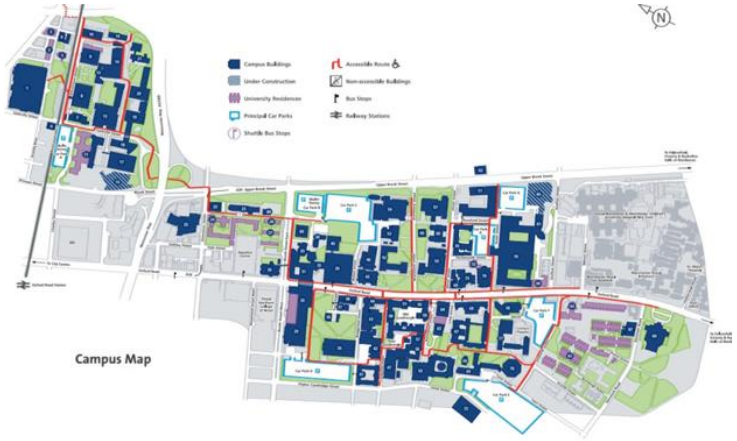

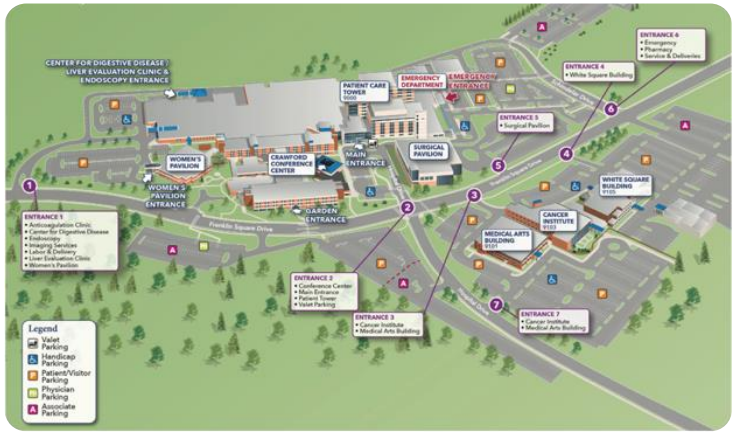

Campus is focused on **User Access**

Campus = Geography

Buildings are spread out. Multiple floors per building

www.cisco.com/c/en/us/solutions/cisco-on-cisco/enterprise-networks.html

Campus Networks

Building MDF/IDF and Wiring Closets

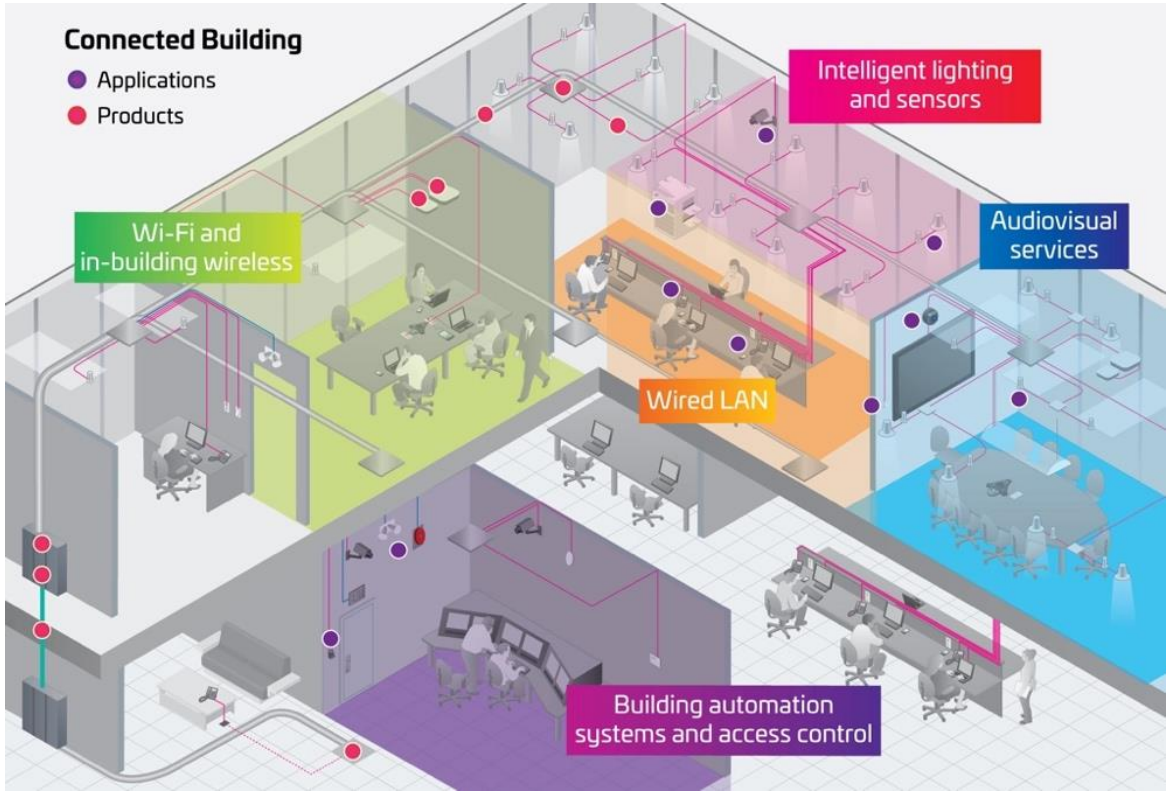

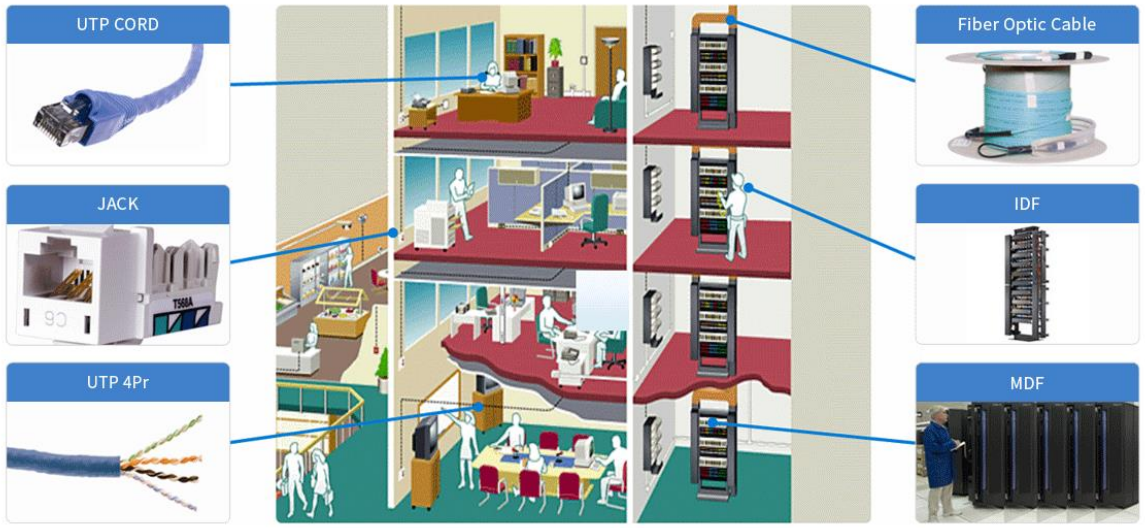

MDF = Main Distribution Framework (Core & Edge)
IDF = Intermediate Distribution Framework (Distro & Access)

www.cisco.com/c/en/us/solutions/design-zone/networking-design-guides/campus-wired-wireless.html

Campus ≠ Data-Center

One or few large buildings nearby. Usually a single floor.

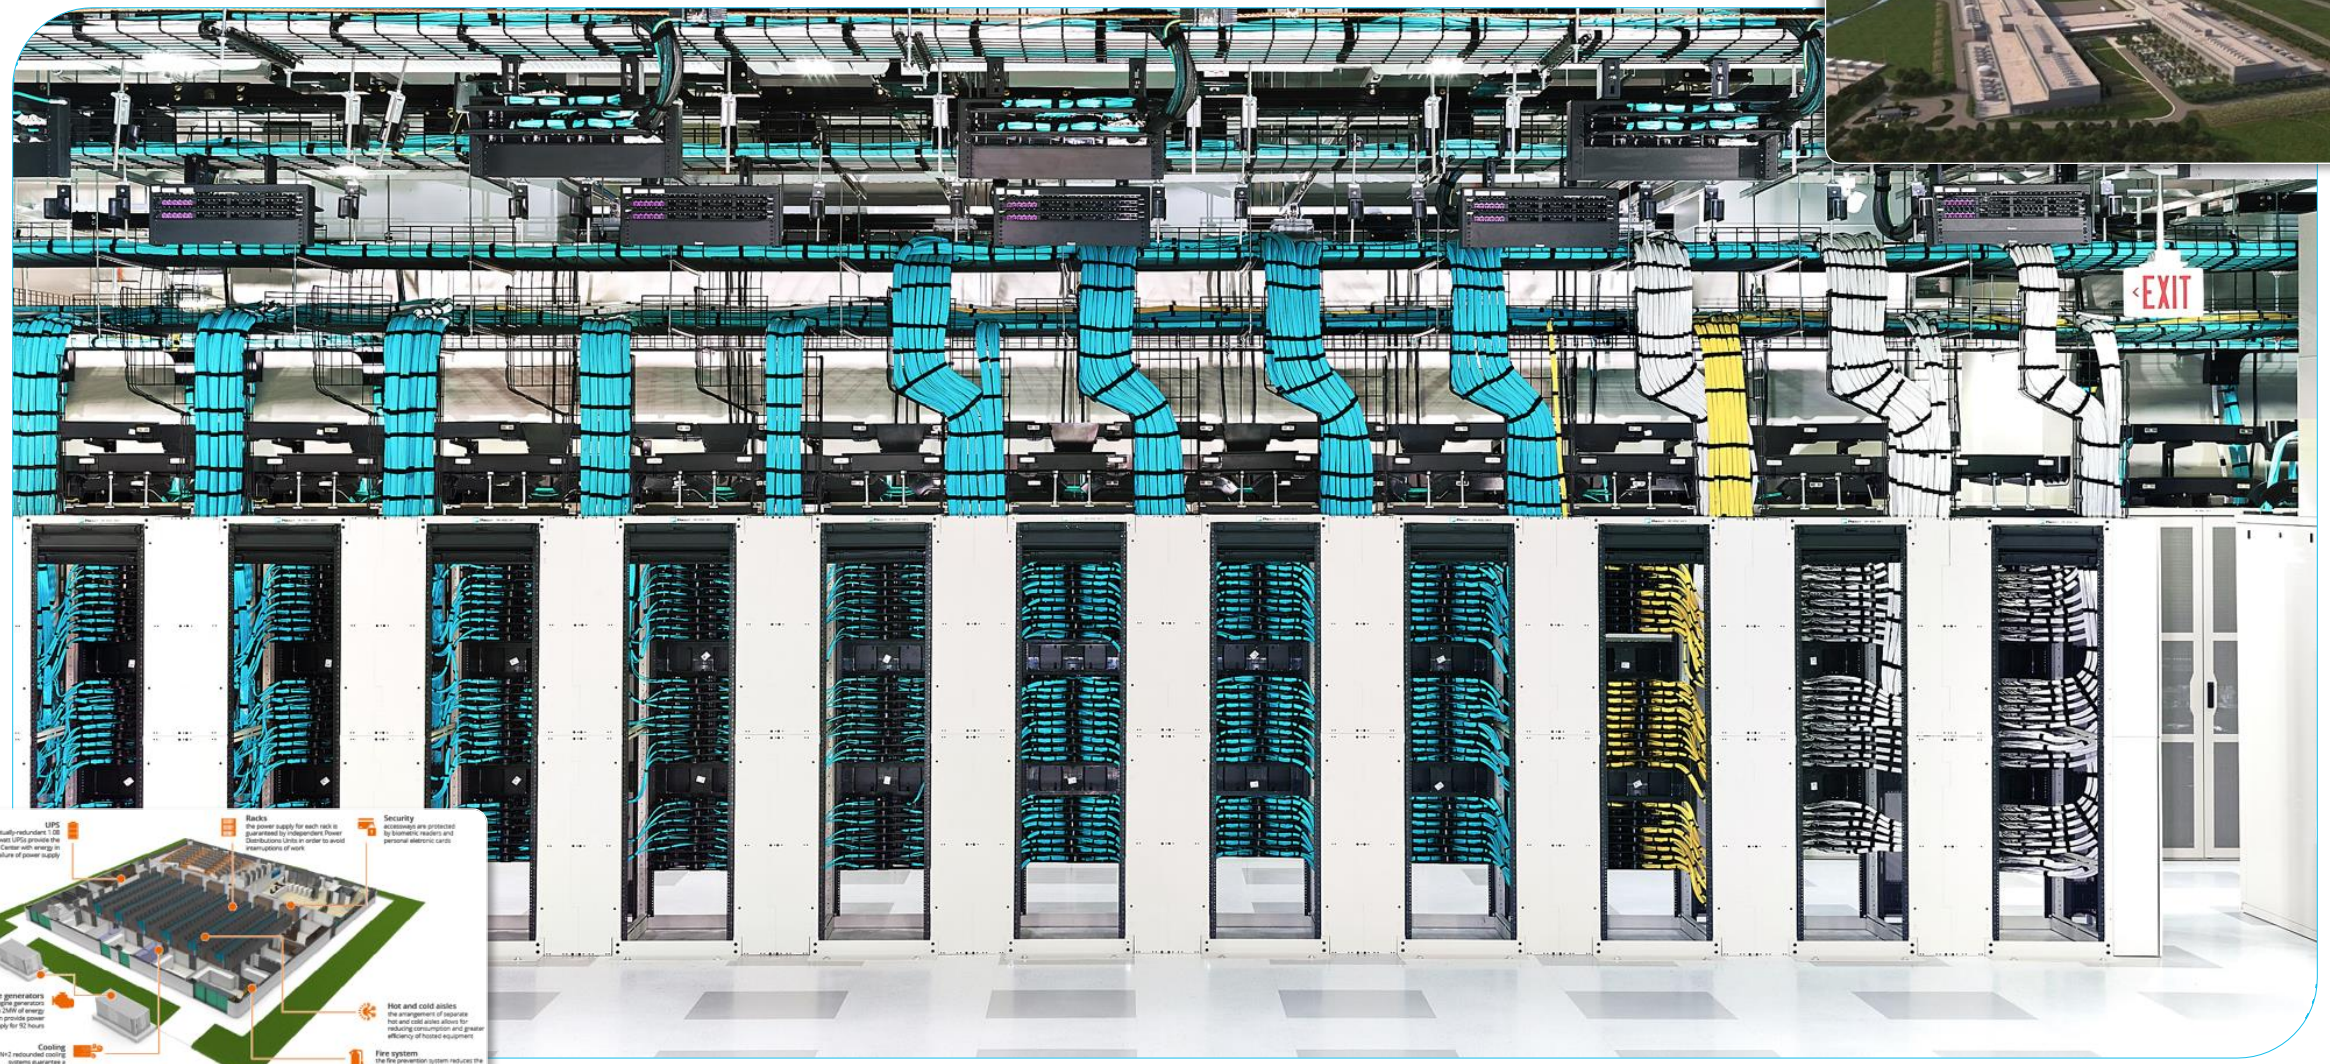

www.cisco.com/c/en/us/solutions/cisco-on-cisco/enterprise-networks.html

Campus Networks - Real Life

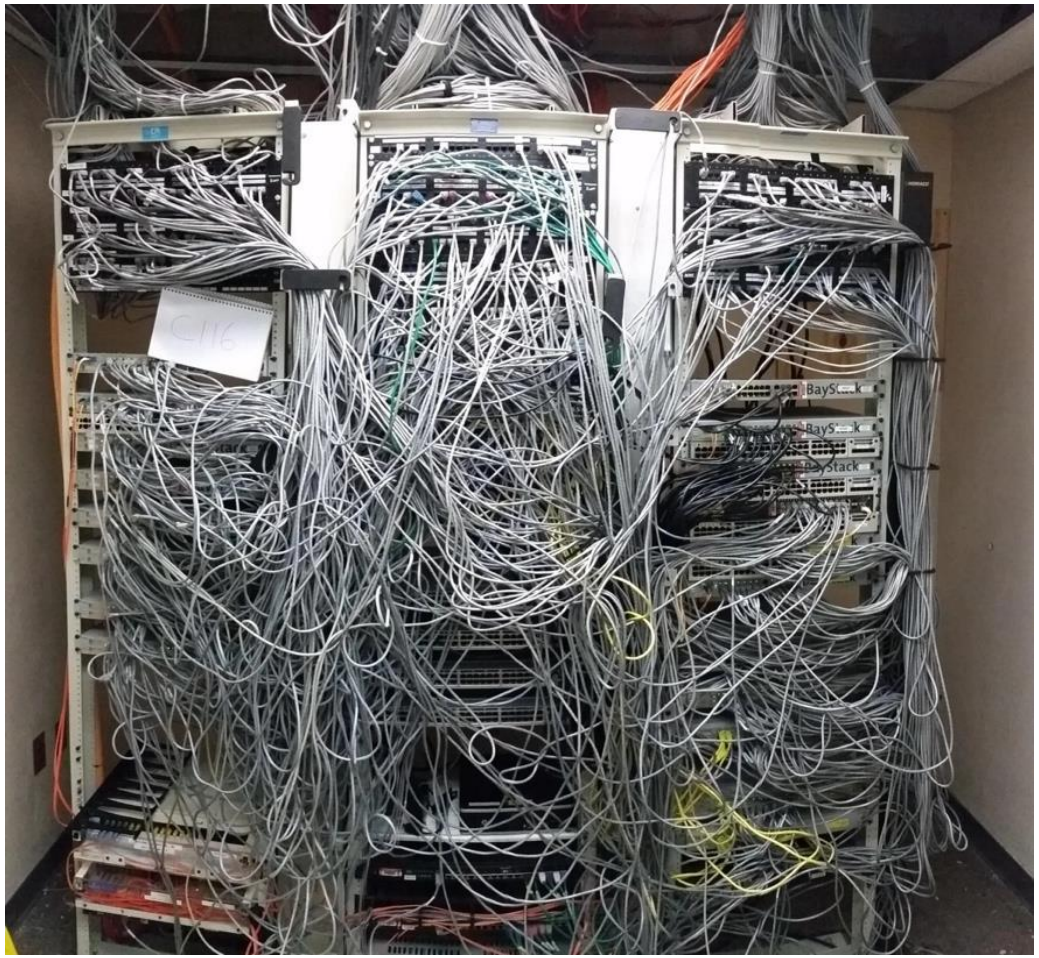

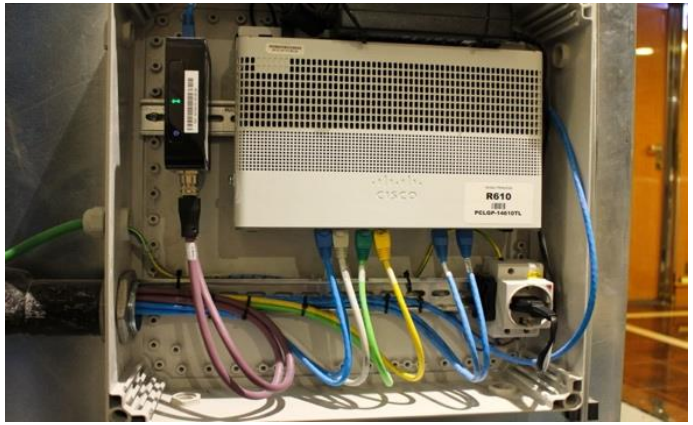

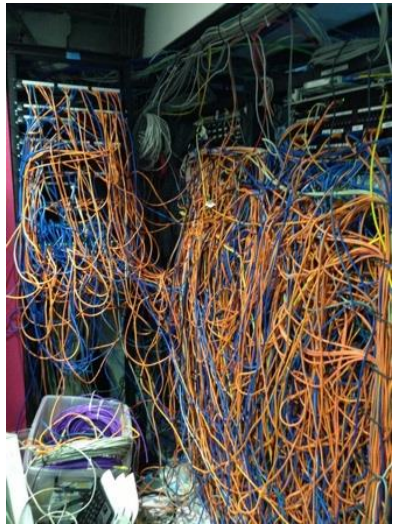

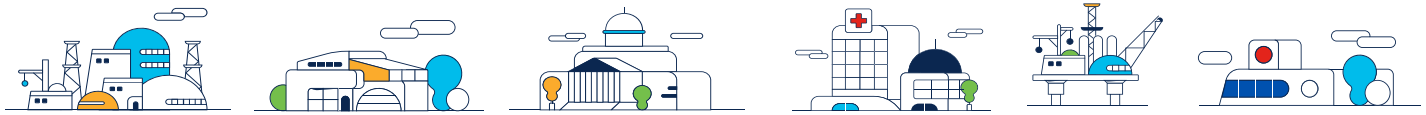

Design Fundamentals

- What is “Campus”?
- Place in Network (PIN)

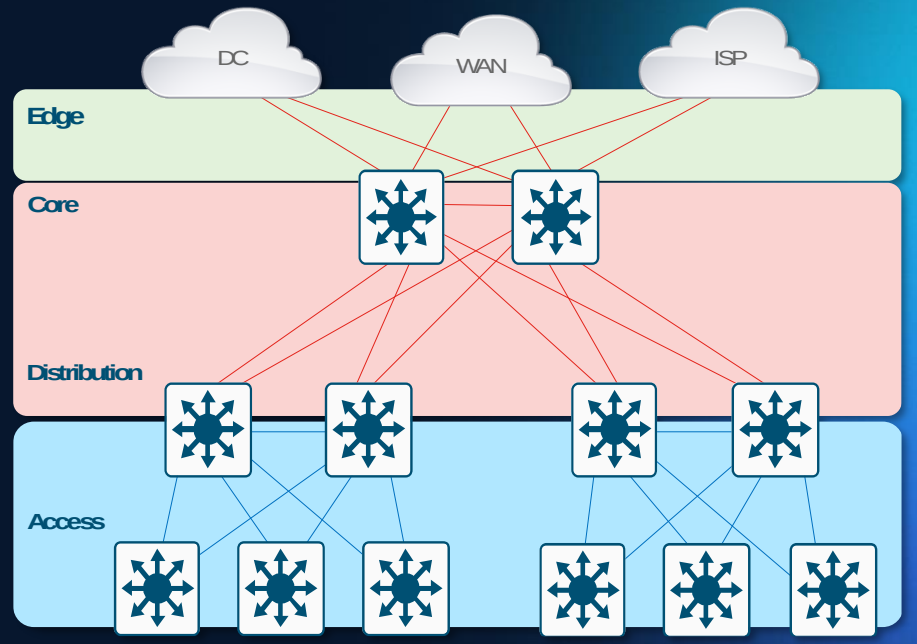

Campus PINs & Topology

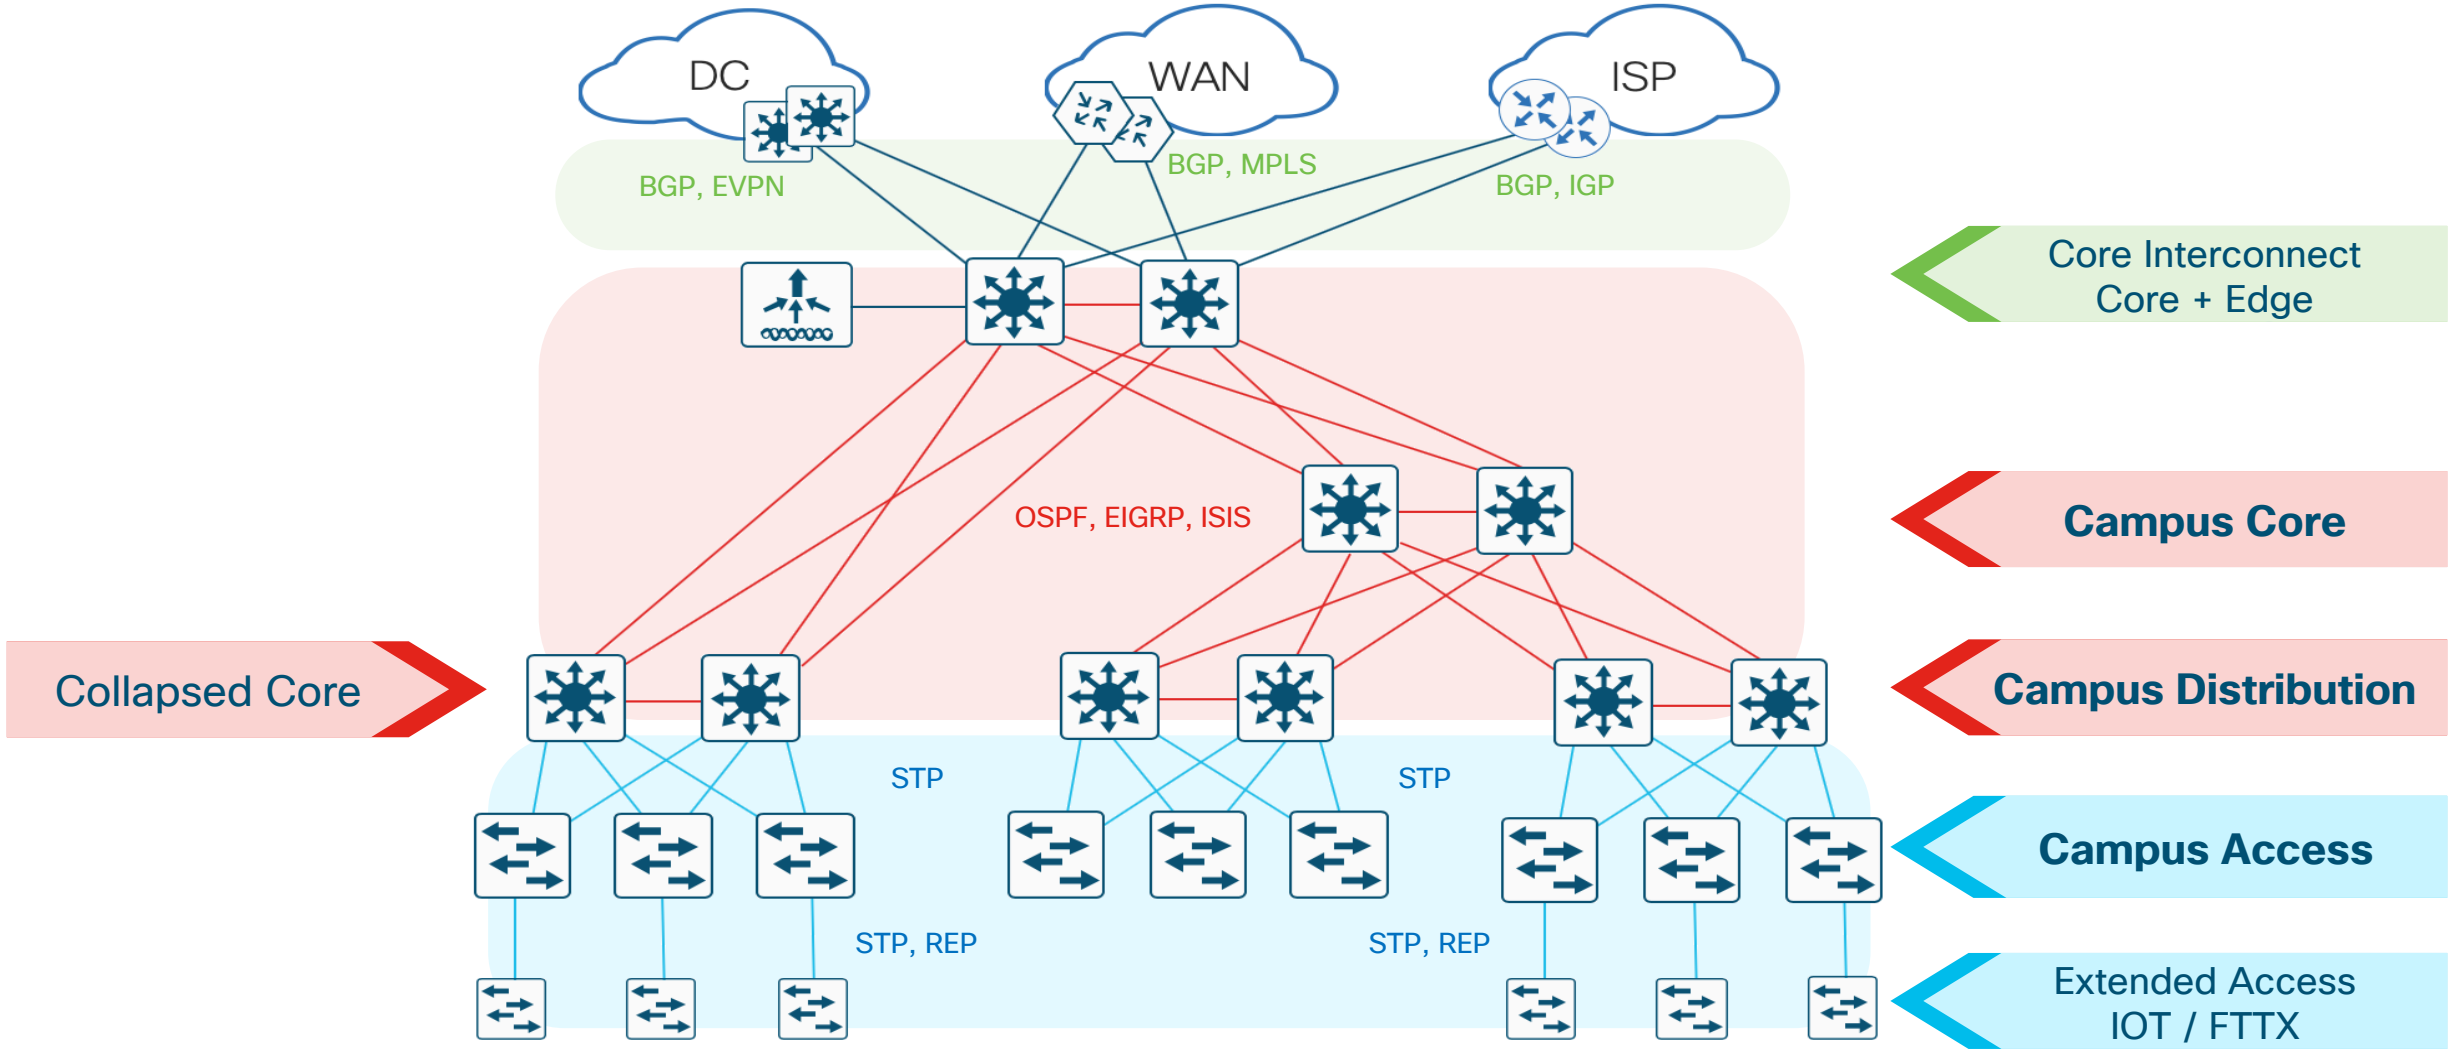

Where do I start?

Cisco Validated Designs

Provide a framework for design and deployment guidance based on common use cases.

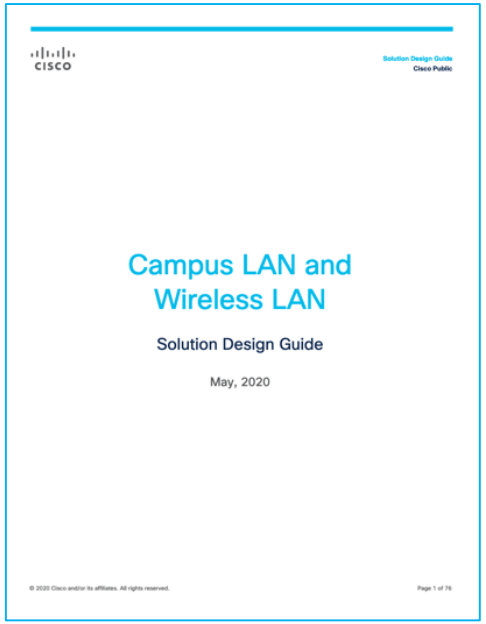

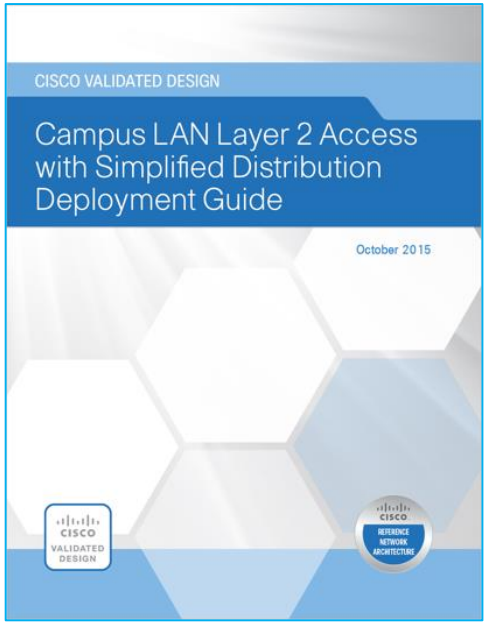

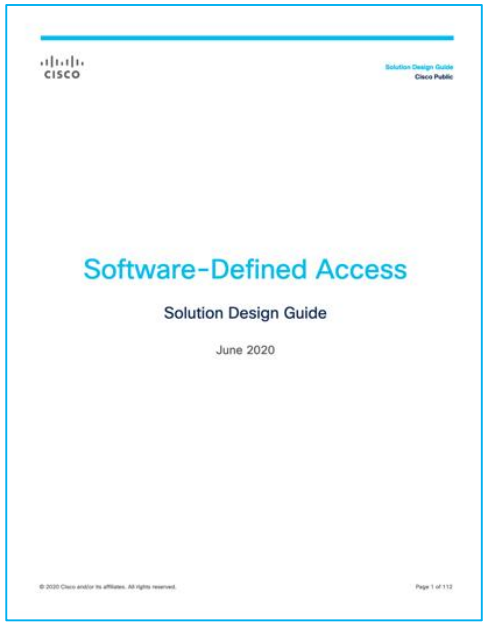

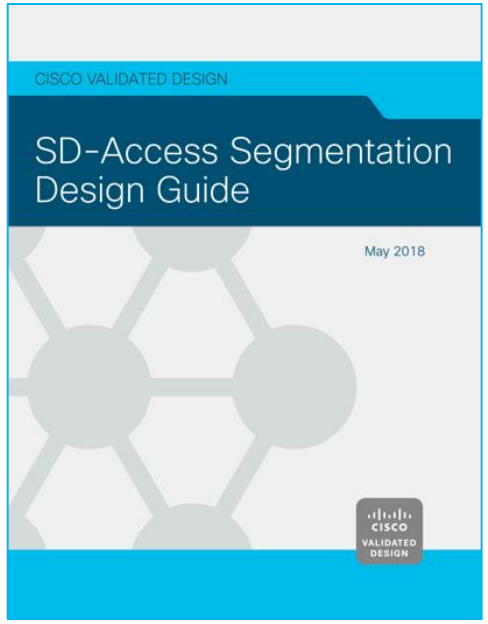

Design Zone: www.cisco.com/go/designzone

Design Zone for Campus: www.cisco.com/go/cvd/campus

Session Agenda

Design Fundamentals | **Design Considerations**

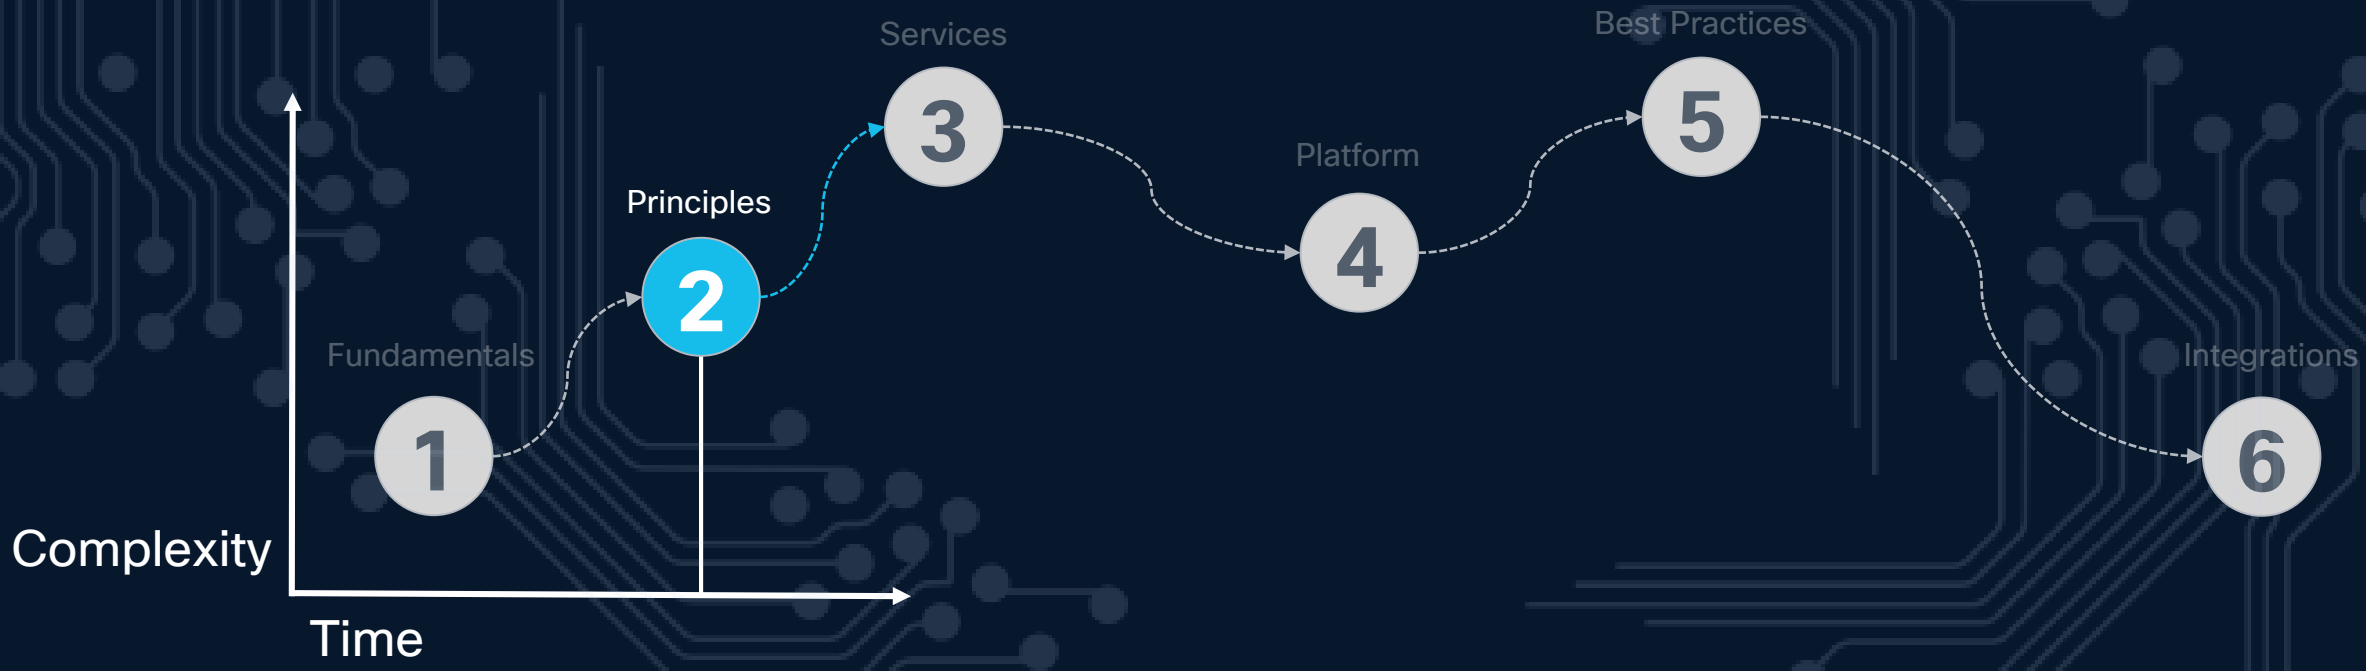

Design Fundamentals

- **Multi-Layer Model**
  - Campus Multi-Layer
  - Hierarchical Design
- **Access Layer**
- **Distribution Layer**
- **Core Layer**

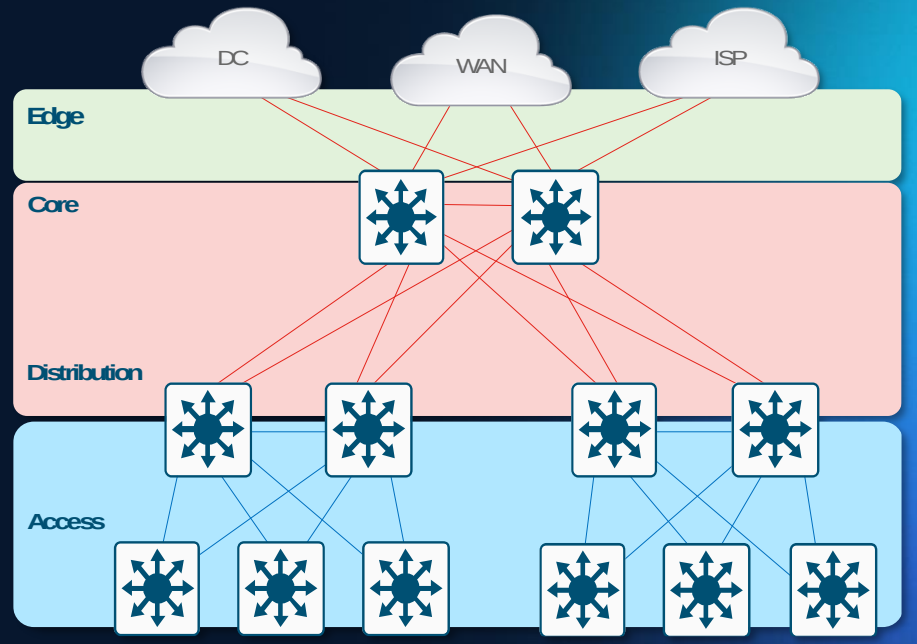

Campus Design Fundamentals

BRKENS-1500

Hierarchical design model – Scalability & Stability

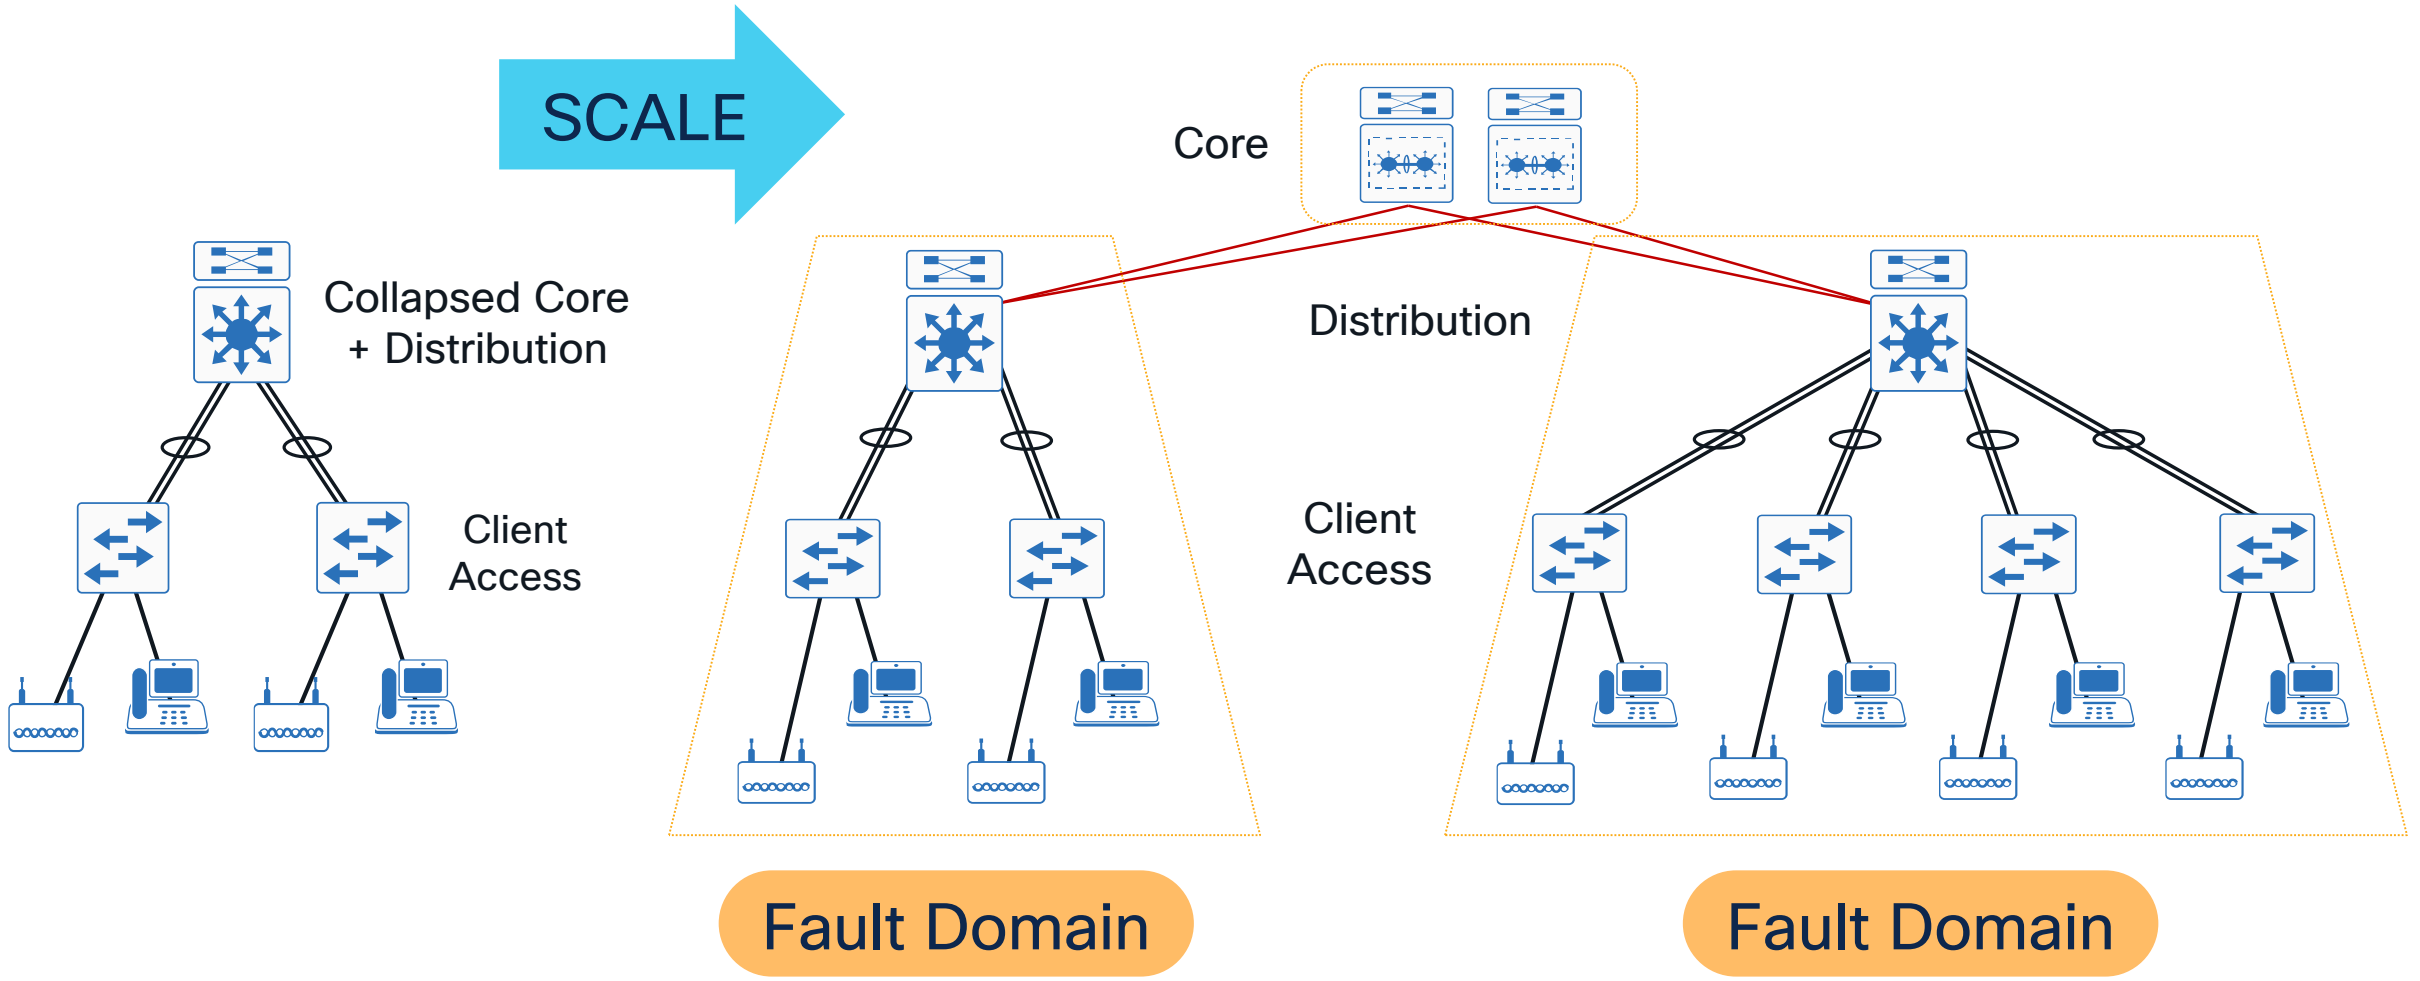

Hierarchical Network Design

BRKENS-1500

Without a Rock-Solid Foundation - the Rest Does not Matter

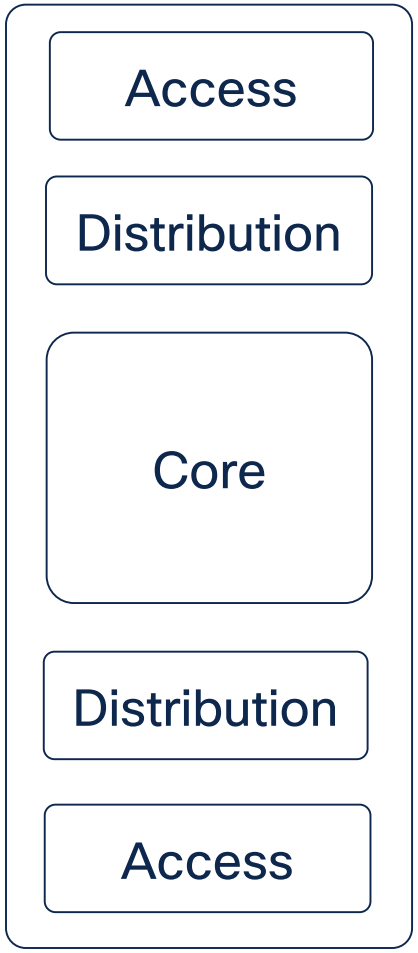

- Offers hierarchy—each layer has specific role
- Modular topology—building blocks
- Easy to grow, understand, and troubleshoot
- Creates small fault domains— clear demarcations and isolation
- Promotes load balancing and redundancy
- Promotes deterministic traffic patterns
- Incorporates balance of both Layer 2 and Layer 3 technology, leveraging the strength of both
- Utilizes Layer 3 routing for load balancing, fast convergence, scalability, and control

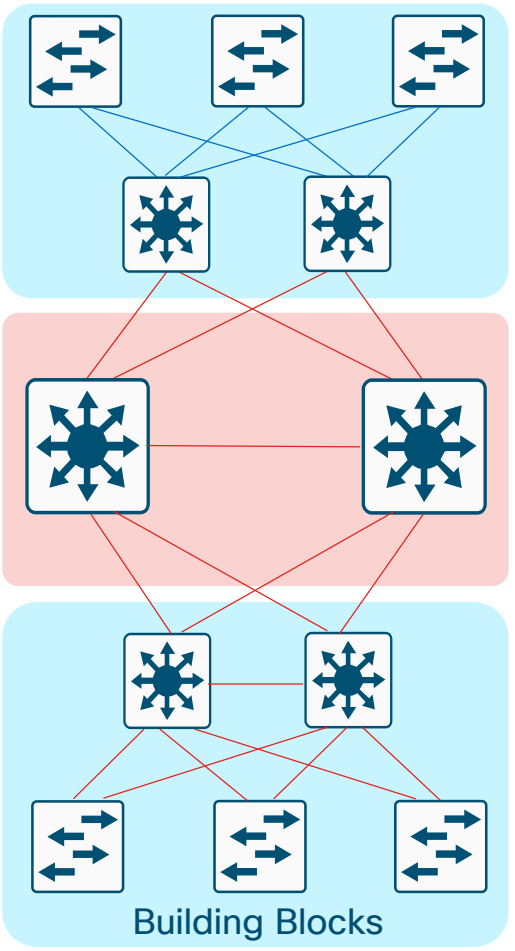

Alternative Designs in Multilayer architecture

Access

Distribution

Core

Distribution

Access

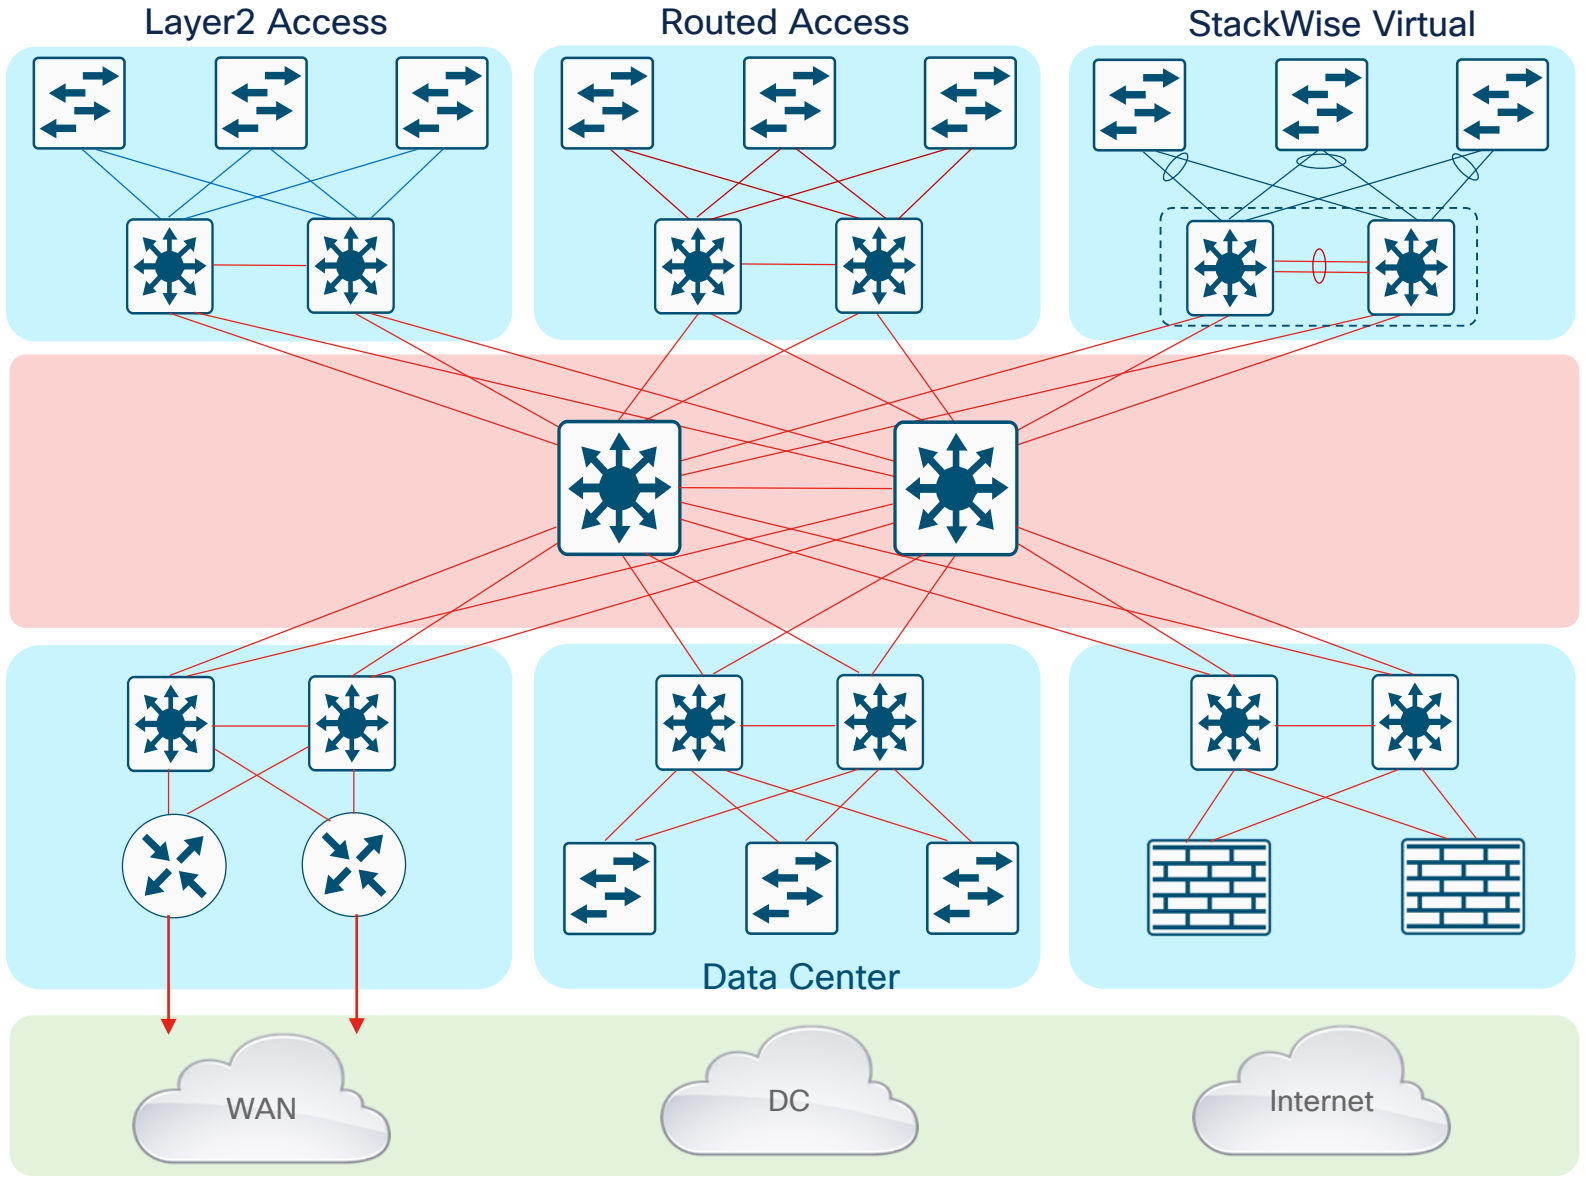

Multilayer Architecture DON'Ts

Do NOT build this!!

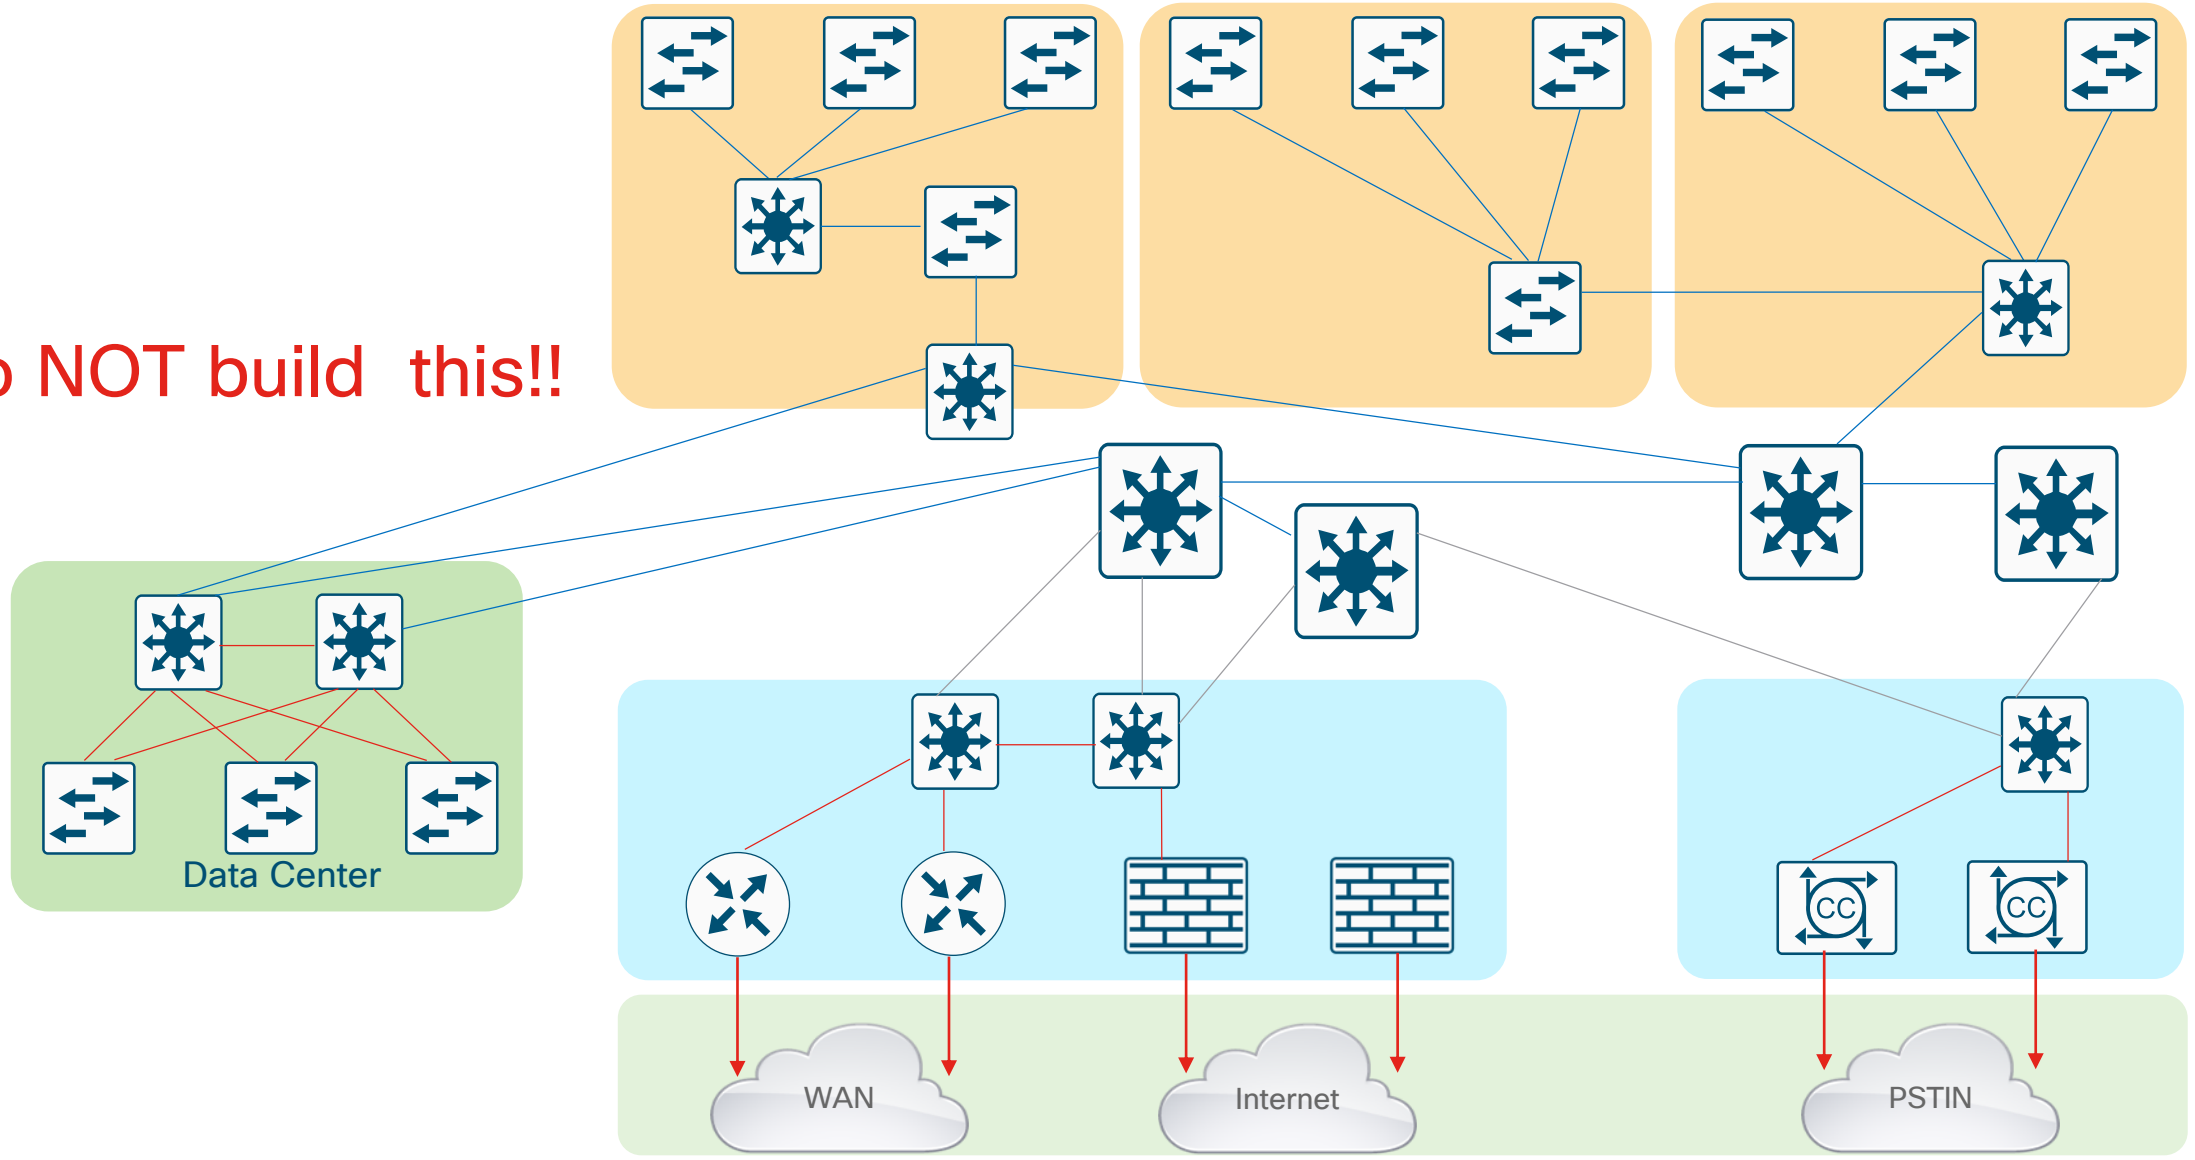

Design Fundamentals

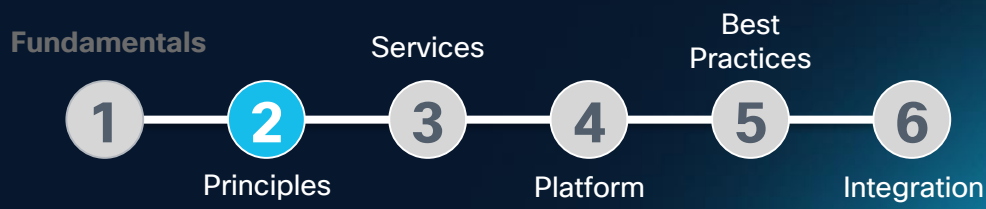

- Multi-Layer Model
- Access Layer
  - Campus Access (Baseline)
  - Extended Access (IOT / FTTX)
  - Routed Access
- Distribution Layer
- Core Layer

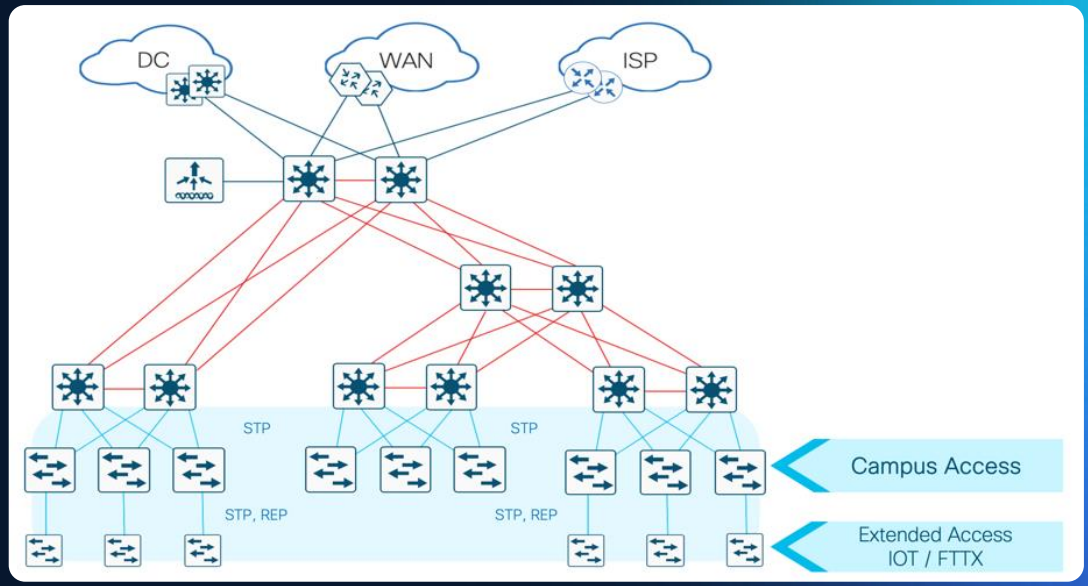

Campus Access (Baseline)

The **Access PIN** (Tier 1) focuses on connecting Users & Devices, or an Extended Access (if applicable), to the Distribution layer

- Other names: IDF, Wiring Closet
- Common in all Campus & Branch networks

Main purpose is to **connect users** to network

Tends to be **L2 switched (north & south)**

- North: VLAN, 802.1Q, STP, MAC, IGMP Snooping
- South: AAA, STP, Portfast, Storm-Control

Tends to use **multiple L2 features & services**

- Access Security (e.g. 802.1x, VACLs, PACLs, etc)
- Access QoS (e.g. L2 CoS, Classification & Marking)
- Access NetFlow (e.g. AVC, FNF, EPA & ETA)

Tends to require **low-med L2 & feature scale**

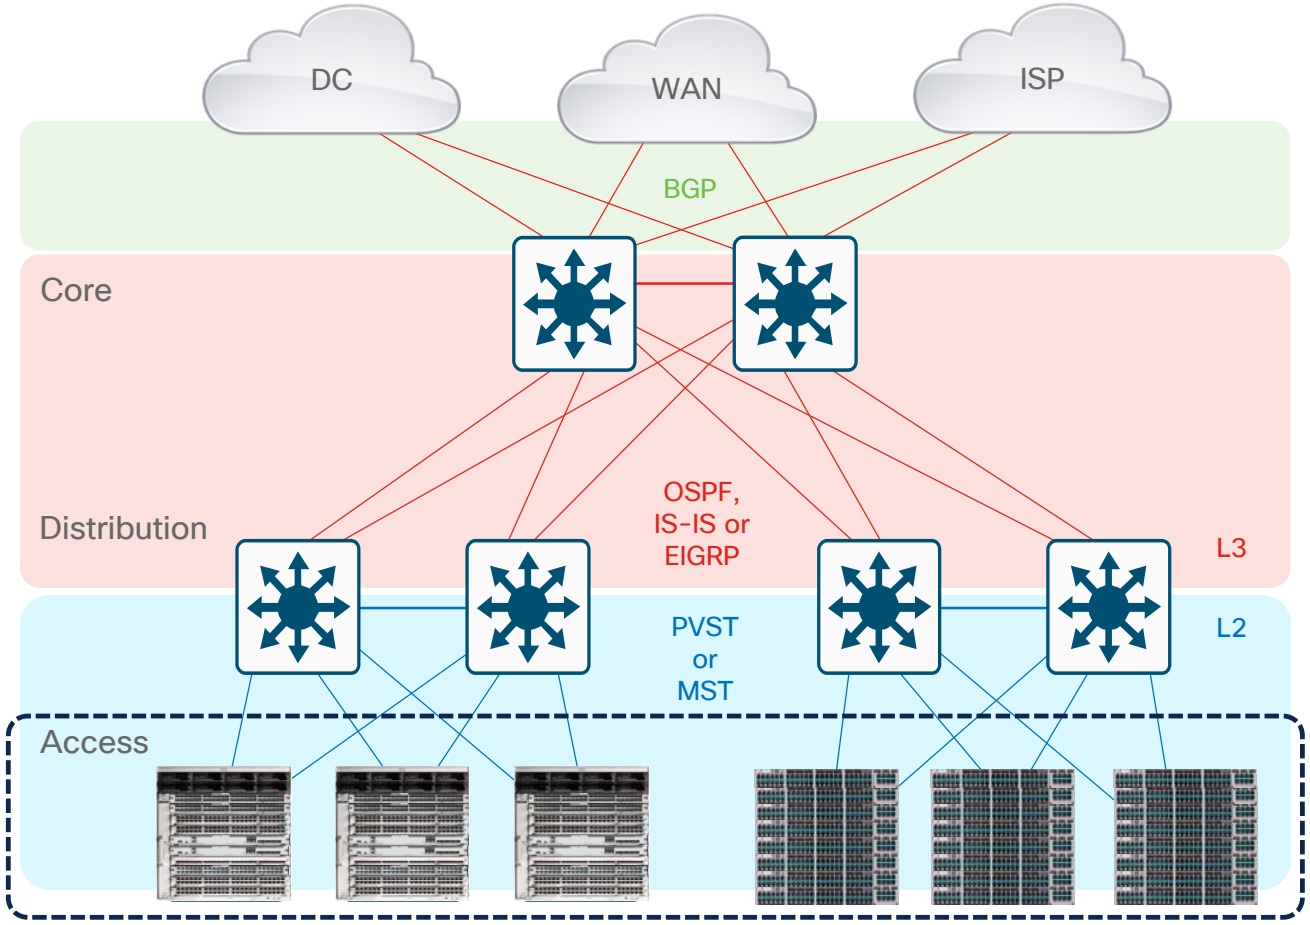

Campus Design Fundamentals

Access Layer – Oversubscription Ratios

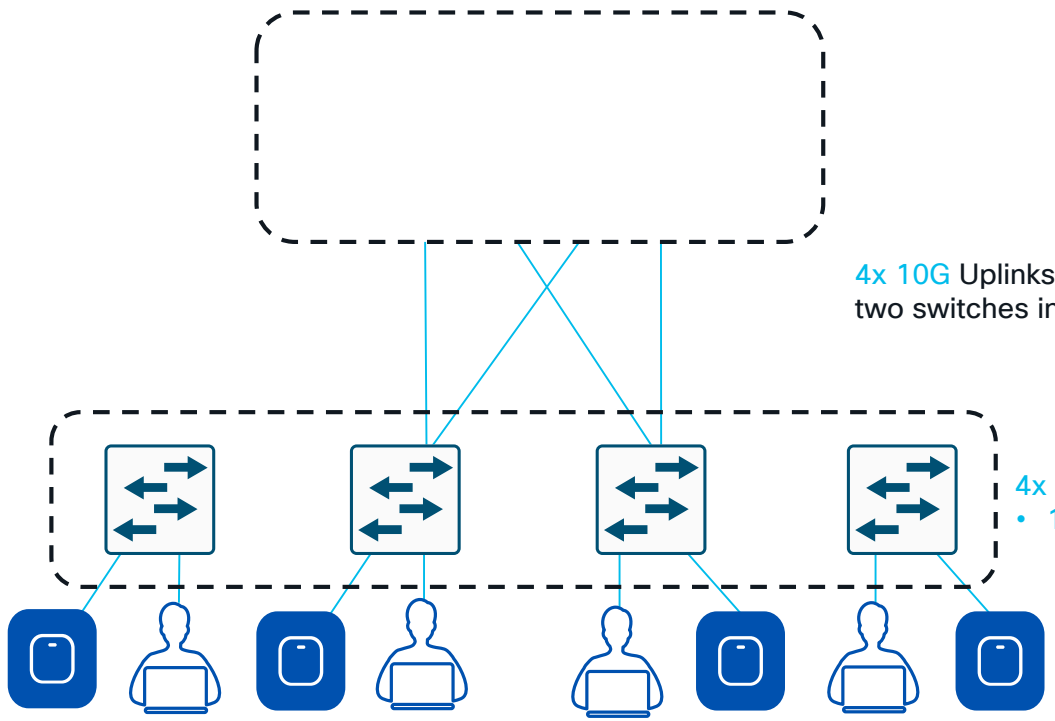

Soft recommendation for Access to Distribution ≤ 20:1

Access Uplinks: **40 Gbps**

Potential Downlinks:

48 x 10 Gbps

\+ **144 x 1 Gbps**

SUM: **624 Gbps**

Oversubcription ratio:

~15.6 : 1

Extended Access (IOT / FTTX)

The **Extended Access PIN** (Tier 1) is an extension of the Access, to connect multiple Access layers (areas) to the Distribution layer

- Other names: High-End Access, IOT, FTTX
- Common in Very-Large Campus or Large Branch

Main goal is to **extend** the size and scale of the Access layer and connect more hosts

Tends to be **L2 switched (north & south)**

- North: VLAN, 802.1Q, STP/REP, MAC, IGMP Snooping
- South: AAA, STP/REP, Portfast, Storm-Control

Tends to use **multiple L2 features & services**

- Access Security (e.g. 802.1x, VACLs, PACLs, etc)
- Access QoS (e.g. L2 CoS, Classification & Marking)
- Access NetFlow (e.g. AVC, FNF, EPA & ETA)

Tends to require **med-high L2 & feature scale**

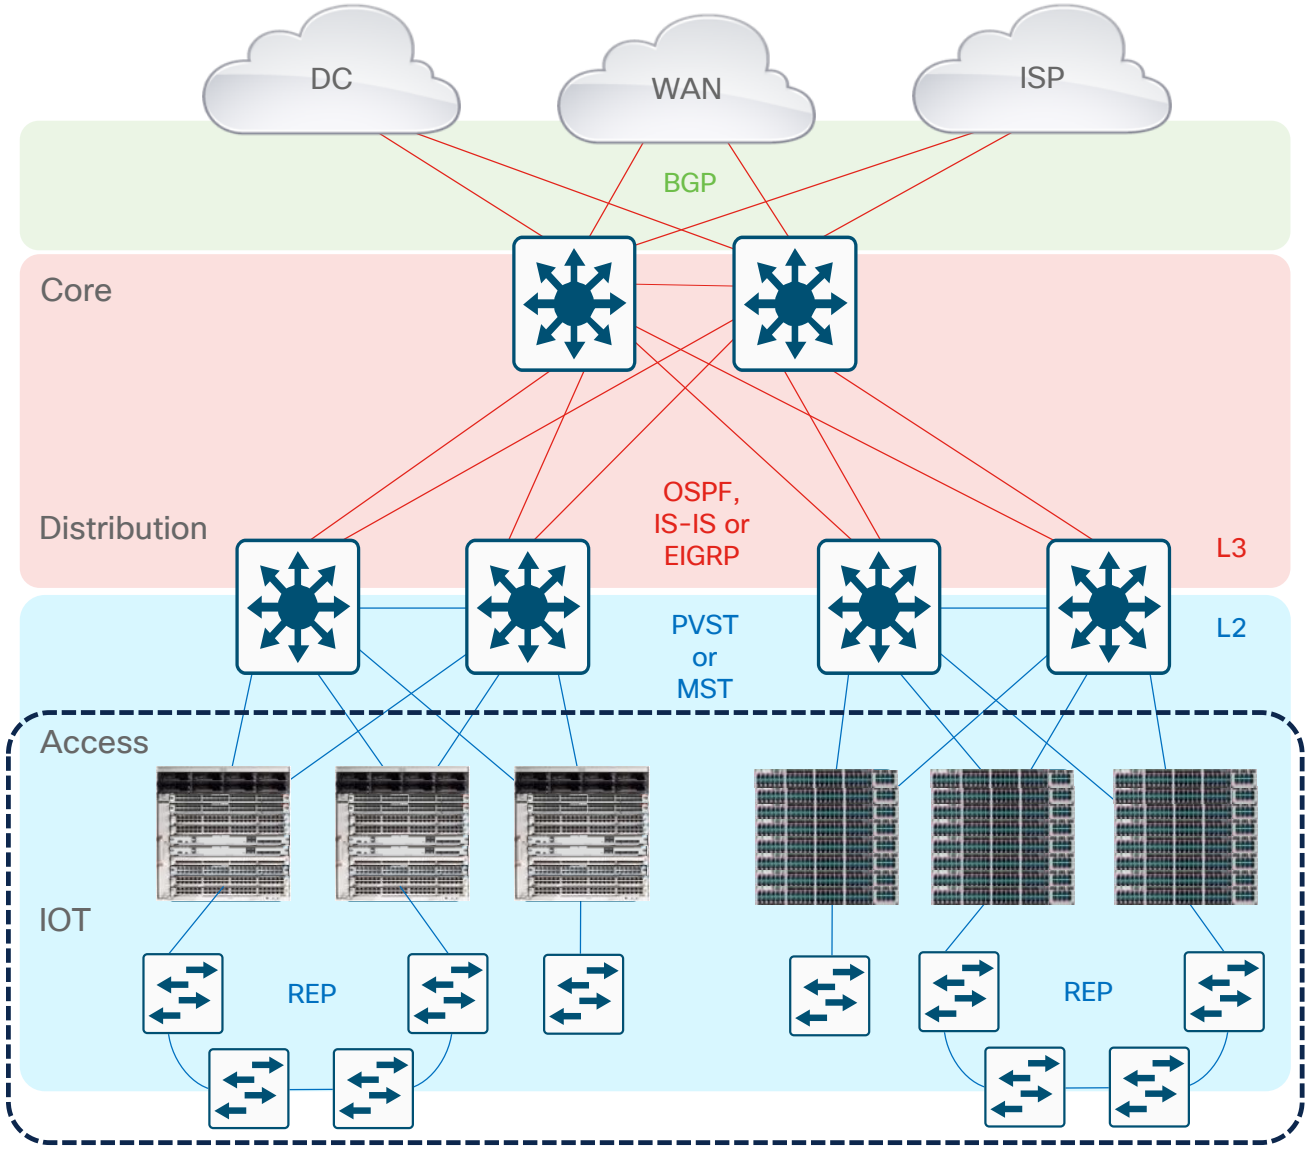

Fiber To The Active Consolidation Point

EcoFlex'ITTM

Upgrading Fiber/Ethernet cabling can be costly

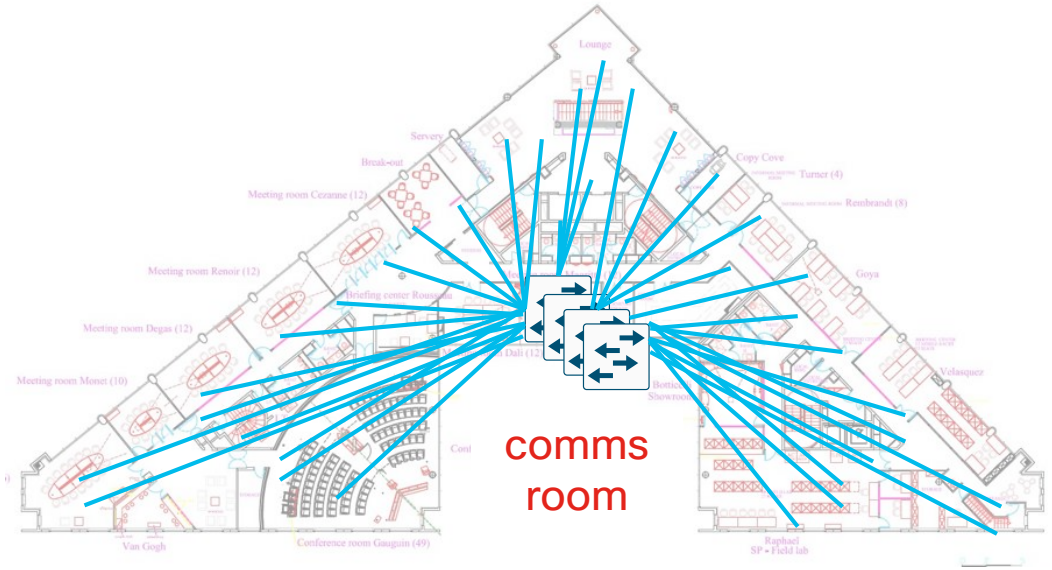

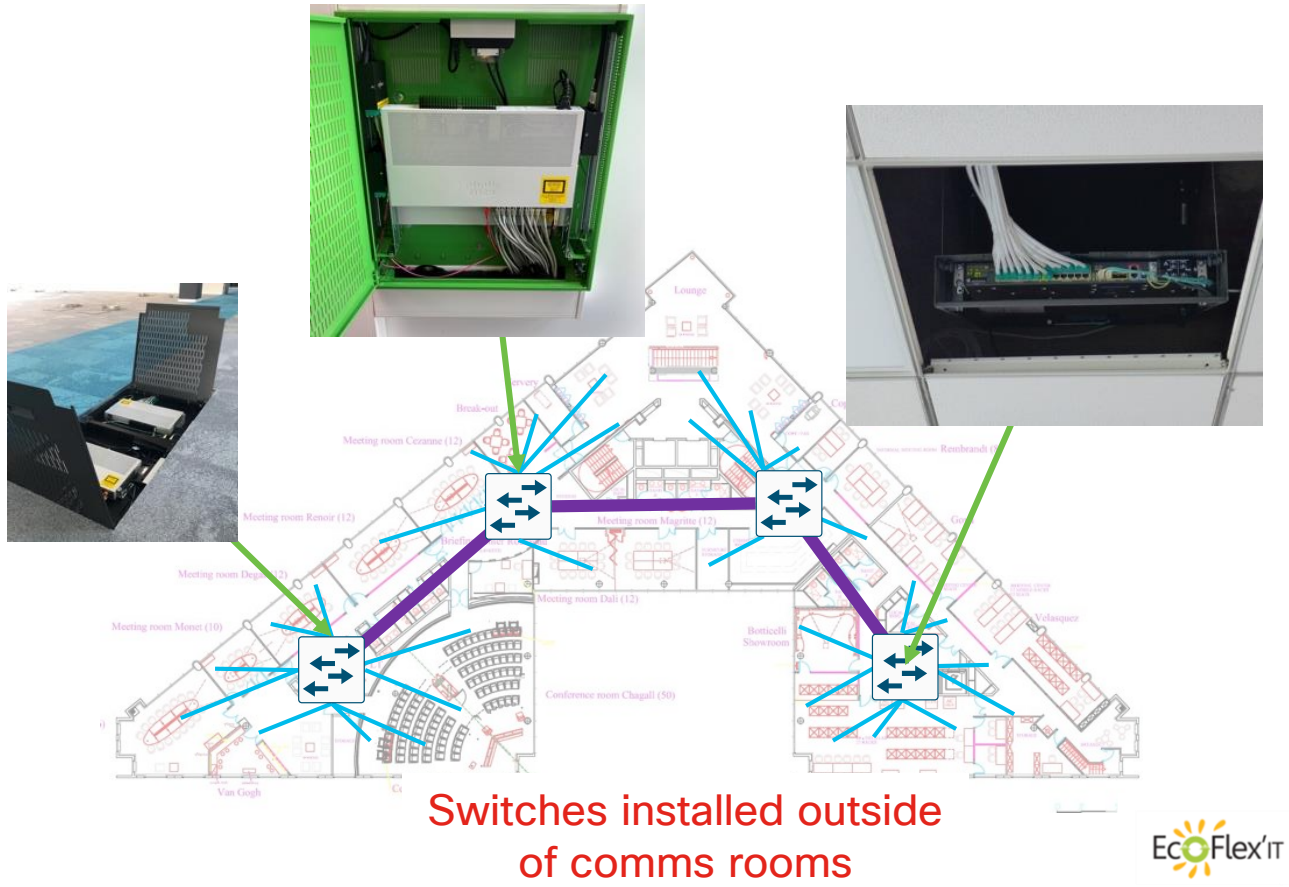

Traditional deployment

FTTACP-EcoFlex'IT™ deployment

https://osi.rosenberger.com/fileadmin/content/osi/EN/News/Whitepaper/Rosenberger_OSI_Whitepaper_FTT-ACP_EN.pdf

Routed Access

The **Routed Access PIN** (Tier 1) has the same purpose, but uses L3 IP routing to limit L2 scale

- Other names: IDF, Wiring Closet
- Semi-common in Campus & Branch networks

Main purpose is to connect users to network using L3 protocols to **reduce L2 challenges**.

- Mostly for network stability and simplicity of protocols
- Similar attributes & requirements as Distribution

Tends to be **both L3 routed (north) and L2 switched (south)**

- North: SVI, HSRP/VRRP, ARP/ND, IGP, PIM
- South: VLAN, AAA, MAC, IGMP, STP Portfast

Tends to use **multiple L2 & L3 features**

- Access Security (e.g. IPDT/SISF, VACLs, PACLs, etc)
- Access QoS (e.g. NBAR, Classification & Marking)
- Access NetFlow (e.g. AVC, FNF, EPA & ETA)

Tends to require **low-med L2 & L3 feature scale**

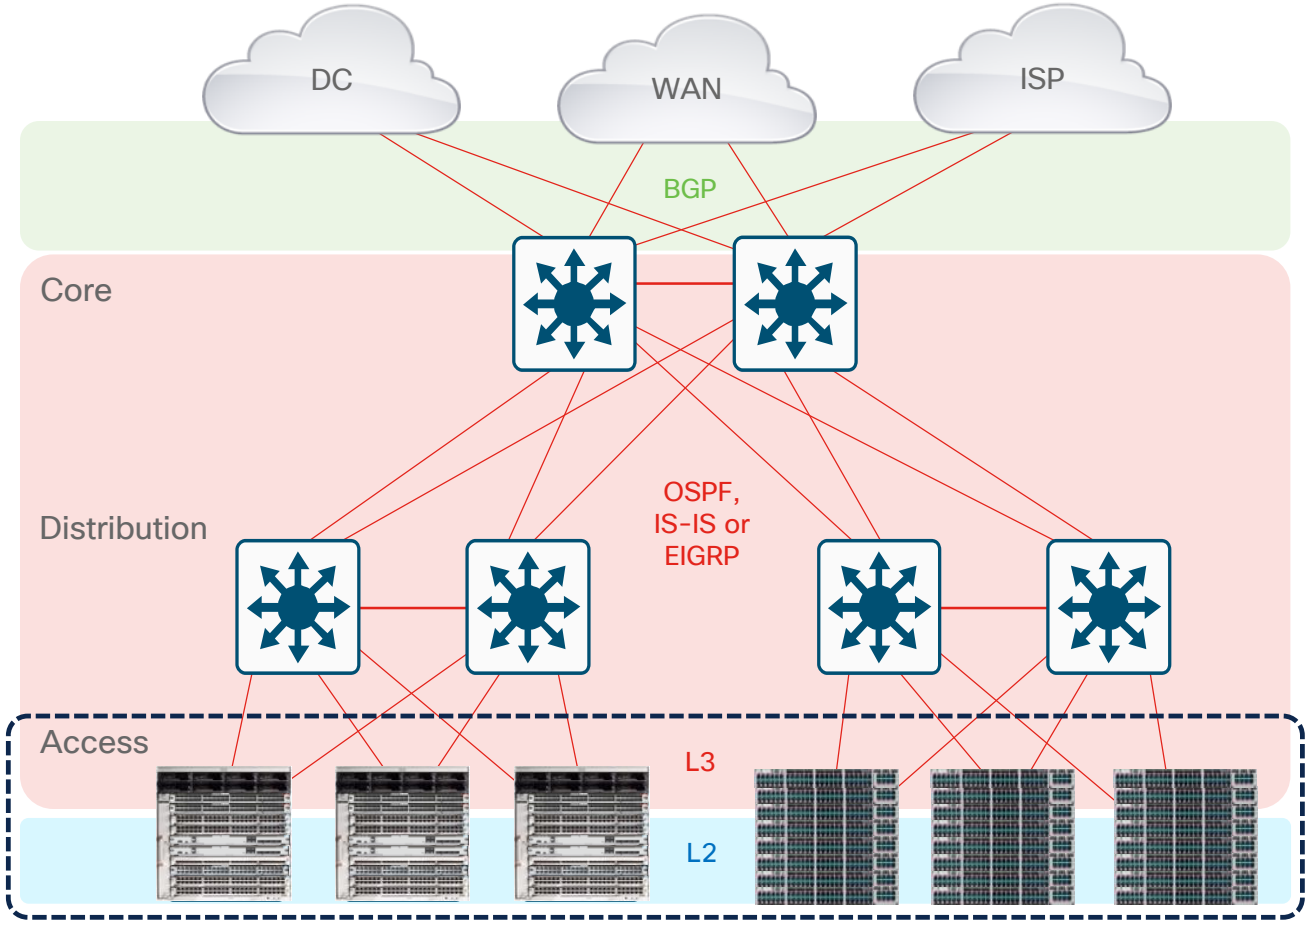

Routing in the Access

Pros:

- Improved network convergence
- Simplified multicast configuration
- Dynamic traffic load balancing
- Single set of troubleshooting (e.g. ping)
- Ease migration towards SDA/EVPN

Cons:

- Different set of VLANs on different access switches
- Lower flexibility
- License for routing (Essentials might not be enough)

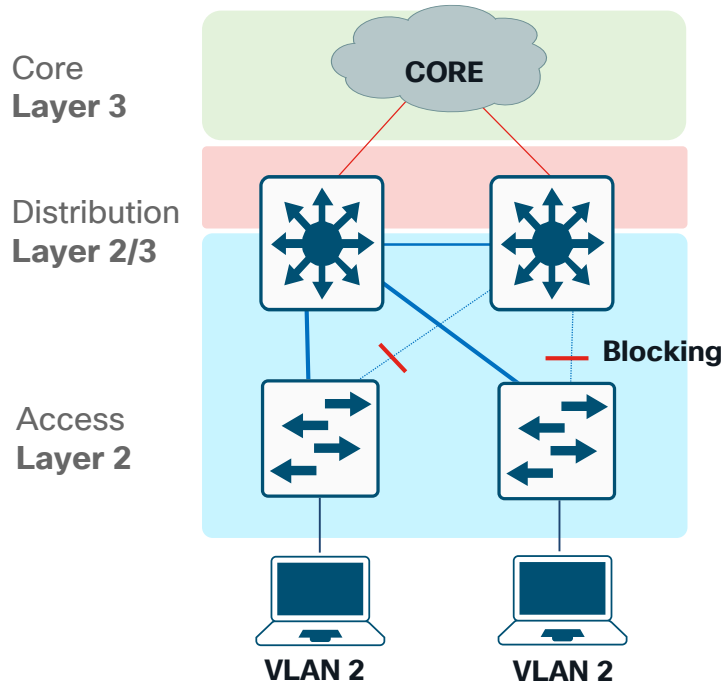

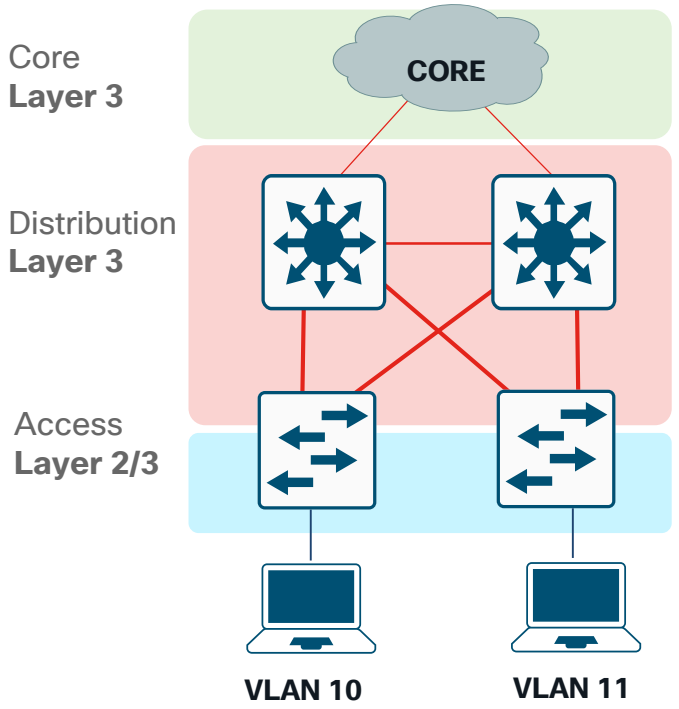

The ability to reduce convergence times to **sub-200** msec range

Design Fundamentals

- Multi-Layer Model
- Access Layer
- Distribution Layer
  - Campus Distribution (Baseline)
  - Collapsed Core + Distribution
  - Collapsed Distro + Ext. Access
- Core Layer

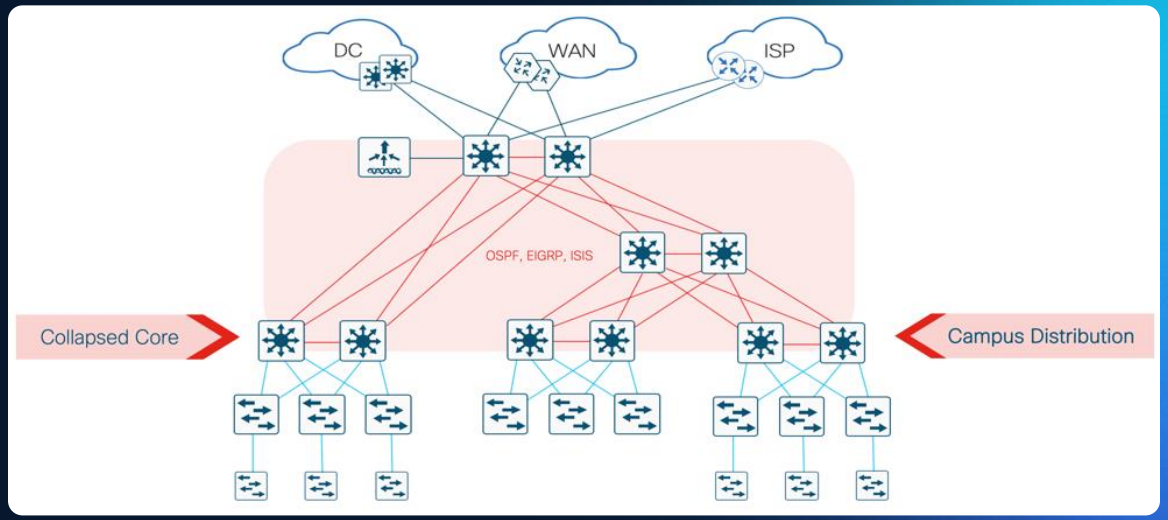

Campus Distribution (Baseline)

The **Distribution PIN** (Tier 2) focuses on connecting multiple Access layers and the Core layer.

- Other names: Collapsed Core, Aggregation, IDF
- Common in Small to Large Campus

Main purpose is to **distribute** connectivity (fan-out) from the Core/WAN to the Access

- Reduces need for high port-density in Core layer
- Also applicable to L3 Routed Access

Tends to be **both L3 routed (north) and L2 switched (south)**

- North: SVI, HSRP/VRRP, ARP/ND, IGP, PIM
- South: VLAN, 802.1Q, STP, MAC, IGMP

Tends to use **multiple L2 & L3 features**

- Access Security (e.g. IPDT/SISF, VACLs, PACLs, etc)
- Access QoS (e.g. NBAR, Classification & Marking)
- Access NetFlow (e.g. AVC, FNF, EPA & ETA)

Tends to require **med-high L2/L3 & feature scale**

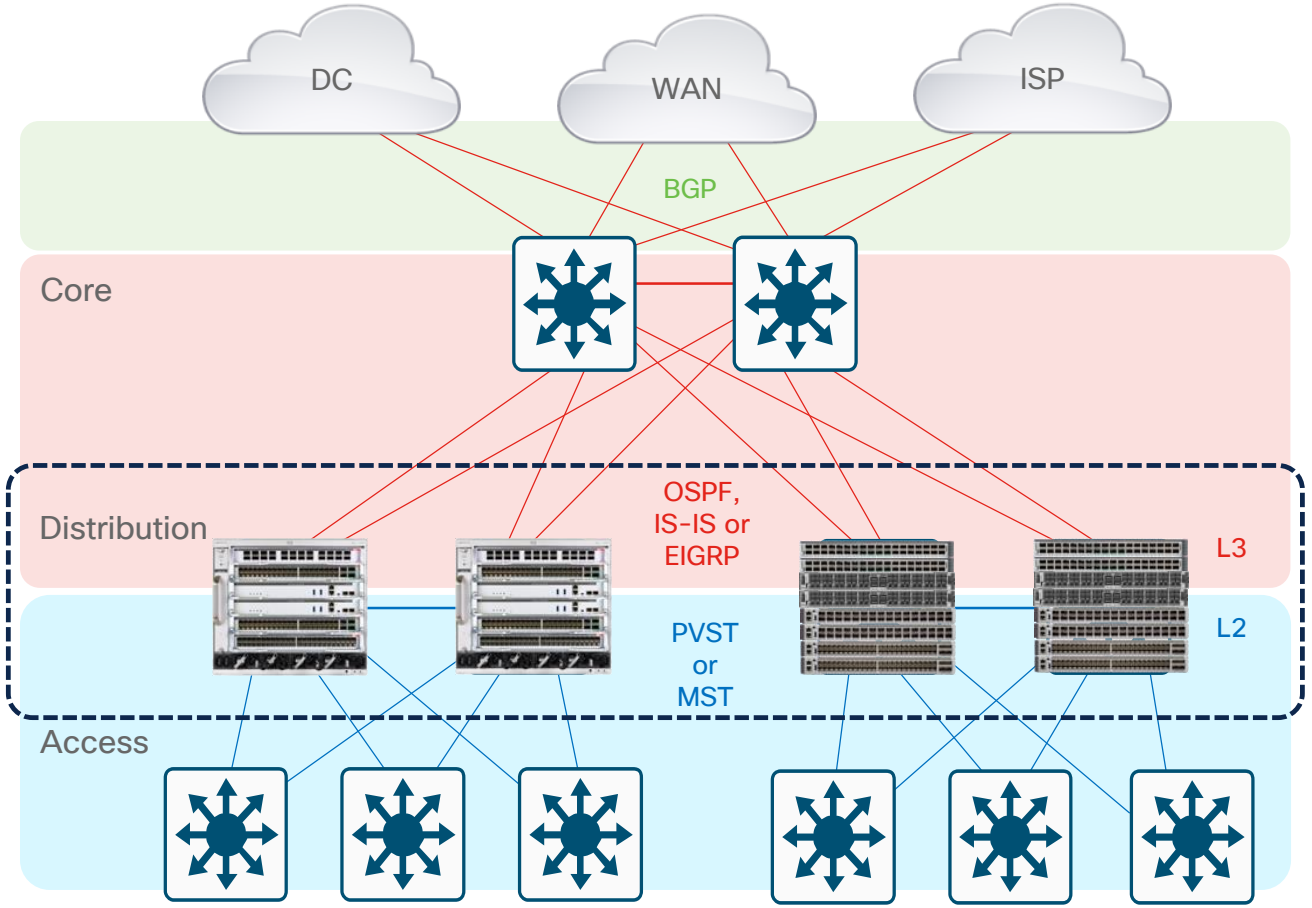

Design Fundamentals

Distribution Layer – Different setups

Two-tier remote site:

- Aggregates LAN Access Layer and connects to WAN routers

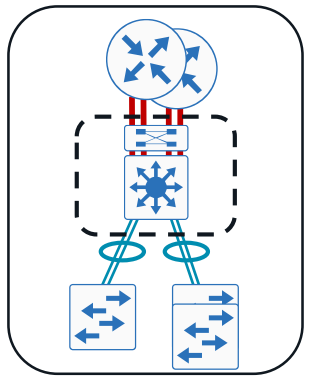

Collapsed Core:

Two tier campus LAN and WAN Core

- LAN Access Layer aggregation
- Central connect point for all services

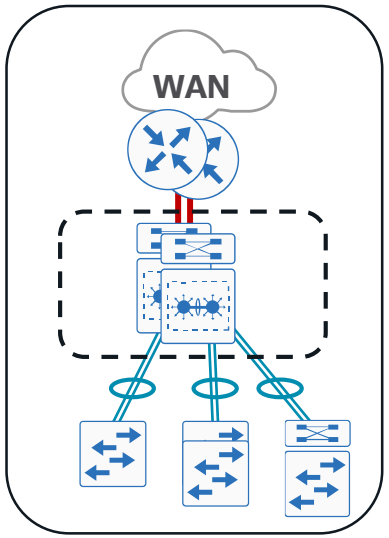

Large LAN Services Block:

- Connection point for services
- Drives modular building block design

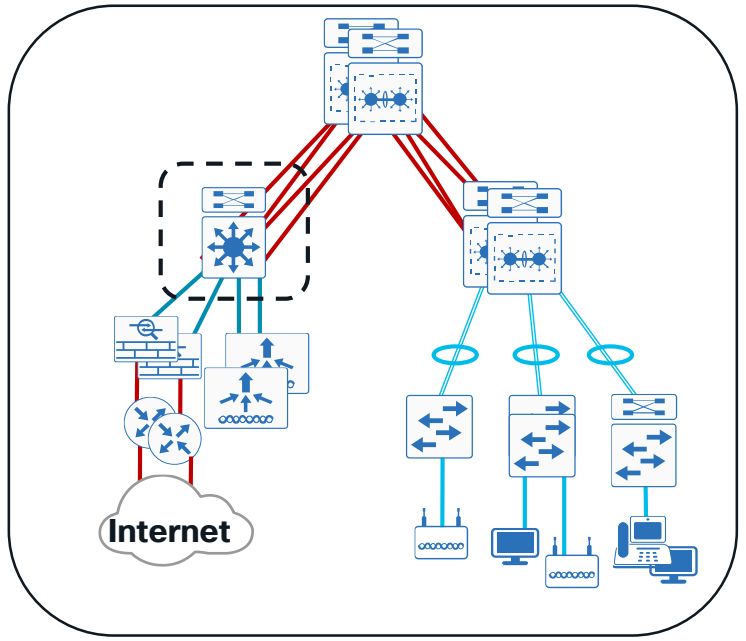

Design Fundamentals

The Standard '1000 Port' (Endpoint) Building Block

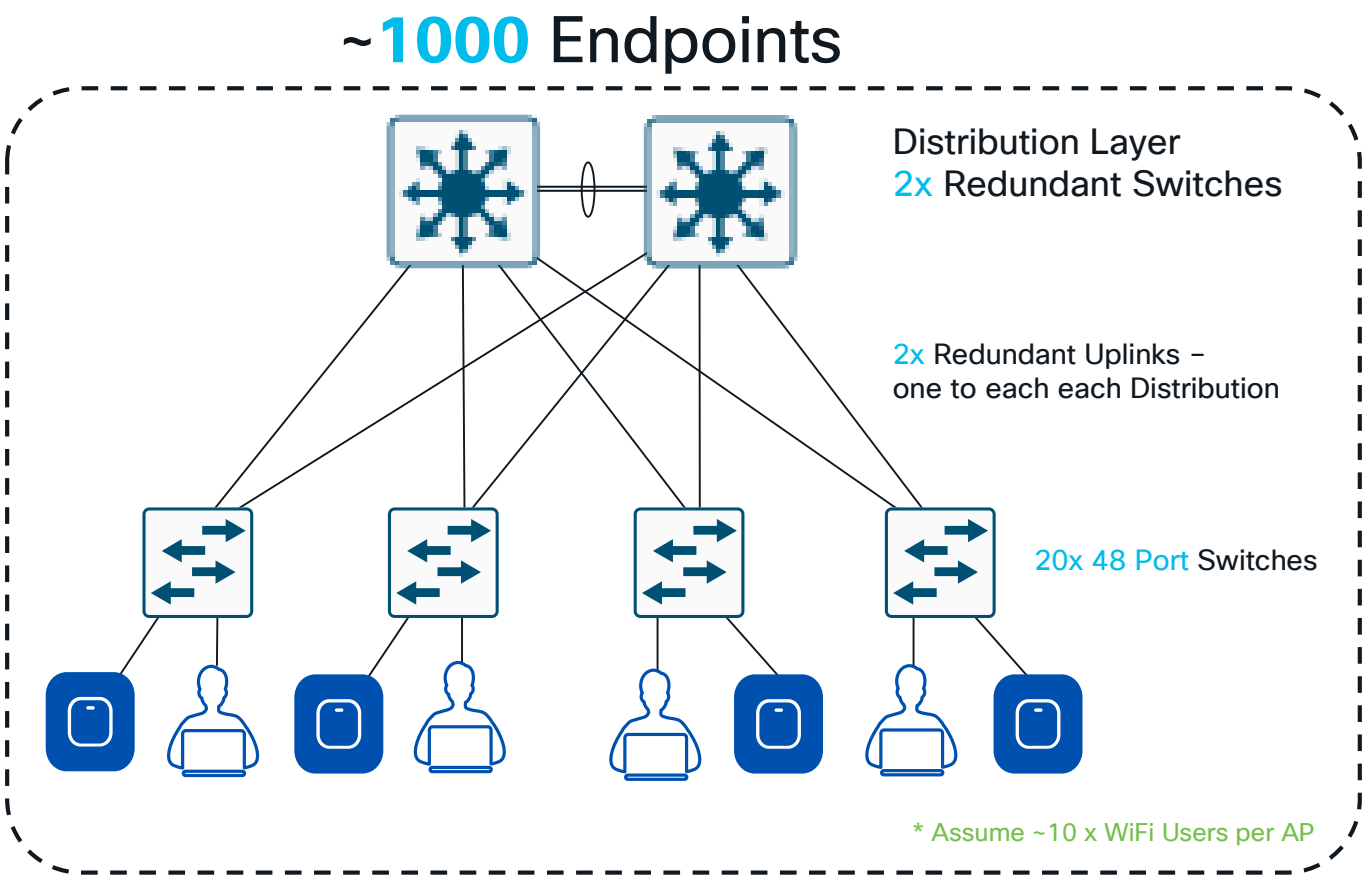

Soft recommendation for Access to Distribution ≤ 20:1

Distribution Switches: **2**

Access Switches: **20** (to 40)

Access Host Ports: **48** (or 24)

- Access Points: **2-4x***

20 x 48 Ports

SUM: **960 Ports**

Oversubcription ratio:

20 : 1

Campus Distro + Ext. Access

The **Distribution + Ext. Access PIN** (Tier 2+) focuses on connecting multiple Access layers, including an Extended Access (IOT/FTTX) layer, to the Core layer.

- Other names: Distribution, BDF
- Common in Very-Large Campus or Large Branch

Main purpose is to **distribute** connectivity (fan-out) from the Core/WAN to the Access **+ Ext. Access**

- Reduces need for high port-density in Core layer

Tends to be **both L3 routed (north) and L2 switched (south)**

- North: VRF, SVI, HSRP/VRRP, ARP/ND, IGP, PIM
- South: VLAN, 802.1Q, STP, MAC, IGMP

Tends to use **multiple L2 & L3 features**

- Access Security (e.g. IPDT/SISF, VACLs, PACLs, etc)
- Access QoS (e.g. NBAR, Classification & Marking)
- Access NetFlow (e.g. AVC, FNF, EPA & ETA)

Tends to require **highest L2/L3 & feature scale**

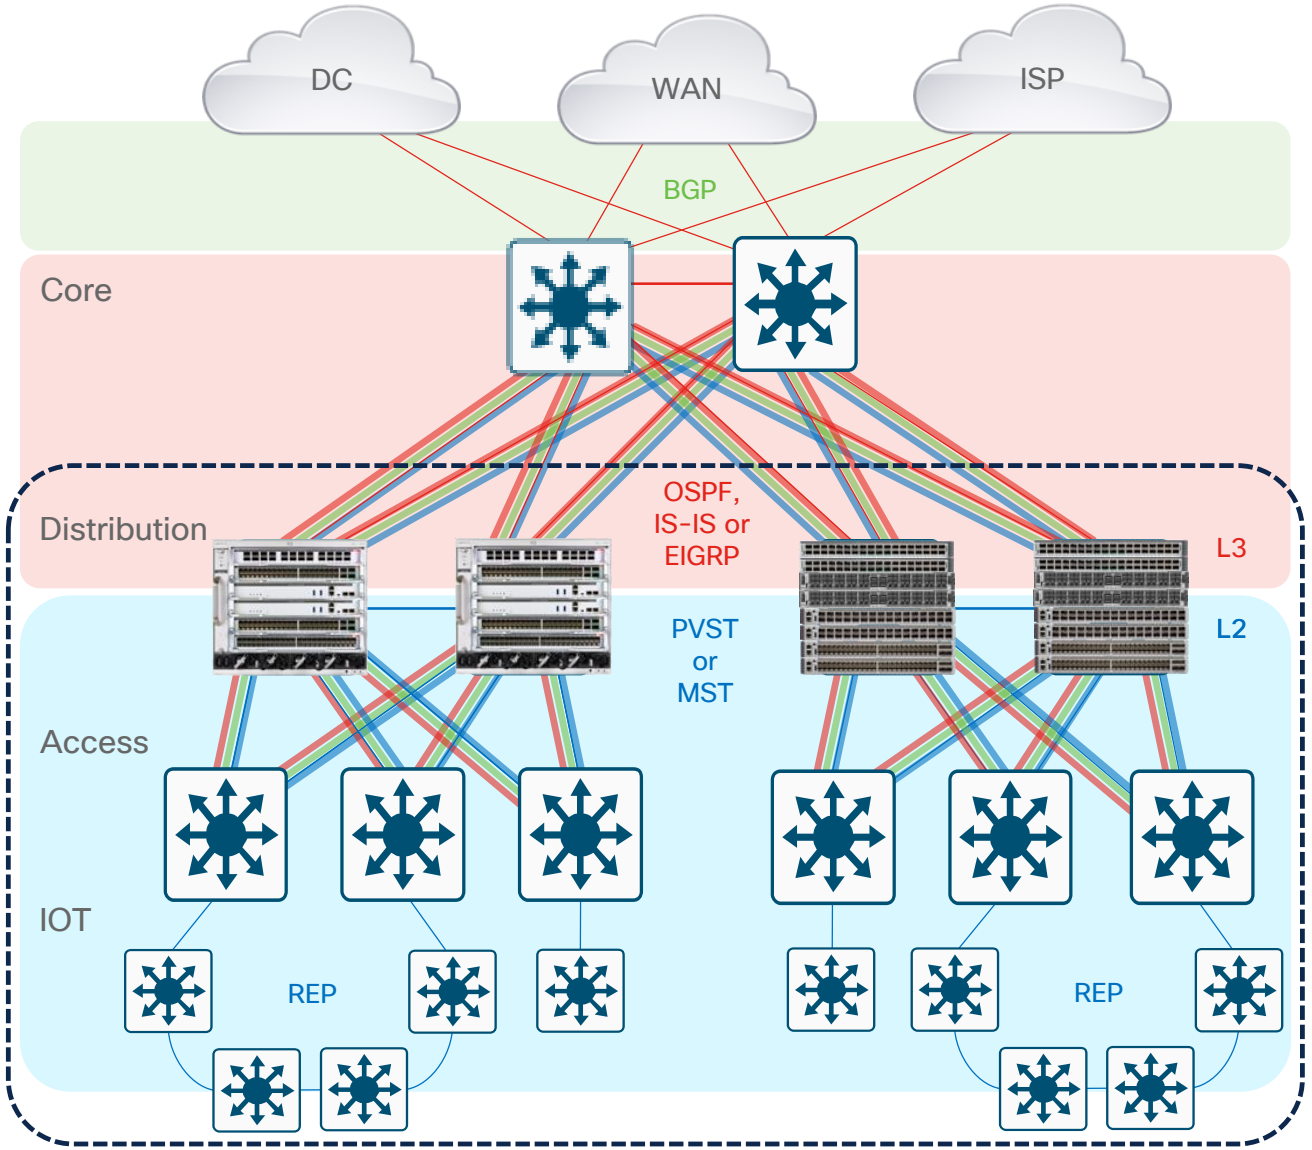

Campus Collapsed Core

The **Collapsed Core** (Tier 2) focuses on connecting multiple Access layers and the WAN/Edge layer.

- Other names : Distribution, BDF
- Common in Small Campus or Medium Branch

Main purpose is to **collapse Core & Distribution** layers

- Mostly for small(er) sites, with low(er) port density
- Similar attributes & requirements as Core + Distribution
- Also applicable to L3 Routed Access

Tends to be **both L3 routed (north) and L2 switched (south)**

- North: SVI, HSRP/VRRP, ARP/ND, IGP, PIM
- South: VLAN, 802.1Q, STP, MAC, IGMP

Tends to use **multiple L2 & L3 features**

- Access Security (e.g. IPDT/SISF, VACLs, PACLs, etc)
- Access QoS (e.g. NBAR, Classification & Marking)
- Access NetFlow (e.g. AVC, FNF, EPA & ETA)

Tends to require **high L2/L3 & feature scale**

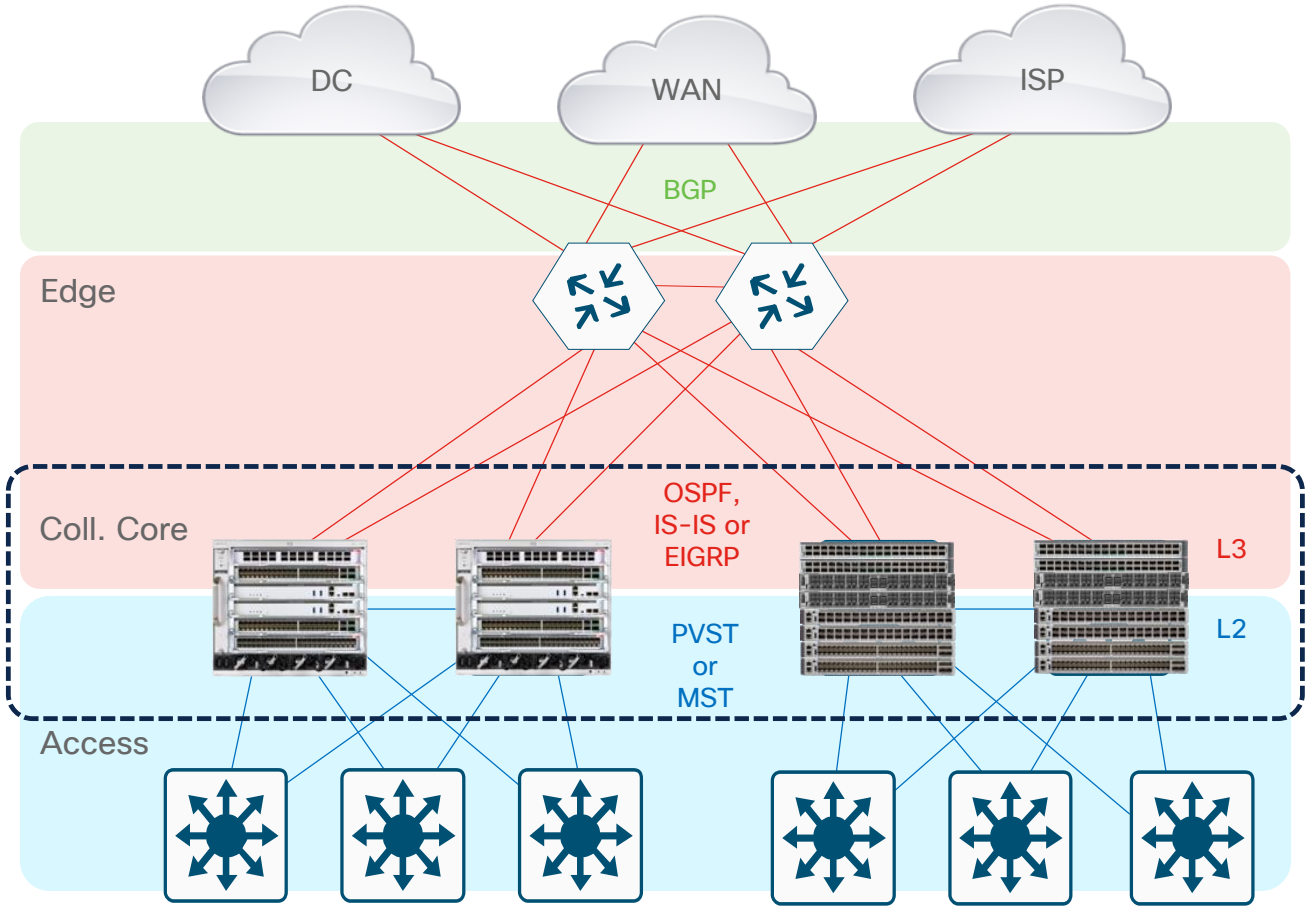

Design Fundamentals

- Multi-Layer Model
- Access Layer
- Distribution Layer
- Core Layer
  - Campus Core (Baseline)
  - Campus Core + Interconnect
  - Campus Core + Edge

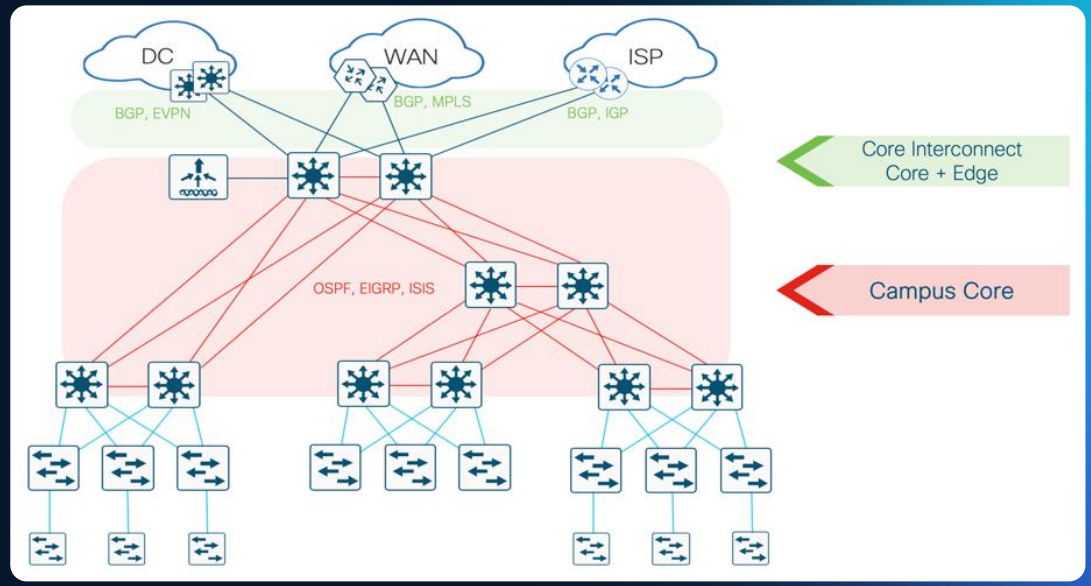

Campus Core (Baseline)

The **Core PIN** (Tier 3) focuses on connecting multiple Distribution layers to an Interconnect (if applicable) and/or other network domains

- Other names: MDF, BDF
- Common in Medium & Large Campus

Main goal is a simple, high-bandwidth, **L3 transport** between other network layers

Tends to be **L3 routed (north & south)**

- North: BGP or IGP (ABR), PIM + MSDP
- South: OSPF, IS-IS or EIGRP, PIM

Tends to use **minimal L3 features**

- Limited ACLs (e.g. inter-area route-maps, remote access)
- Limited QoS (e.g. many-to-one WRED, aggregate policers)
- Limited NetFlow (e.g. inter-area, aggregate flows)

Tends to require **high L3 forwarding scale**

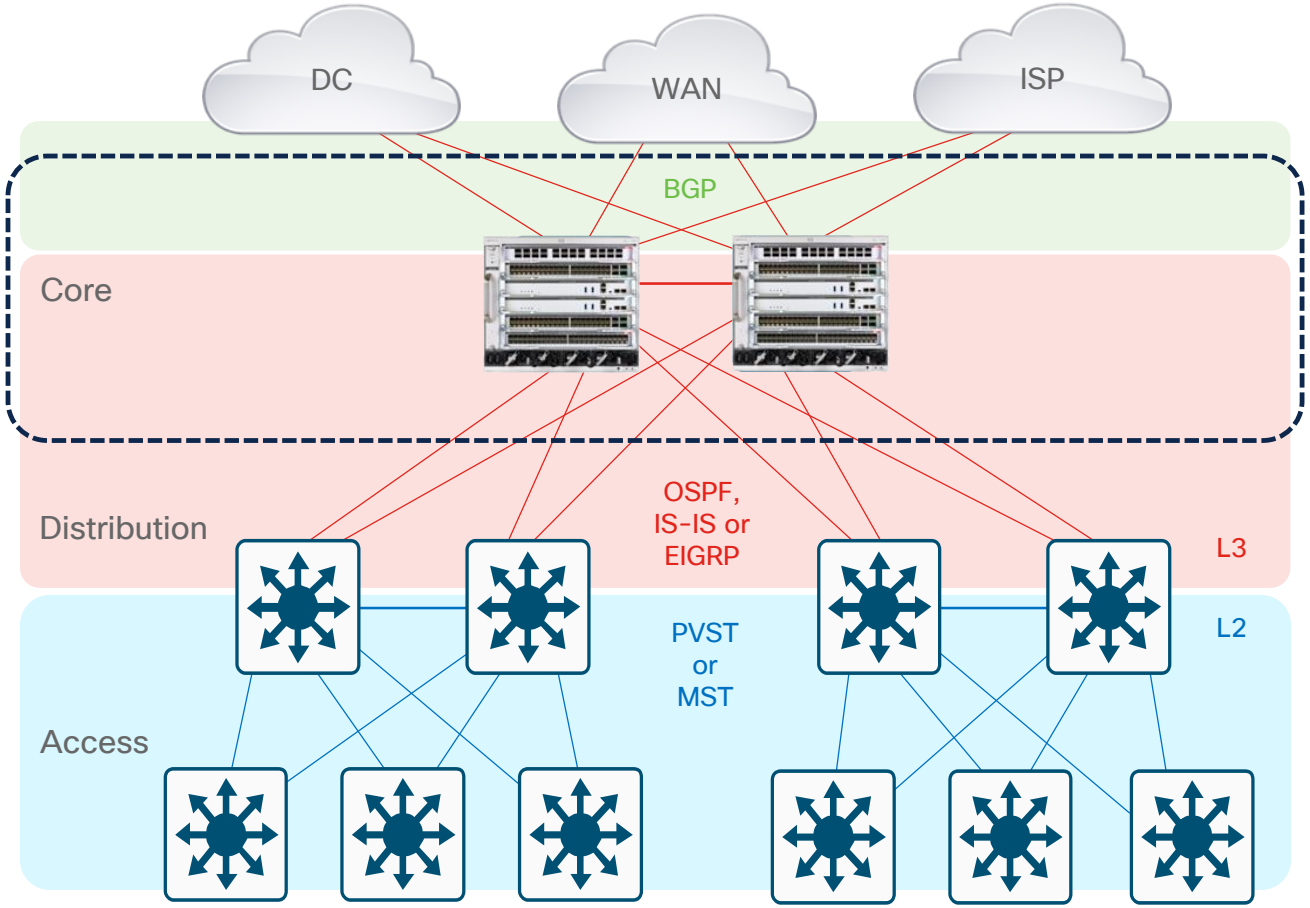

Design Fundamentals

Distribution Layer – Oversubscription Ratios

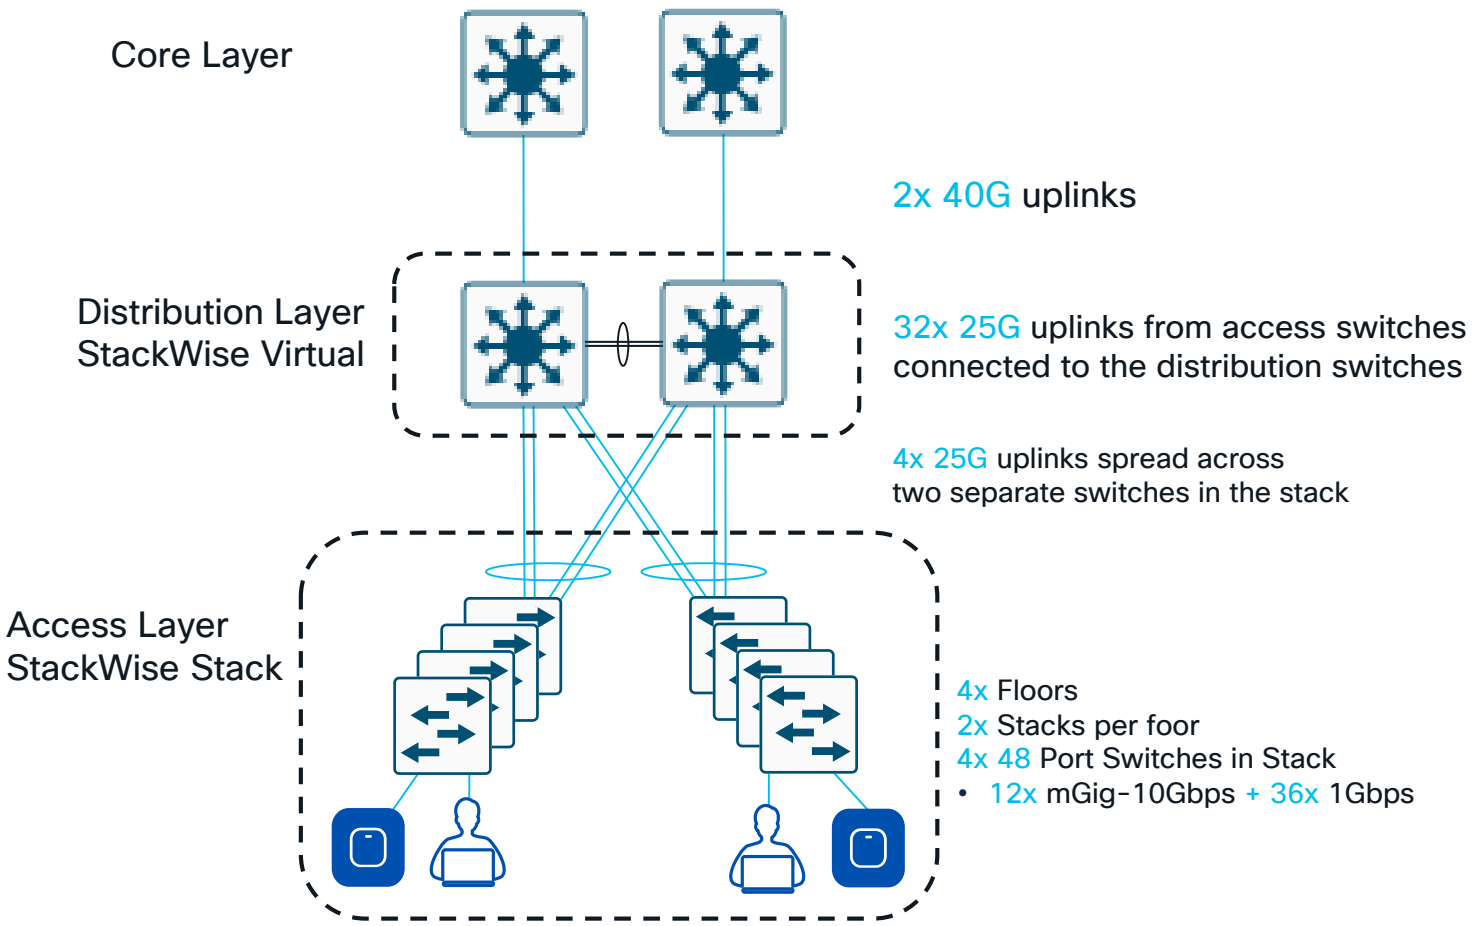

Soft recommendation for Distribution to Core ≤ 4:1

Distribution to Core: **80 Gbps**

Access to Distribution:
4 x 2 x 4 x 25 Gbps

SUM: **800 Gbps**

Oversubcription ratio:
10 : 1

Do I need a Core Layer?

It's Really a Question of Scale, Complexity, and Convergence

No Core (2-Tier)

- Fully-meshed distribution layers
- Difficult to add new blocks
- More physical cabling (2n-2)
- Routing complexity
  - More routing peers
  - More ECMP paths

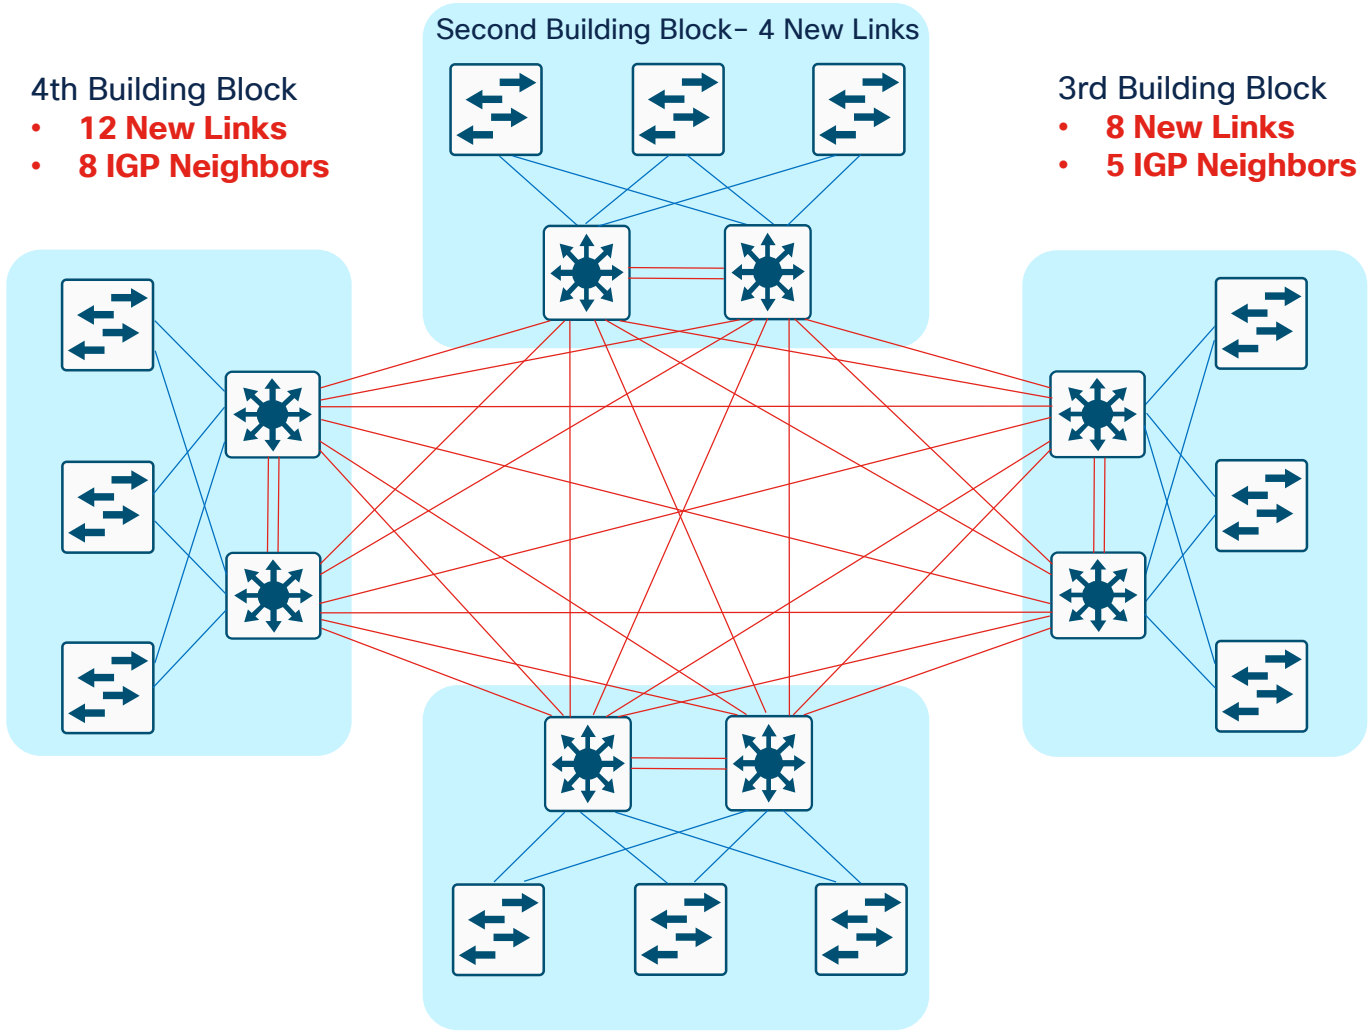

Remember - 1K Ports & 20:1 Rules

1K Endpoints = 1x Distro Block

- Do you need 5K? = 5x Distro
- Do you need 10K? = 10x Distro

Do I need a Core Layer?

It's Really a Question of Scale, Complexity, and Convergence

Dedicated Core (3-Tier)

- Easier to add a block
- Fewer links in the Core
- Easier bandwidth upgrade
- Fewer routing peers
- Fewer ECMP paths
- Best for convergence

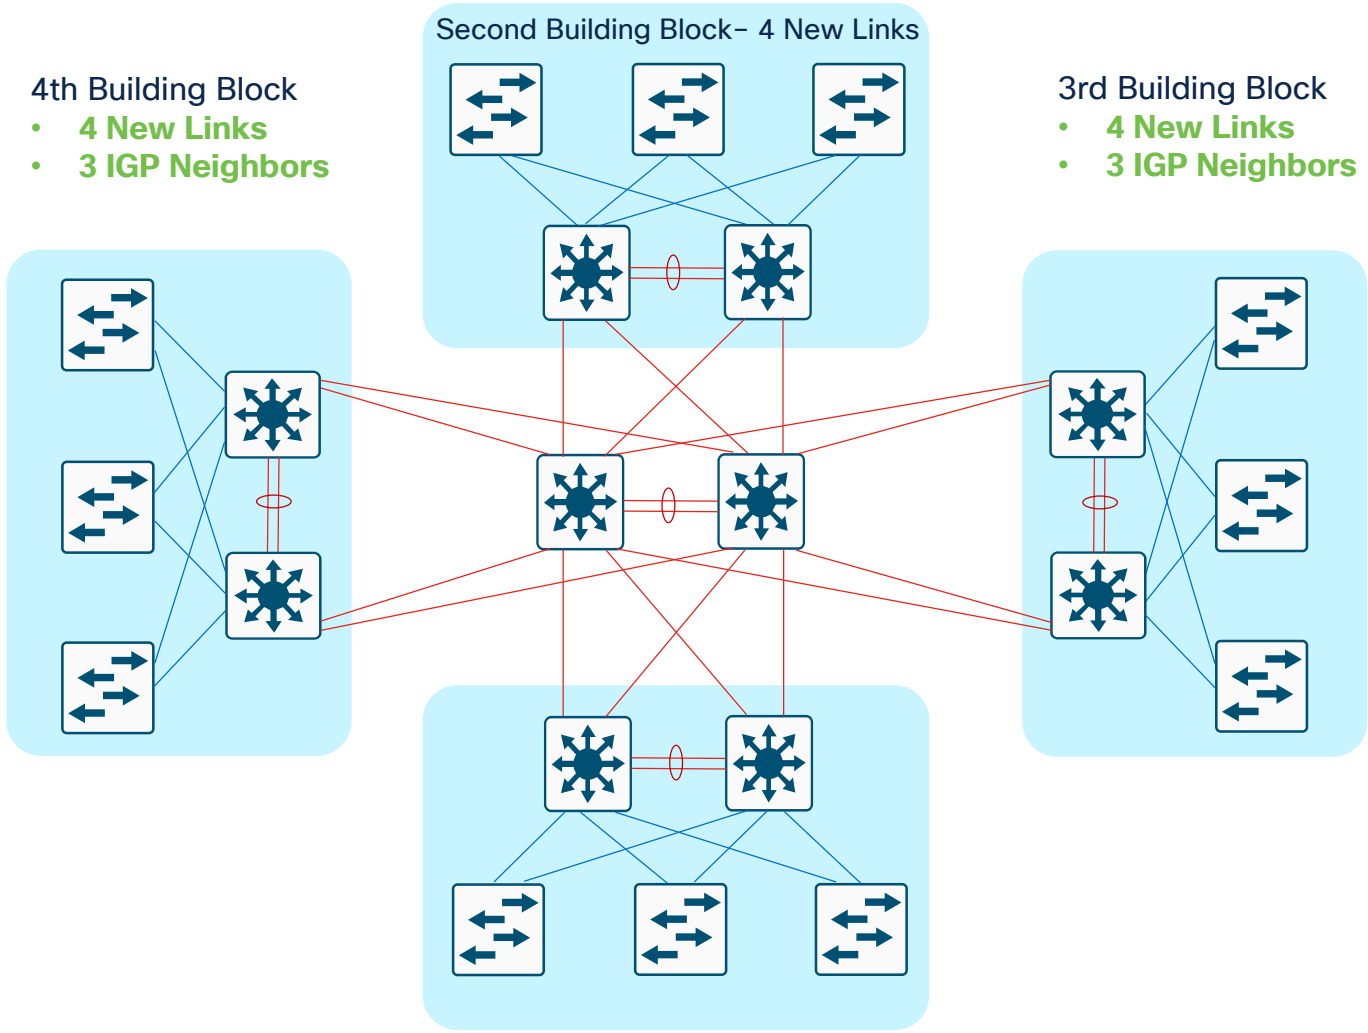

Remember - 1K Ports & 20:1 Rules

1K Endpoints = 1x Distro Block

- Do you need 5K? = 5x Distro
- Do you need 10K? = 10x Distro

Campus Core Interconnect

The **Interconnect PIN** (Tier 4) is an extension of the Core, used to connect multiple Core layers (areas) and/or other network domains.

- Other names: Backbone, Super Core, MAN, DCI
- Common in Large & Very-Large Campus

- Main goal is to **distribute bandwidth and density** requirements of multiple Core layers
  - Similar attributes & requirements as Core PIN
- Tends to be **L3 routed (north & south)**
  - North: BGP or IGP (ABR/ASBR), PIM + MSDP
  - South: OSPF, IS-IS or EIGRP, PIM
- Tends to use **minimal L3 features**
  - Limited ACLs (e.g. inter-area route-maps, remote access)
  - Limited QoS (e.g. many-to-one WRED, aggregate policers)
  - Limited NetFlow (e.g. inter-area, aggregate flows)
- Tends to require **higher L3 scale**

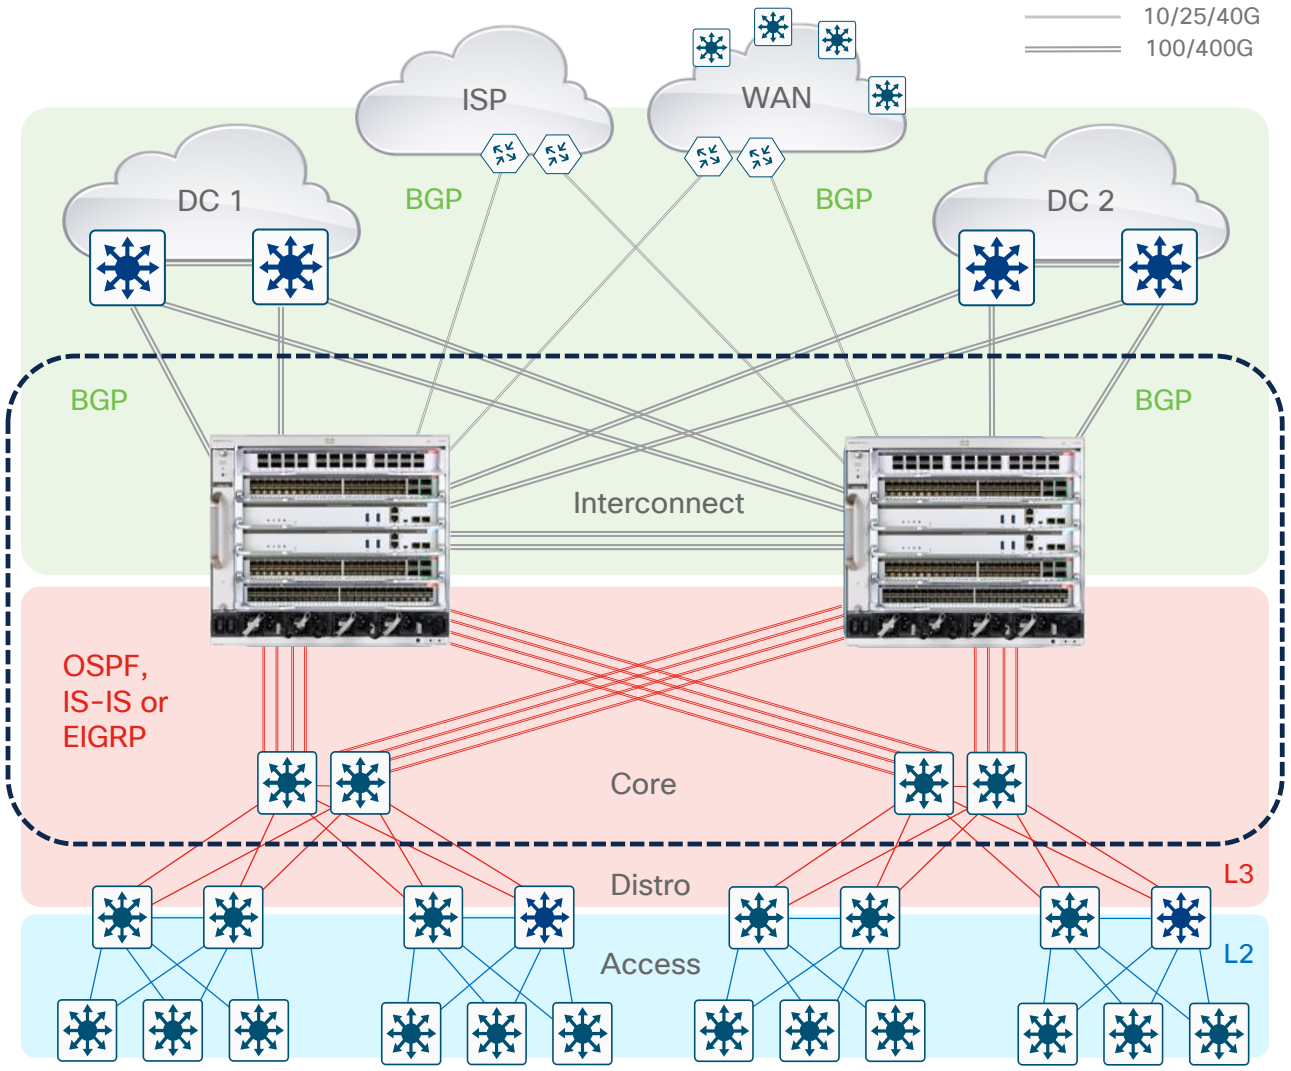

Campus Core + Edge (SP/WAN)

The **Core-Edge PIN** (Tier 4) focuses on connecting multiple Campus areas to remote domains (SP/WAN) and/or to the Internet.

- Other names: Edge Device, Internet Edge
- Common in Medium to Very-Large Campus

Main purpose is to **collapse Core & Edge** layers

Tends to be **L3 routed (north & south)**

- North: MP-BGP + Inter-AS, NAT/PAT, PIM + MSDP
- South: BGP or IGP (ABR/ASBR), PIM + MSDP

Tends to use **Virtualization & Tunnels**

- VRF-Lite, MPLS/VPLS, SR, MVPN
- GRE/MGRE, IPsec, DMVPN
- QinQ, L2oMGRE, OTV, EVPN

Tends to use **multiple L3/VRF features**

- Edge Security ACLs (e.g. RACL, CBAC, ZBFW)
- Hierarchical QoS (e.g. Class-based Queuing, Shaping)
- Policy Based Routing (e.g. WAAS & WCCP)
- WAN NetFlow (e.g. L3/VRF FNF, WAN ETA)

Tends to require **highest L3/VRF & feature scale**

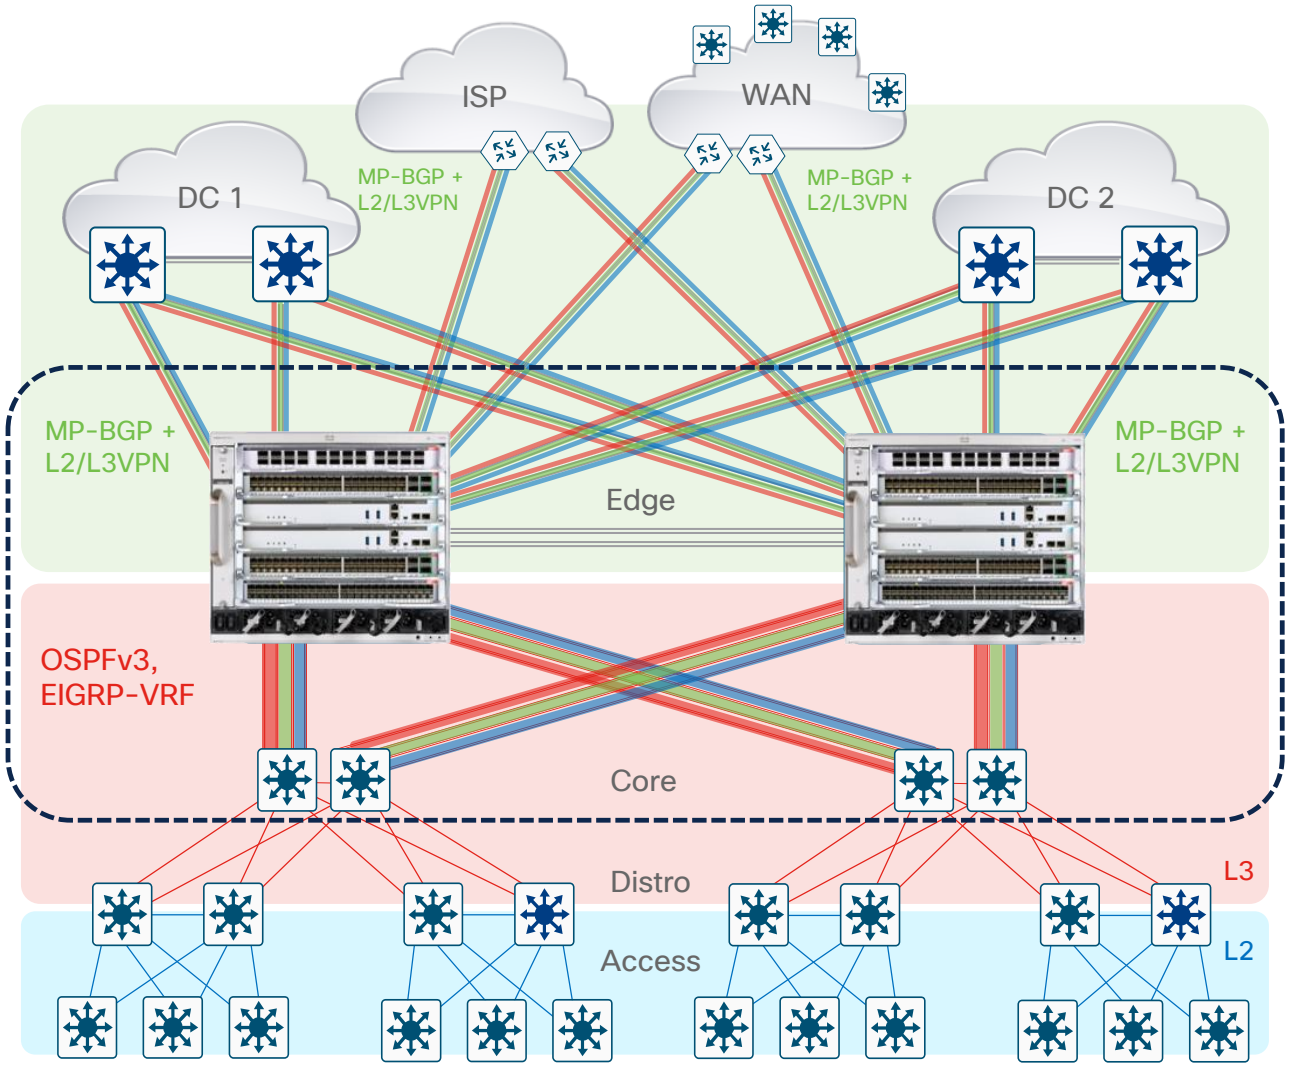

Session Agenda

Design Fundamentals | **Design Considerations**

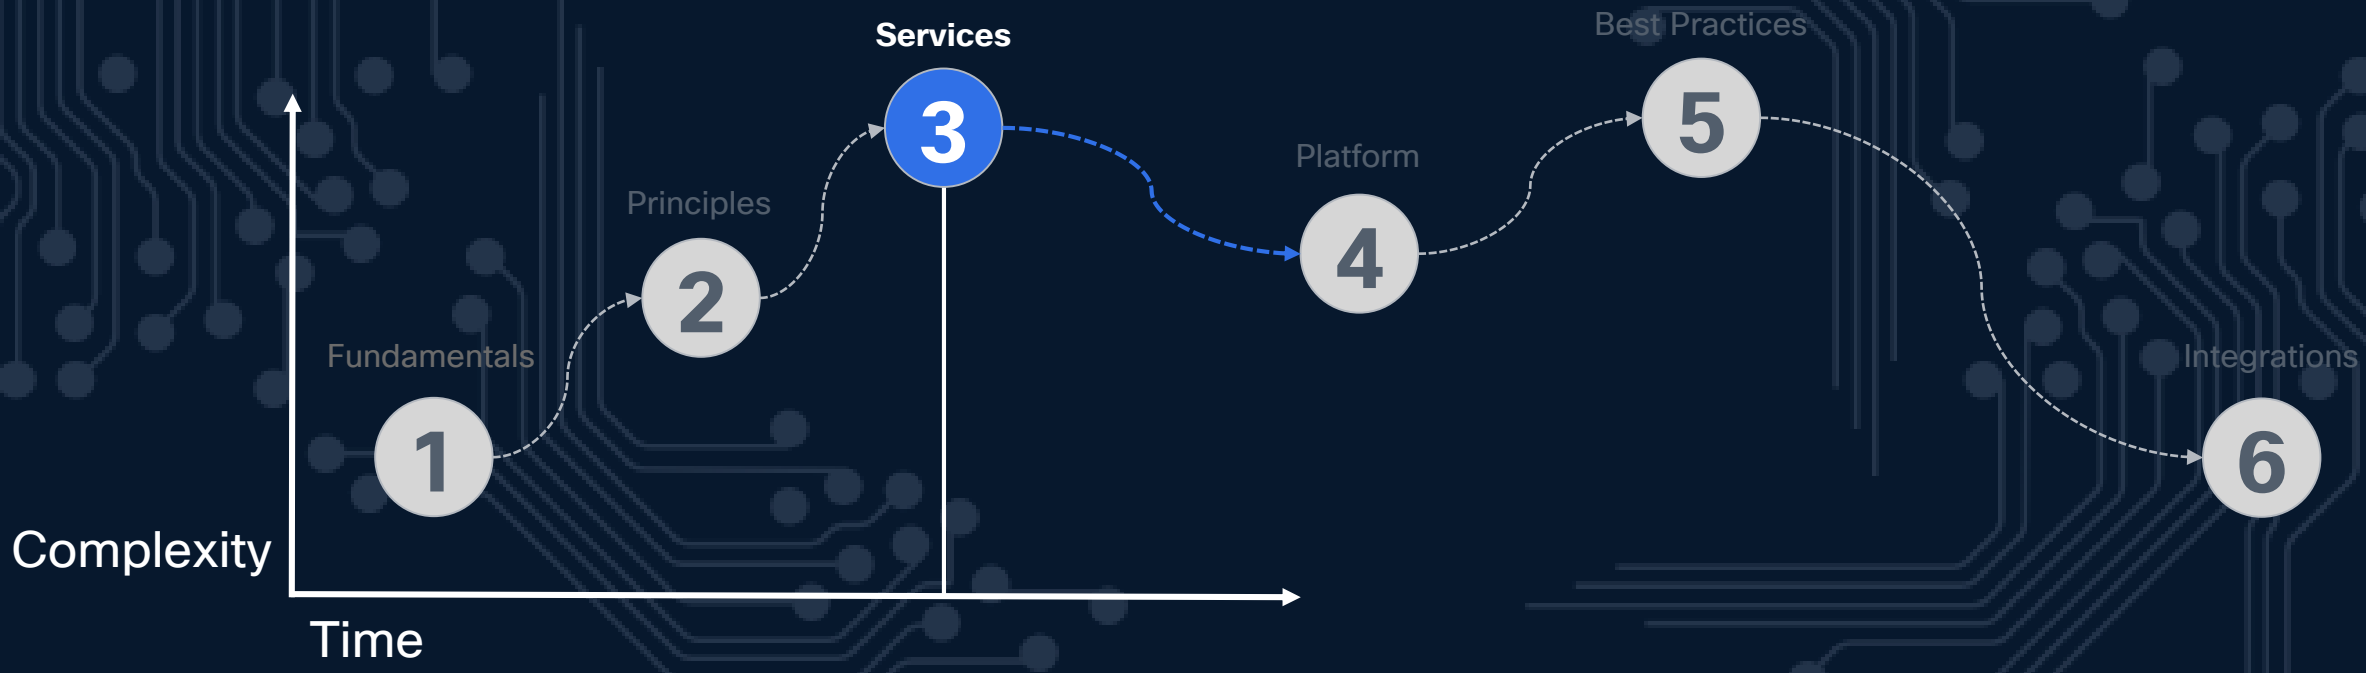

Campus Services

- **Layer 1 physical layer & links**
  - **Copper vs. Fiber Cabling**
  - **L1.5 with UDLD**
  - **L1.5 with EtherChannel**
- **Layer 2 switching protocols**
- **Layer 3 routing protocols**

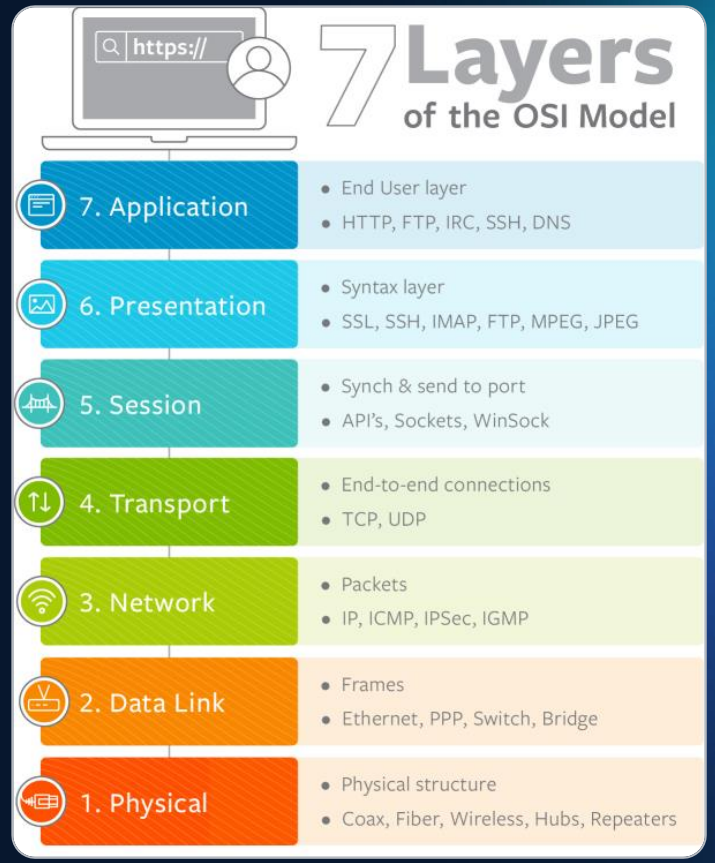

Copper vs. Fiber Media

BRKENS-1500

www.cisco.com/c/en/us/products/interfaces-modules/transceiver-modules/

Category 5, 6 and 7

Unshielded (UTP) | Shielded (STP)

RJ45 (Access to Endpoints)

Cat6A (Offset Wires) | Cat5E (Flush Wires)

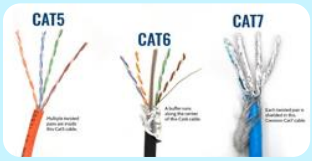

Category	Frequency	Distance	Data Rate	Shielding
5E	100-350 MHz	100m	1000 Mbps	UTP or STP
6	250-550 MHz	1G - 100m 10G - 50m	1 Gbps 10 Gbps	UTP or STP
6A	500-550 MHz	100m	10 Gbps	UTP or STP
7	600 MHz	100m	10 Gbps	Shielded only

10M | 100M | 1G | 2.5G | 5G | 10G

NBASE-T ALLIANCE

OM3, OM4 and OM5

Multi-Mode (MMF) | Single-Mode (SMF) | Wave-Division Multiplex (WDM)

SFP (Access and Distribution) | **QSFP** (Core and Edge)

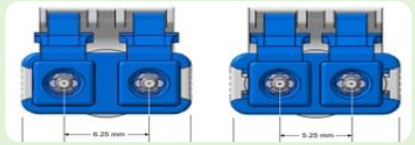

SFP-LC (LC Duplex) | mSFP (Mini LC Duplex)

MPO12 (12 Fibers) | MPO24 (24 Fibers)

Multimode
- Short distance cable runs (less than 1000ft.)
- High bandwidth support
- Higher cable cost
- Lower electronics cost
- Easier to terminate due to larger core size

Single Mode
- Long distance cable runs (greater than 1000ft.)
- Highest bandwidth support
- Lower cable cost
- Higher electronics cost
- Harder to terminate due to smaller core size

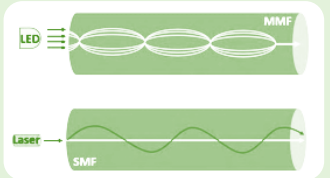

100M | 1G | 10G | 25G | 40G | 50G (NEW) | 100G (NEW) | 200G | 400G (NEW)

www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/nb-06-cat9000-panduit-cables-wp-cte-en.html

Unidirectional Link Detection

Protecting Against One-Way Communication

- UDLD protects against one-way communication or partially failed optics, and the effect it could have on L2 protocols like STP
- Primarily used on fiber optic links where patch or cable errors cause link up/up – with mismatched transmit/receive pairs
- Each switch port configured for UDLD will send UDLD protocol packets (L2) containing the port's own device/port ID
  - The neighbor's device/port IDs seen by UDLD on that port
- Neighboring ports should see their own device/port ID (echo) in the packets received from the other side
- If the port does not see its own device/port ID in the incoming UDLD packets (for a specific duration) – then the link is considered unidirectional and is put into errdisable

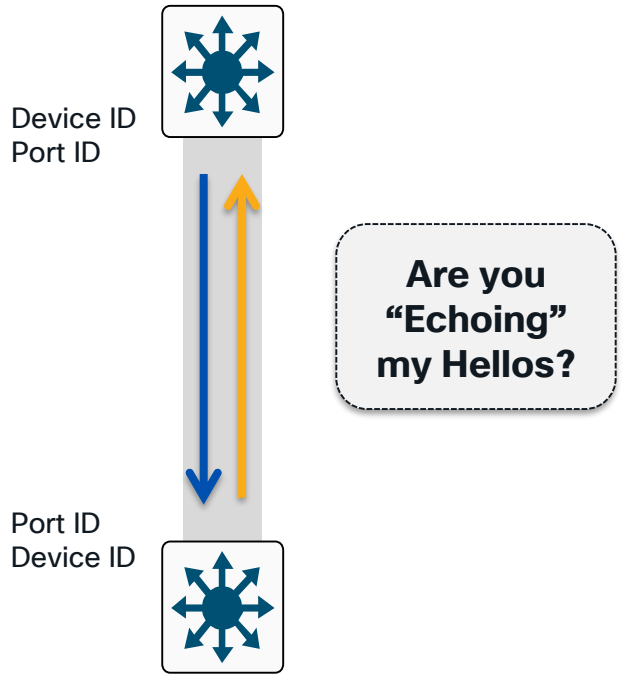

L2 UDLD Configuration

Best practices

- Typically deployed for faster recovery of ‘reliable’ fiber connections.
- Use UDLD aggressive mode for most aggressive protection
  - more incoming packets & CPU usage
  - aggressive may cause false-positive
- Turn on in global config to avoid operational error/misses
- Config example:
  - Cisco IOS: `udld aggressive`

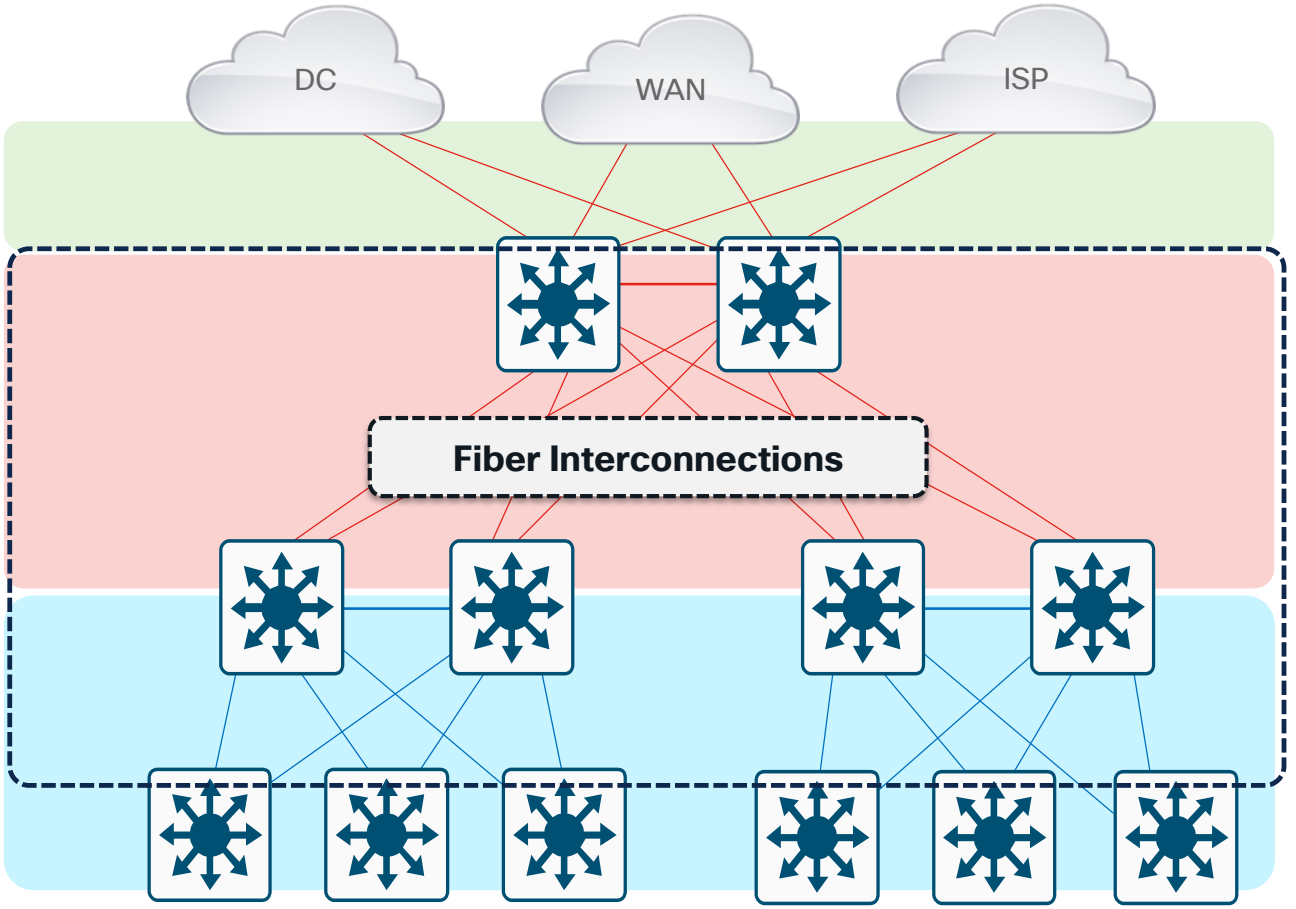

Campus + EtherChannel

Using **EtherChannel** focuses on combining multiple physical links into a single logical link

- Other names: Portchannel, Link-Aggregation (LAG)
- Common in Medium & Large Campus

Main goal is to **increase bandwidth**, and provide **link-level redundancy** between network layers

- Mostly for large(r) sites, with high(er) port density
- Similar attributes & requirements as existing PIN(s)

Can be used for **both L2 & L3 links** (north & south)

- North: BGP or IGP, PIM
- South: STP or REP, IGMP/MLD

Tends to require **special L2/L3 features**

- Portchannel ACLs (e.g. L2/L3 RACL)
- Portchannel QoS (e.g. L2/L3 aggregate policers)
- Portchannel NetFlow (e.g. L2/L3 FNF)

Tends to require **less L2/L3 forwarding scale**

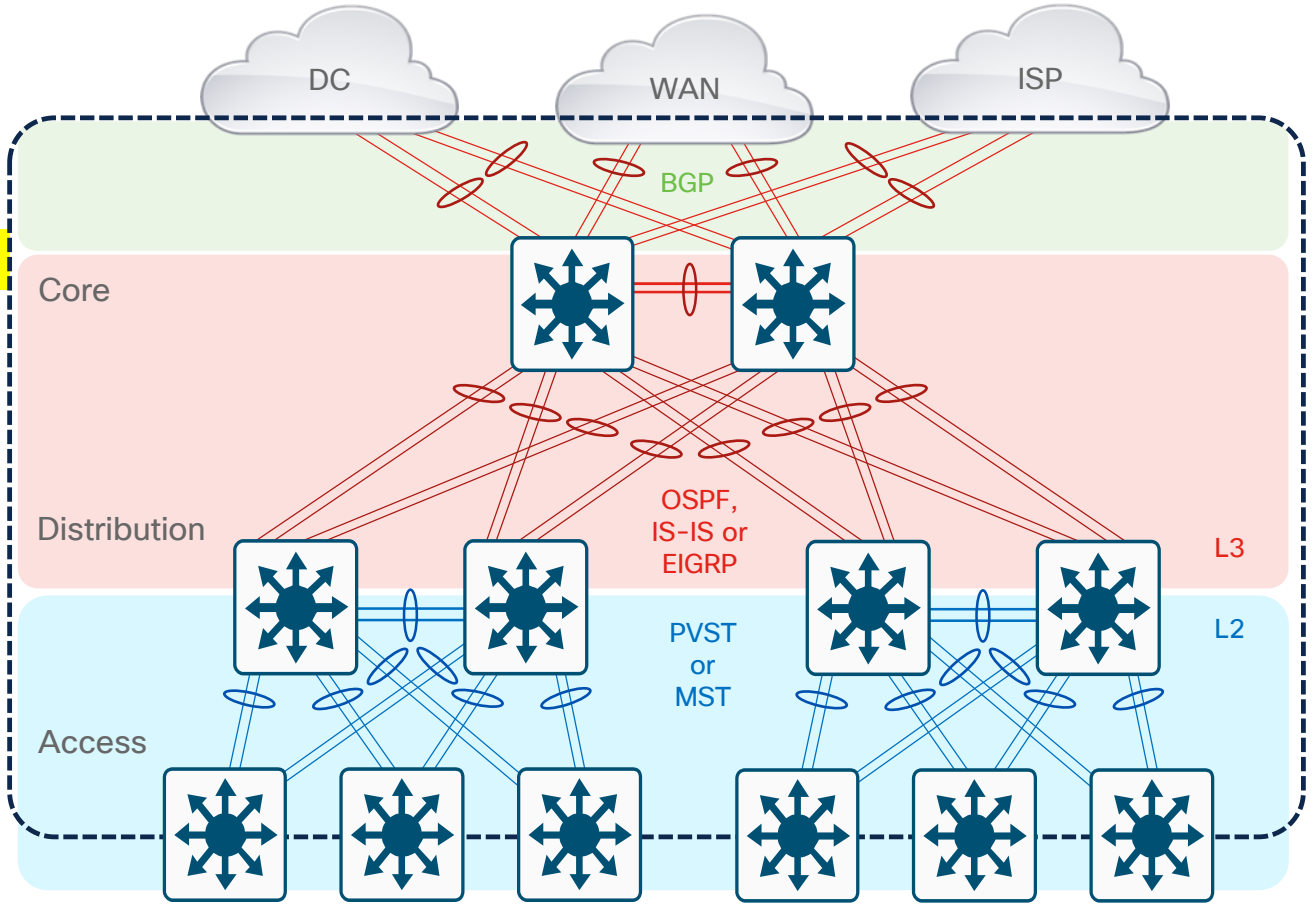

EtherChannels

Reduce Complexity/Peer Relationships

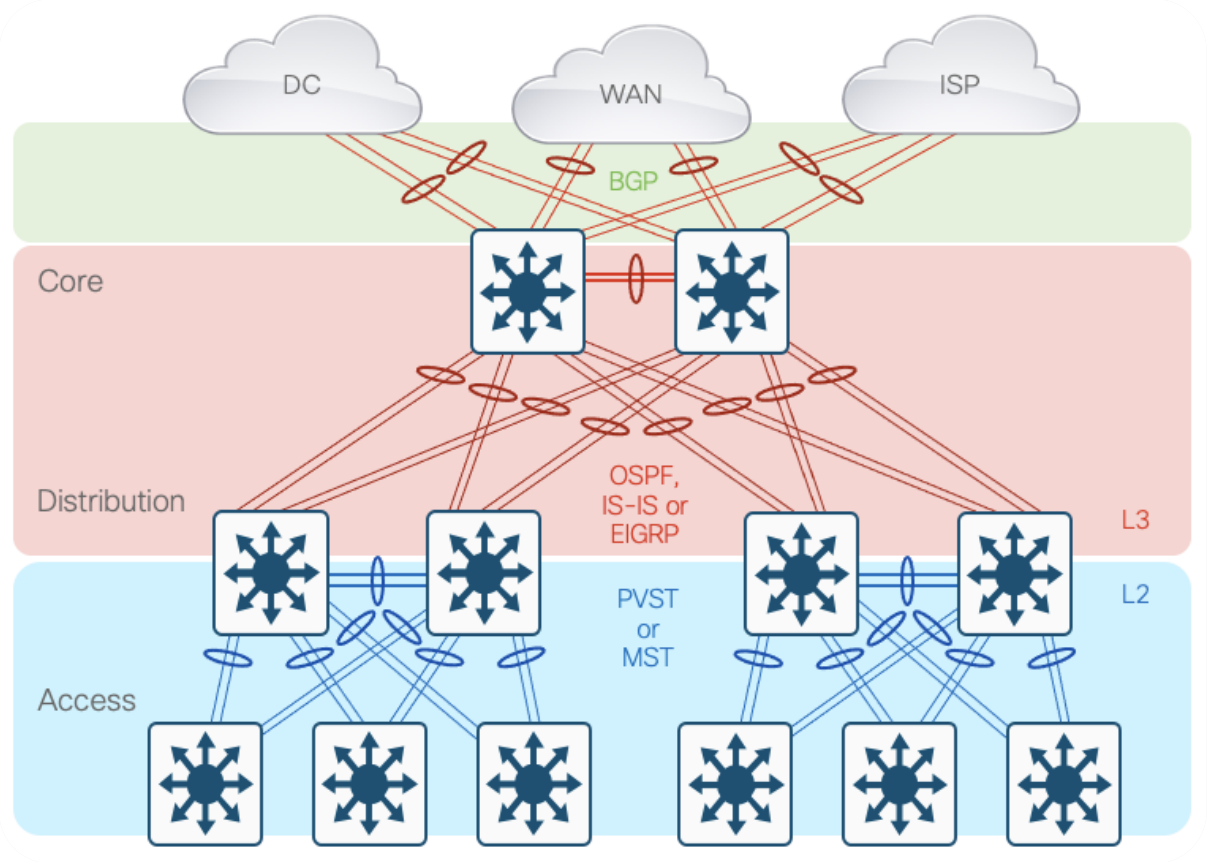

- More links = more protocol neighbor relationships (and associated overhead)
  - more incoming packets & CPU usage
  - multiplied by number of links
  - multiplied by protocol timers
- EtherChannels allow you to reduce peers by creating single logical interface to peer
- When single link-failure in a Channel:
  - OSPF on a Cisco IOS-based switch will reduce link cost (and may reroute traffic)
  - EIGRP can change link cost or not change and overload the remaining links

Understanding Ether Channel

Channel Negotiation Options – PAGP & LACP

Port Aggregation Protocol (PAGP)

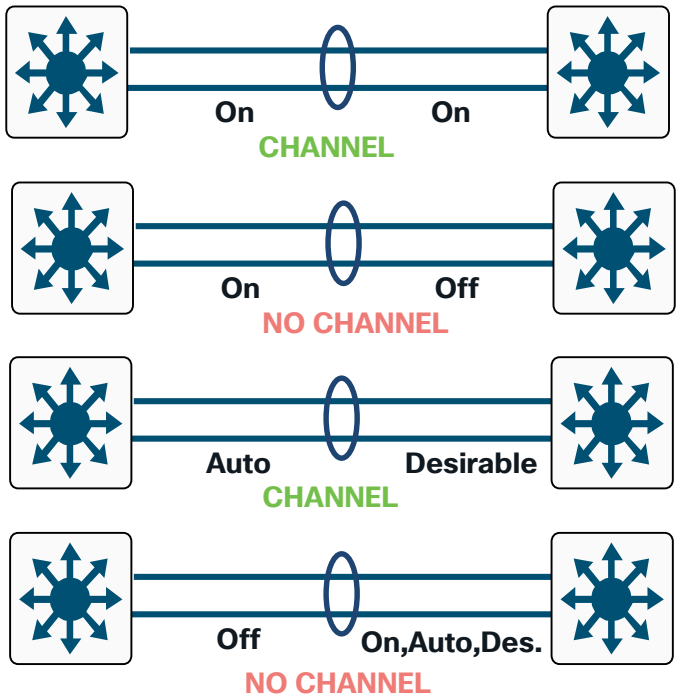

- **On:** always be a channel/bundle member
- **Desirable:** ask if the other side can/will
- **Auto:** if the other side asks - I will
- **Off:** don't become a channel member

Link Aggregation Protocol (LACP)

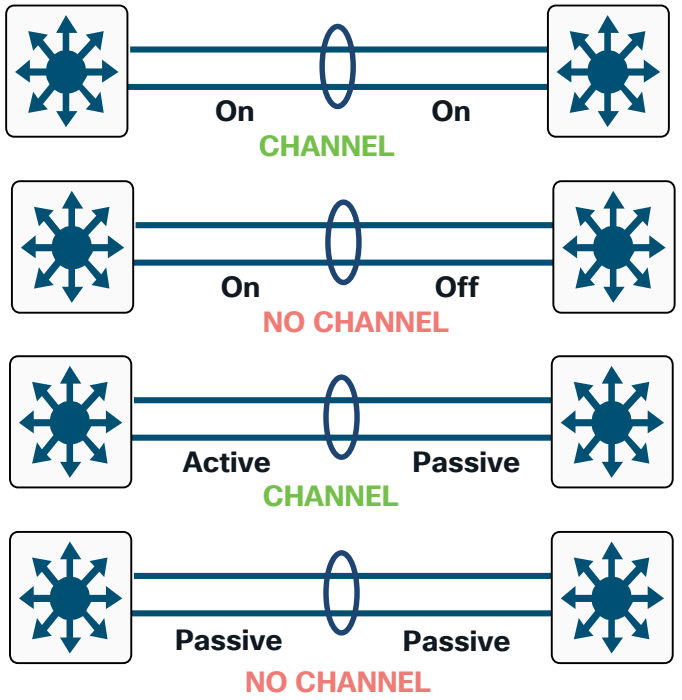

- **On:** always be a channel/bundle member
- **Active:** ask if the other side can/will
- **Passive:** if the other side asks - I will
- **Off:** don't become a channel member

Ether Channel load balancing

Use as much information as possible

- Cisco switches let you tune the hashing algorithm used to select the specific EtherChannel link
- You can use the default source + destination IP header information
- Or you can add an additional fields, such as L2 MAC, L4 TCP/IP port, etc. as input to the algorithm

```
switch(config)#port-channel load-balance ?
  dst-ip                  Dst IP Addr
  dst-mac                 Dst Mac Addr
  dst-mixed-ip-port       Dst IP Addr and TCP/UDP Port
  dst-port                Dst TCP/UDP Port
  extended                Extended Load Balance Methods
  src-dst-ip              Src XOR Dst IP Addr
  src-dst-mac             Src XOR Dst Mac Addr
  src-dst-mixed-ip-port   Src XOR Dst IP Addr and TCP/UDP Port
  src-dst-port            Src XOR Dst TCP/UDP Port
  src-ip                  Src IP Addr
  src-mac                 Src Mac Addr
  src-mixed-ip-port       Src IP Addr and TCP/UDP Port
  src-port                Src TCP/UDP Port
```

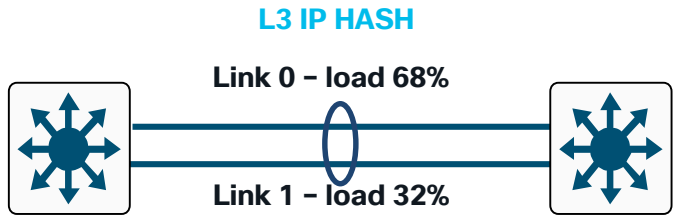

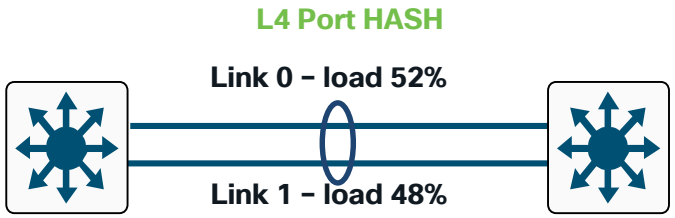

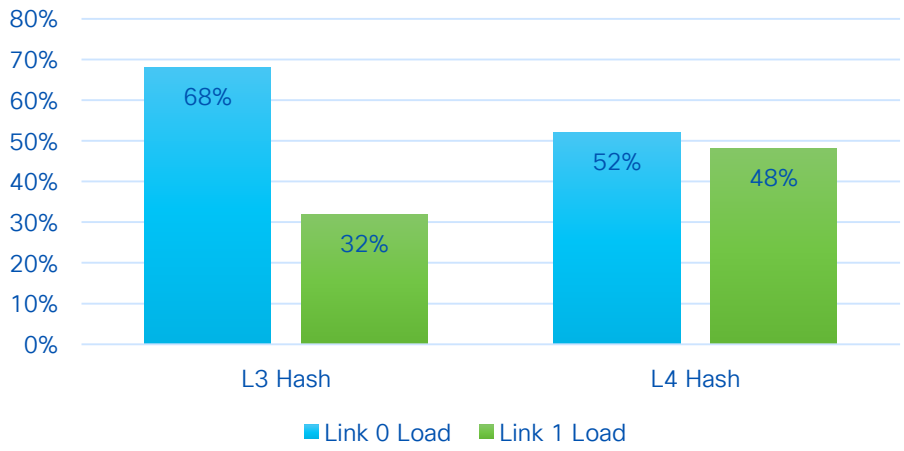

L2/L3 EtherChannel

Best Practices

- Typically deployed in **ALL** Campus multi-layer interconnects
- Used to provide link redundancy and reduce peering complexity
- Tune the load-balancing hash to achieve maximum utilization of all channel members
- Deploy in Exponents of Two (2, 4, or 8)
- Use 802.3ad LACP protocol for standard interop

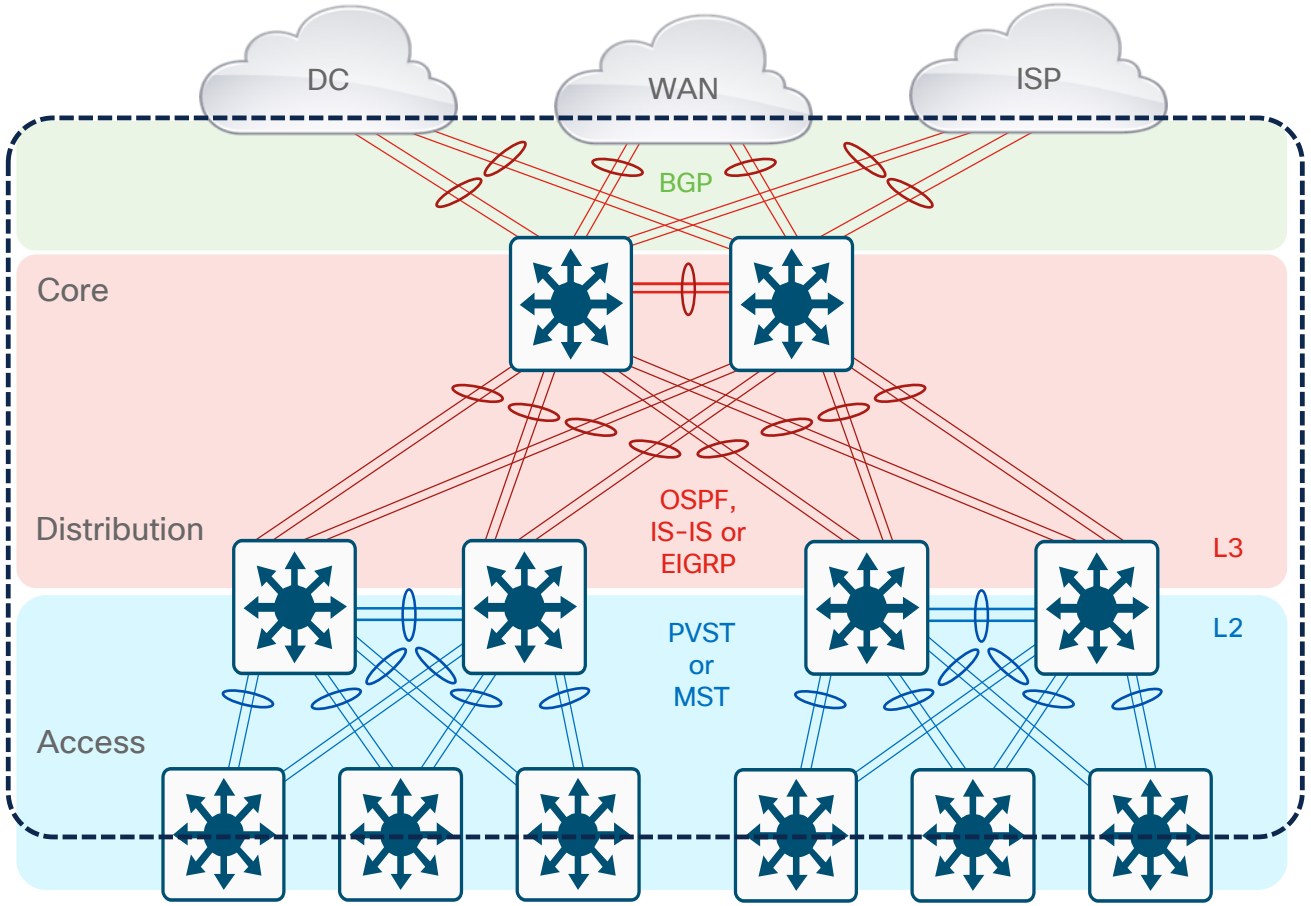

Campus Services

- Layer 1 physical layer & links
- Layer 2 switching protocols
  - Layer 2 Design
  - L2 Spanning-Tree
  - L2 VLAN Trunks
- Layer 3 routing protocols

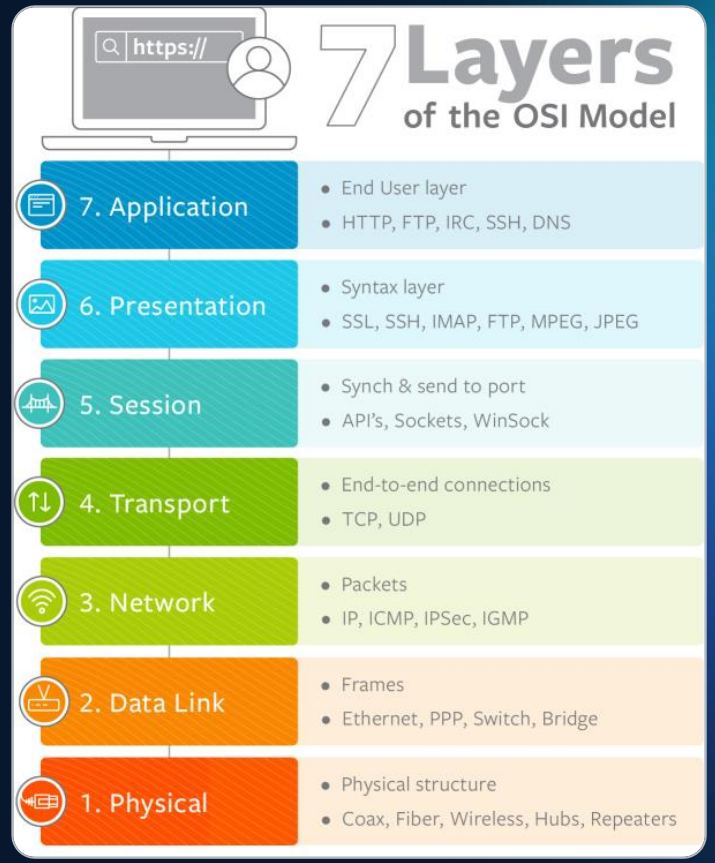

L2 Spanning Tree

Best Practices

- Only extend VLANs across Access & Distribution layers when you must!
- Use PVST for best convergence
  - Rapid-PVST+ (RPVST) is default
- Use MST for best scale
  - Required to protect against access loops
  - Required to protect against operational accidents (misconfig or hardware failure)
  - Take advantage of Spanning Tree toolkit

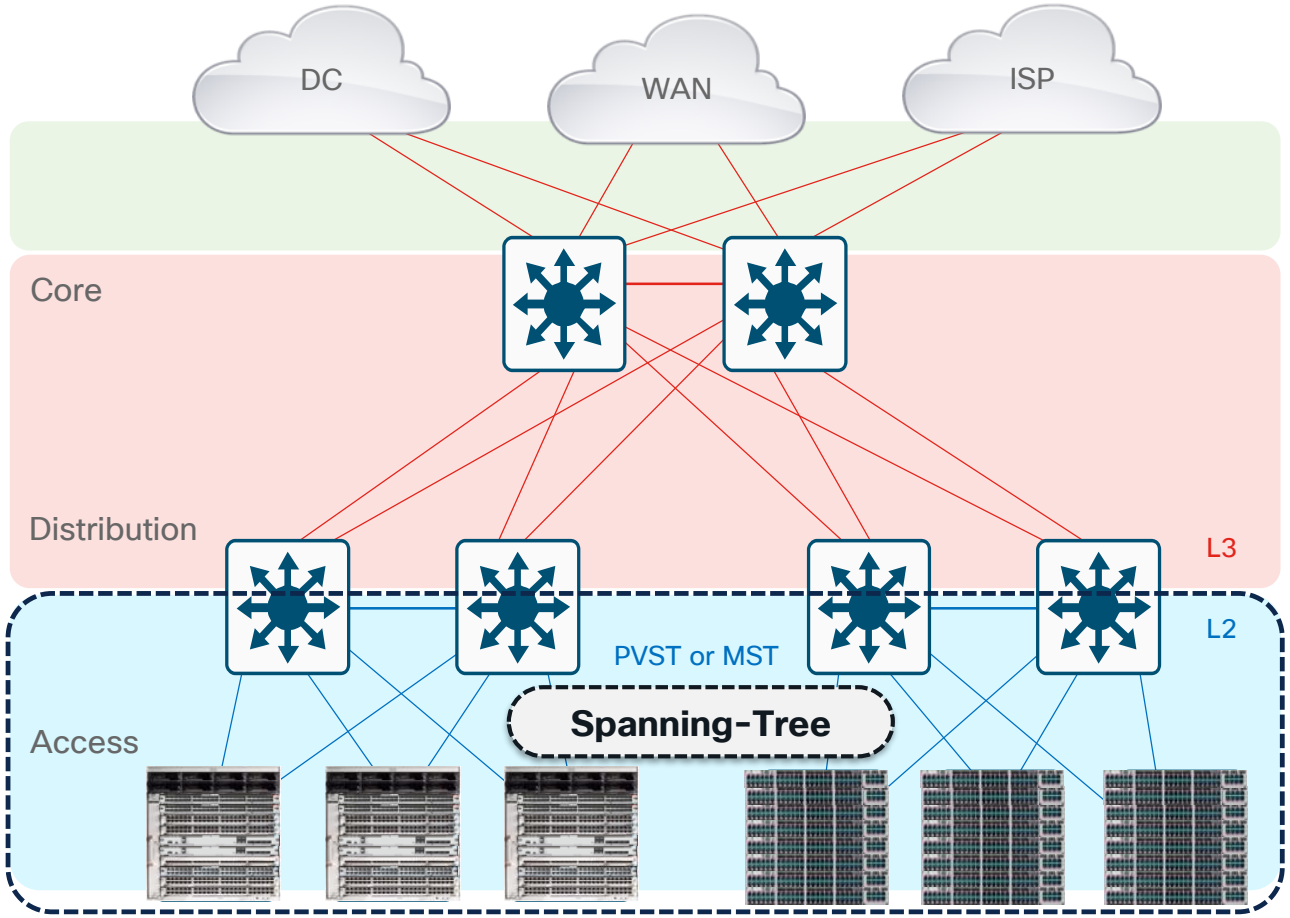

Daisy Chaining Access Layer Switches

Avoid Potential Black Holes

Daisy-Chains: Return Traffic Has a 50/50 Chance of Being ‘Black Holed’

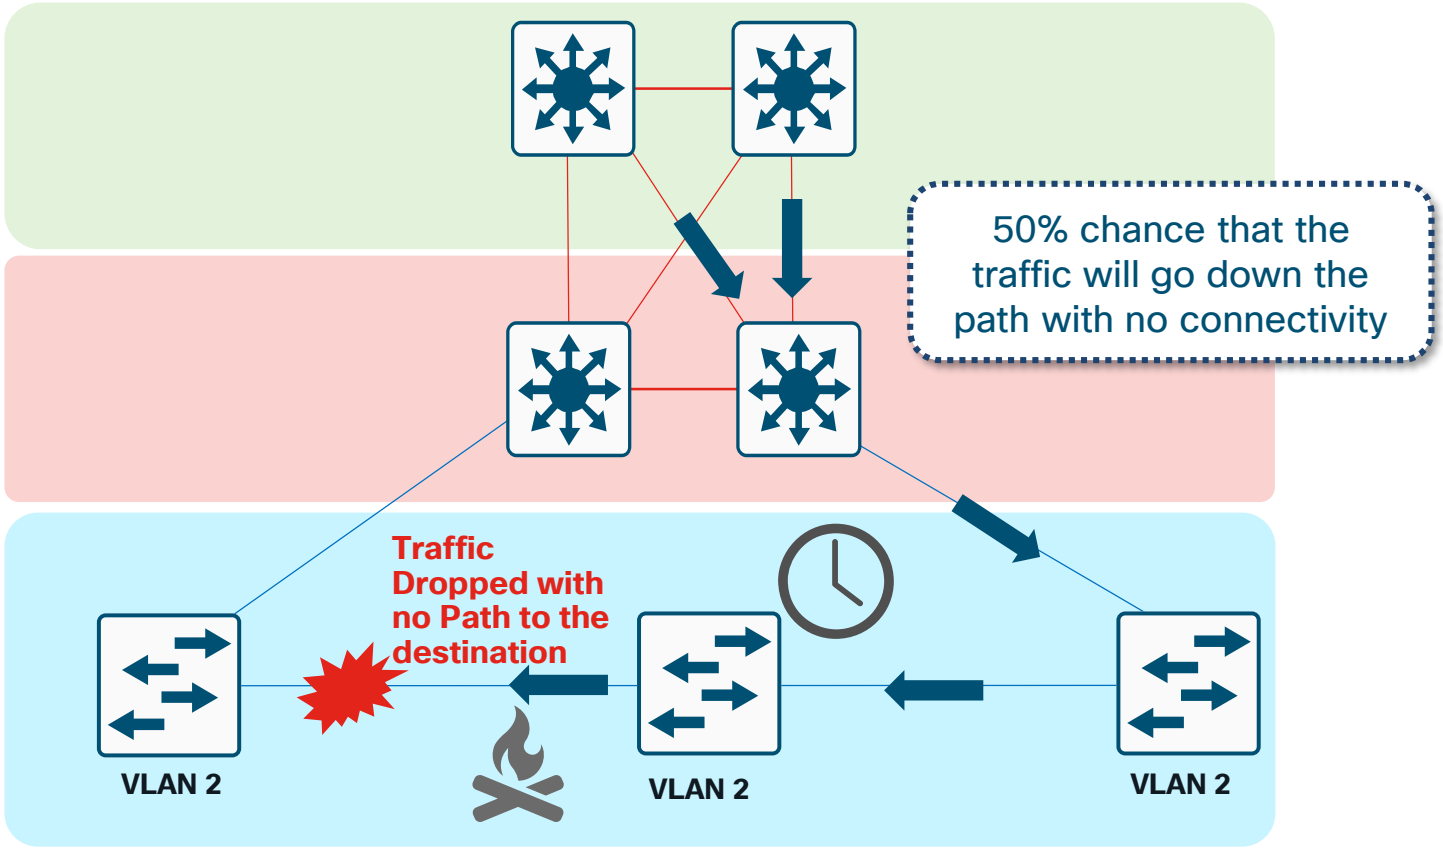

Virtual Trunk Protocol

VTP

- Centralized VLAN management
- VTP server switch propagates VLAN database to VTP client switches
- Runs only on Trunks
- Four modes:
  - Server: updates clients and servers
  - Client: receive updates– cannot make changes
  - Transparent: let updates pass through
  - Off: ignores VTP updates

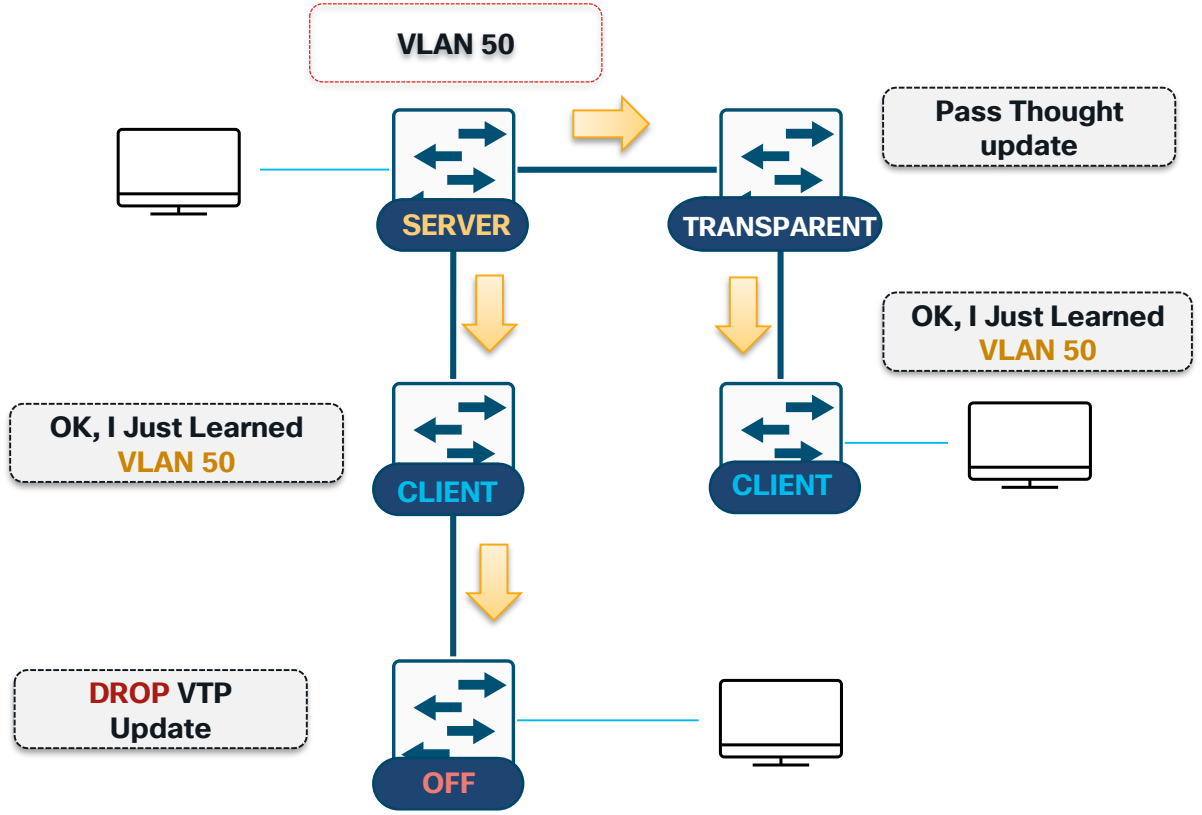

Trunk Auto/Desirable takes some time

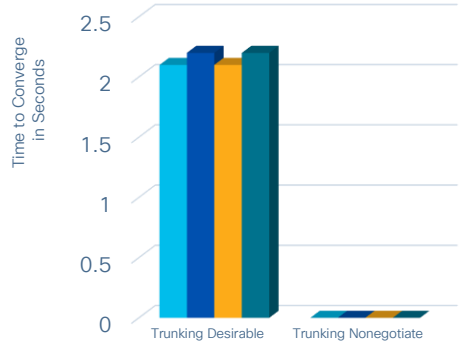

L2 Trunk Configuration

Best Practices

- Typically deployed on interconnection between Access and Distribution layers
- Use VTP transparent mode to decrease potential for operational error
- Hard set trunk mode to ON and encapsulation negotiate off for optimal convergence
- Change the native VLAN to something unused to avoid VLAN hopping
- Manually prune all VLANS except those used
- Disable on host ports*

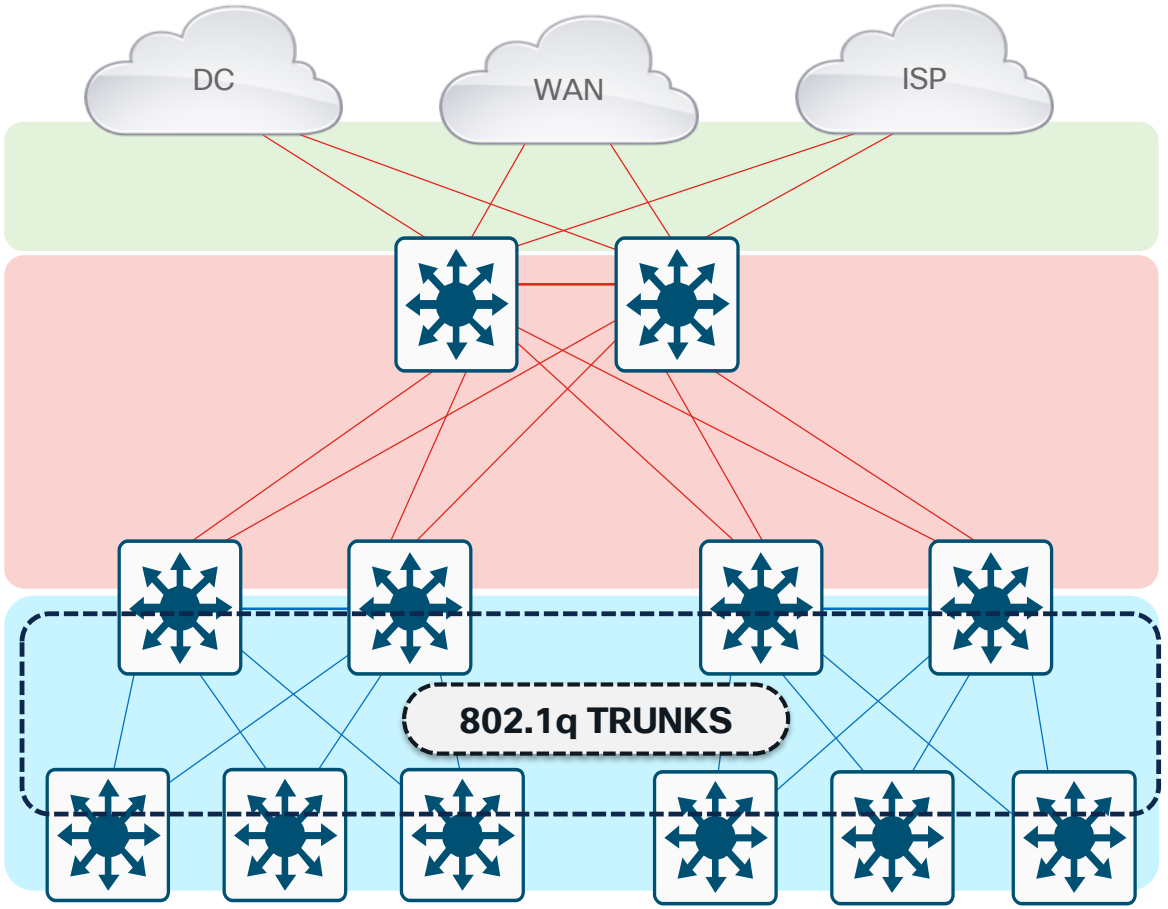

Campus Services

- Layer 1 physical layer & links
- Layer 2 switching protocols
- Layer 3 routing protocols
  - Layer 3 Design
  - L3 First Hop Routing
  - L3 with BFD
  - L3 Best Practices

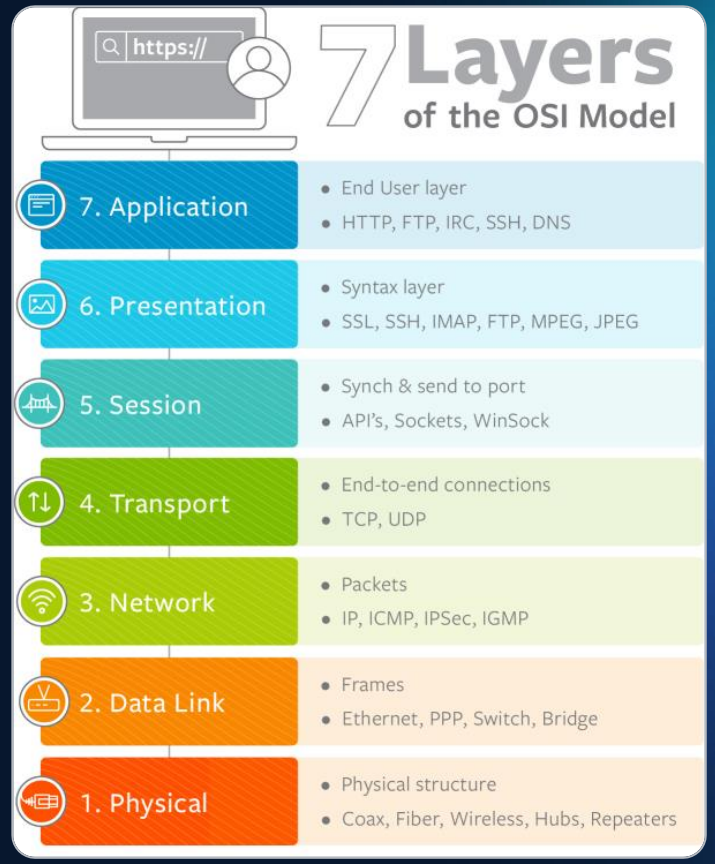

L3 Routing Protocols

Best Practices

- Typically deployed in Distribution-to-Core, and Core-to-Core interconnects
- Used to quickly re-route around failed nodes or links, while providing load balancing over redundant paths
- Build Triangles - Not Squares for deterministic convergence
- Insure redundant L3 paths to avoid black holes
- Only create peers on links that you intend to use as transit

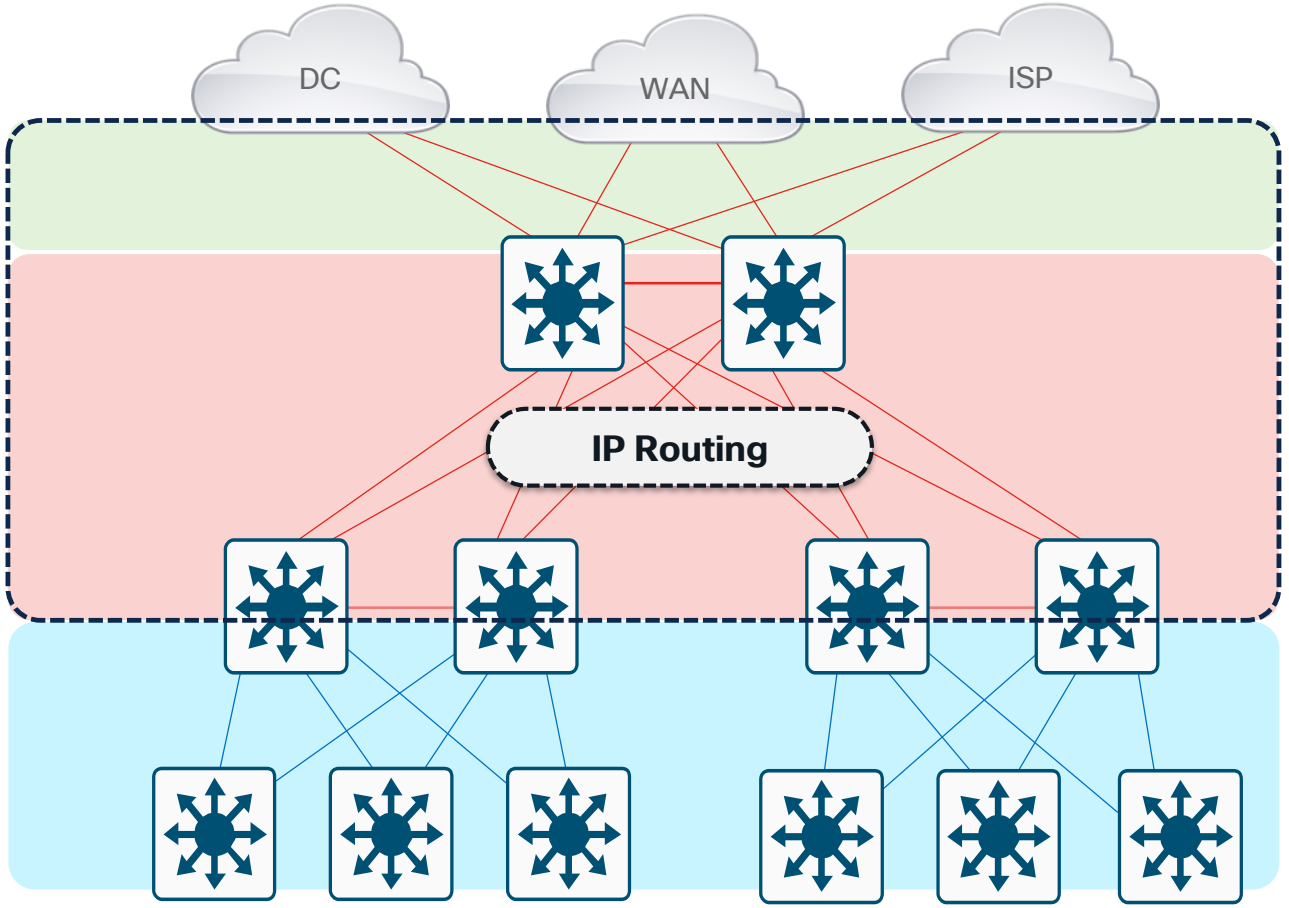

L3 Best Practices – Build Triangles not Squares

Deterministic vs. Non-Deterministic

Squares: Link/Box Failure Requires Routing Protocol Convergence

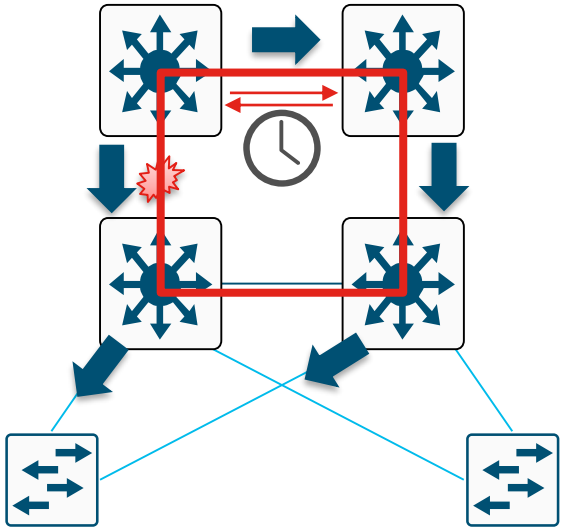

Triangles: Link/Box Failure Does not Require Routing Protocol Convergence

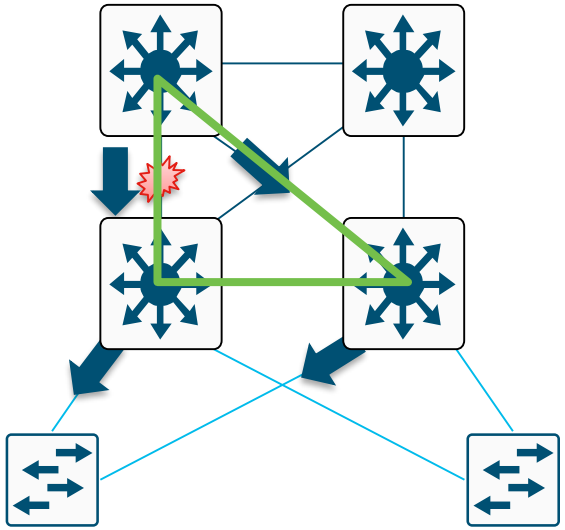

- Layer 3 redundant equal cost links support fast convergence
- Hardware based—fast recovery to remaining path
- Convergence is extremely fast (dual equal-cost paths: no need for IGP to recalculate a new path)

First Hop Redundancy

Hot-Standby Routing Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP)

- A pair of L3 routers function as one virtual router by sharing one virtual IP address and one virtual MAC address
- One L3 router is elected as “Active” and performs packet forwarding for local hosts
- The other routers are elected as “Standby” in case the Active router fails
- Standby routers stay idle and do not participate in packet forwarding
  - Use alternating Active/Standby routers for different VLANs (known as Load-Splitting)

- www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhrp/configuration/15-mt/fhp-15-mt-book/fhp-hsrp-v2.html
- www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhrp/configuration/15-mt/fhp-15-mt-book/fhp-vrrp.html

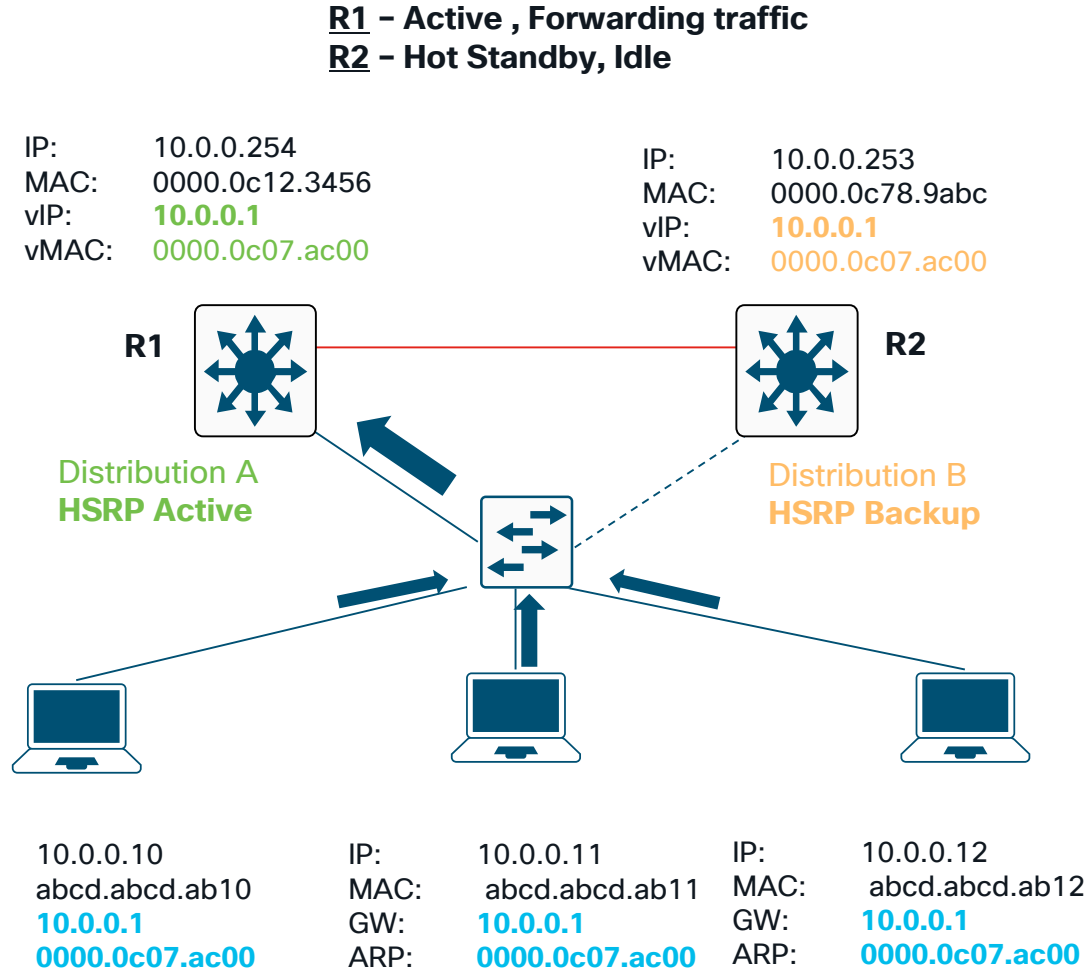

First Hop Redundancy with Load Balancing

Cisco Gateway Load Balancing Protocol (GLBP)

- Each member of a GLBP redundancy group owns a unique virtual MAC address for a common IP address/default gateway
- When end-stations ARP for the common IP address/default gateway they are given a load-balanced virtual MAC address
- Host A and host B send traffic to different GLBP peers but have the same default gateway

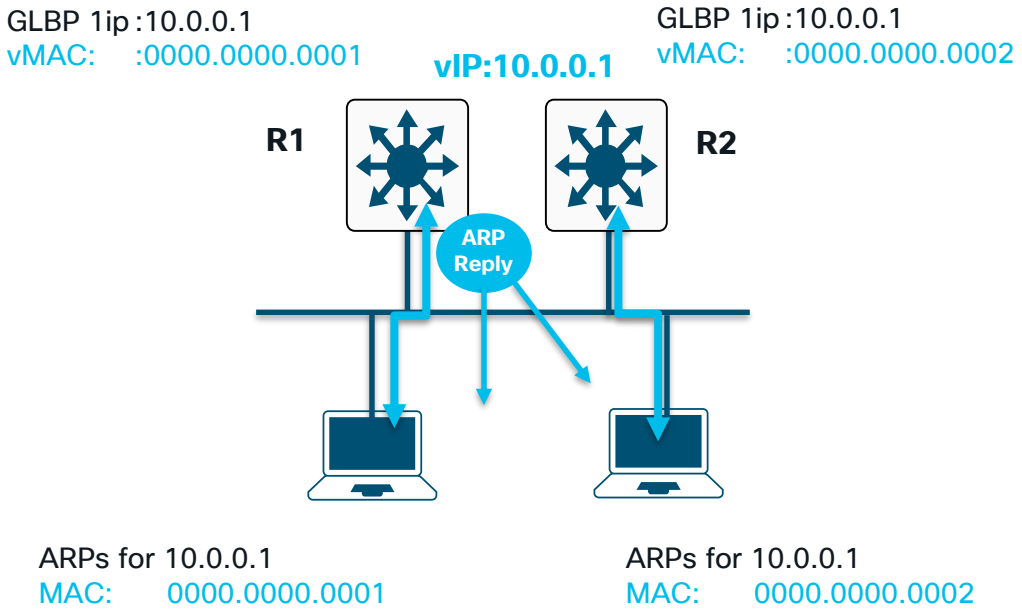

Redundancy & Protocol Interaction

Layer 2 and 3 - Why Use Routed Interfaces

L3 routed interface provides faster convergence than L2 switch port with an associated L3 SVI

1. Link Down
2. Interface Down
3. Routing Update

1. Link Down
2. Interface Down
3. Autostate
4. SVI Down
5. Routing Update

~ 8 msec loss

```
21:38:37.042 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet3/1, changed state to down

21:38:37.050 UTC: %LINK-3-UPDOWN: Interface GigabitEthernet3/1,
changed state to down

21:38:37.050 UTC: IP-EIGRP(Default-IP-Routing-Table:100): Callback:
route_adjust GigabitEthernet3/1
```

~ 150-200 msec loss

```
21:32:47.813 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet2/1, changed state to down

21:32:47.821 UTC: %LINK-3-UPDOWN: Interface GigabitEthernet2/1, changed
state to down

21:32:48.069 UTC: %LINK-3-UPDOWN: Interface Vlan301, changed state to down

21:32:48.069 UTC: IP-EIGRP(Default-IP-Routing-Table:100): Callback: route,
adjust Vlan301
```

Bidirectional Forwarding Detection

BFD

- Detect faults between 2 routers
  - Fast (reaction time in milliseconds)
  - Single mechanism to signal upper-layer routing protocols (ISIS, BGP, OSFP, Static) that link is down
    - faster than the DEAD timer of that protocol
  - Works on directly-connected (single hop) routers, as well as routers separated by an L2 overlay (Metro Ethernet, MPLS, VPLS/Pseudowire, etc.)
  - Uses fast exchange of IP/UDP packets
    - port 3784 for control
    - port 3785 for echo
- Supports single-hop and multi-hop

The official recommendation for Catalyst 9000 switches:

- **250ms x3 for physical interfaces**
- **750ms x3 for SVI**

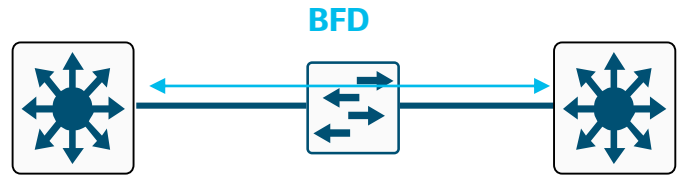


```
interface Gig1/0/1
 ip address 1.1.1.1 255.255.255.0
 bfd interval 250 min_rx 250 multiplier 3
 ip ospf 1 area 0

router ospf 1
 bfd all-interfaces
```

Why You Want to Summarize

Reduce the Complexity of IGP Convergence

- It is important to **force summarization** at the distribution **towards the core**
- For return path traffic an OSPF or EIGRP re-route is required
- By **limiting the number of peers** an EIGRP router must query or the number of LSAs an OSPF peer must process we can optimize his reroute
  - For **EIGRP** if we summarize at the Distribution, we stop queries at the core boxes for an Access layer flap
  - For **OSPF** when we summarize at the Distribution (area border or L1/L2 border), flooding of LSAs is limited to the Distribution: SPF now deals with one LSA not three.

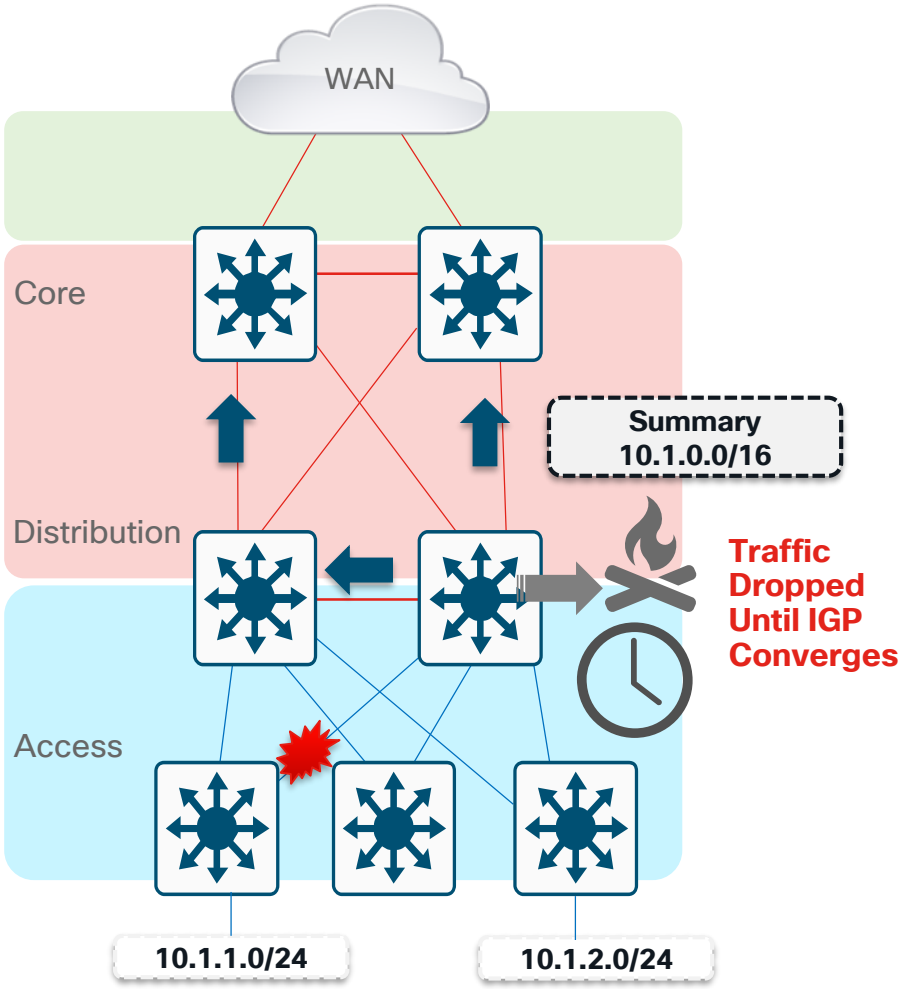

Asymmetric Routing (Unicast Flooding)

Affects redundant topologies with shared L2 access

- One path upstream and two paths downstream
- CAM table entry ages out on standby HSRP
- Without a CAM entry packet is flooded to all ports in the VLAN

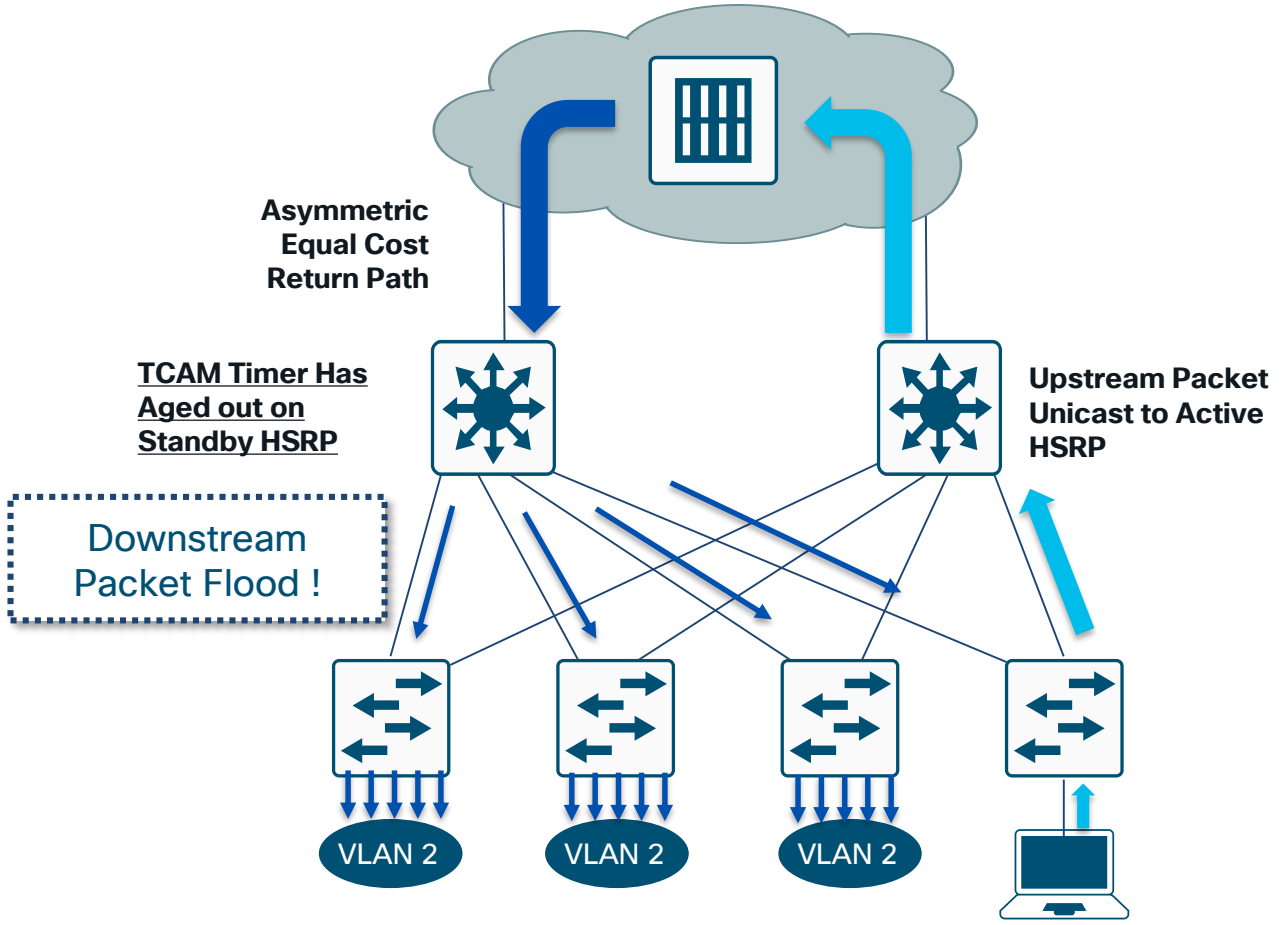

Best Practices Prevent Unicast Flooding

- Assign one unique data and voice VLAN to each access switch
- Traffic is now only flooded down one trunk
- Access switch unicasts correctly; no flooding to all ports
- If you have to:
  - Tune ARP and CAM aging timers; CAM timer exceeds ARP timer
  - Bias routing metrics to remove equal cost routes

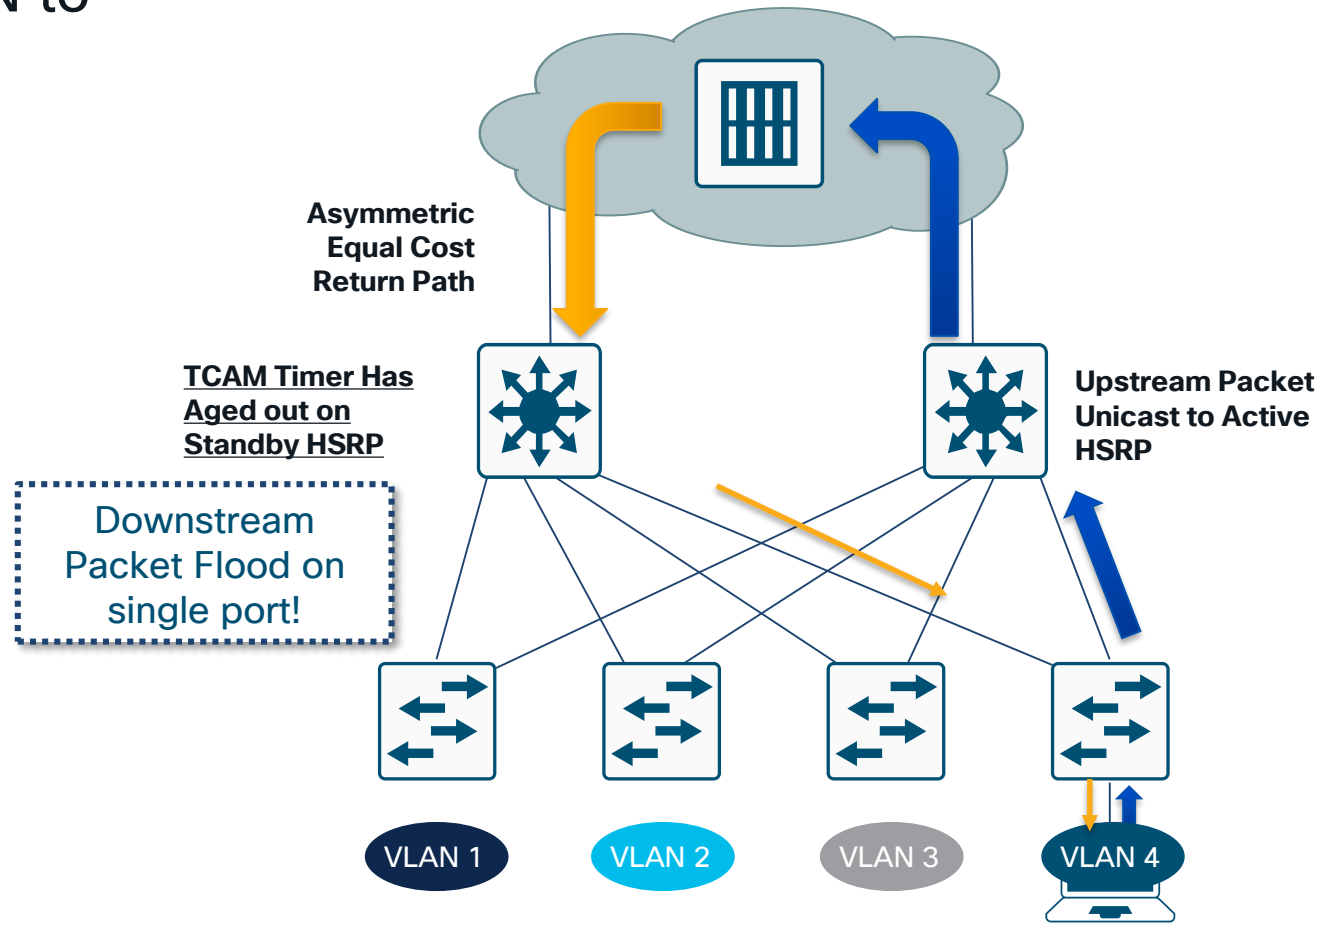

Distribution Interconnect

Best Practices - Summary

Core

Distribution

Access

- ✓ Summarize routes towards Core
- ✓ Limit redundant IGP peering
  - User EtherChannels
  - Passive interfaces to Access
- ✓ HSRP Active tuning
- ✓ Set Trunk mode on/no-negotiate
- ✓ Set EtherChannel mode on/auto
- ✓ STP Root tuning
- ✓ RootGuard or BPDU-Guard
- ✓ Limit protocols on Access ports:
  - Enable PortFast
  - Disable Trunking
  - Disable EtherChannel
- ✓ Use Port Security features

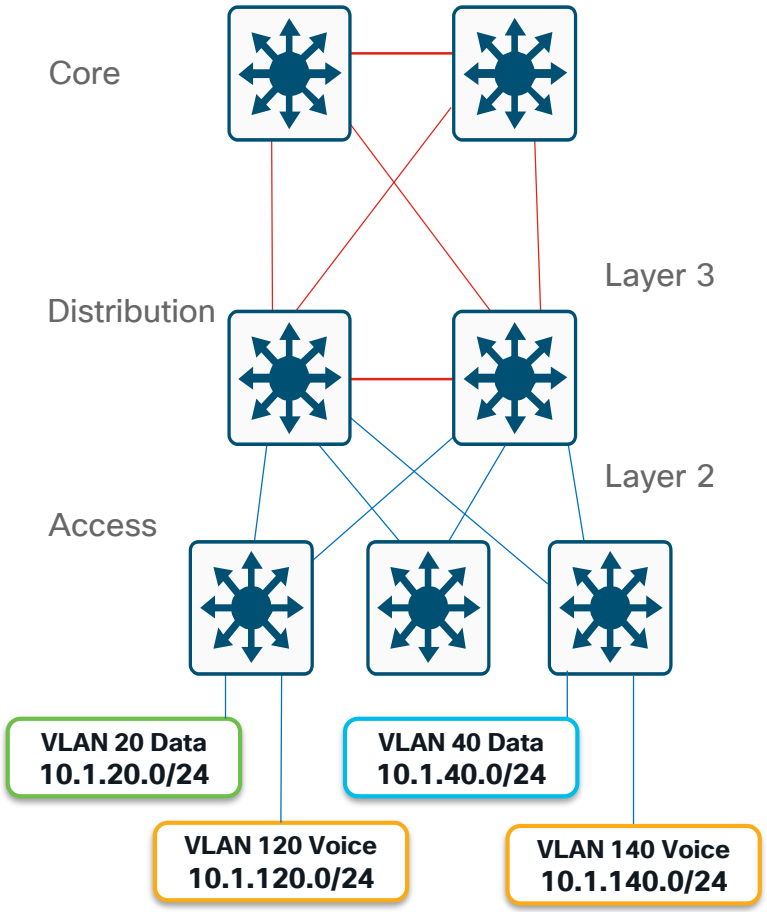

Campus Design Considerations

Chapter 2

Shawn Wargo
Principal Tech Marketing Engineer

Session Agenda

Design Fundamentals | **Design Considerations**

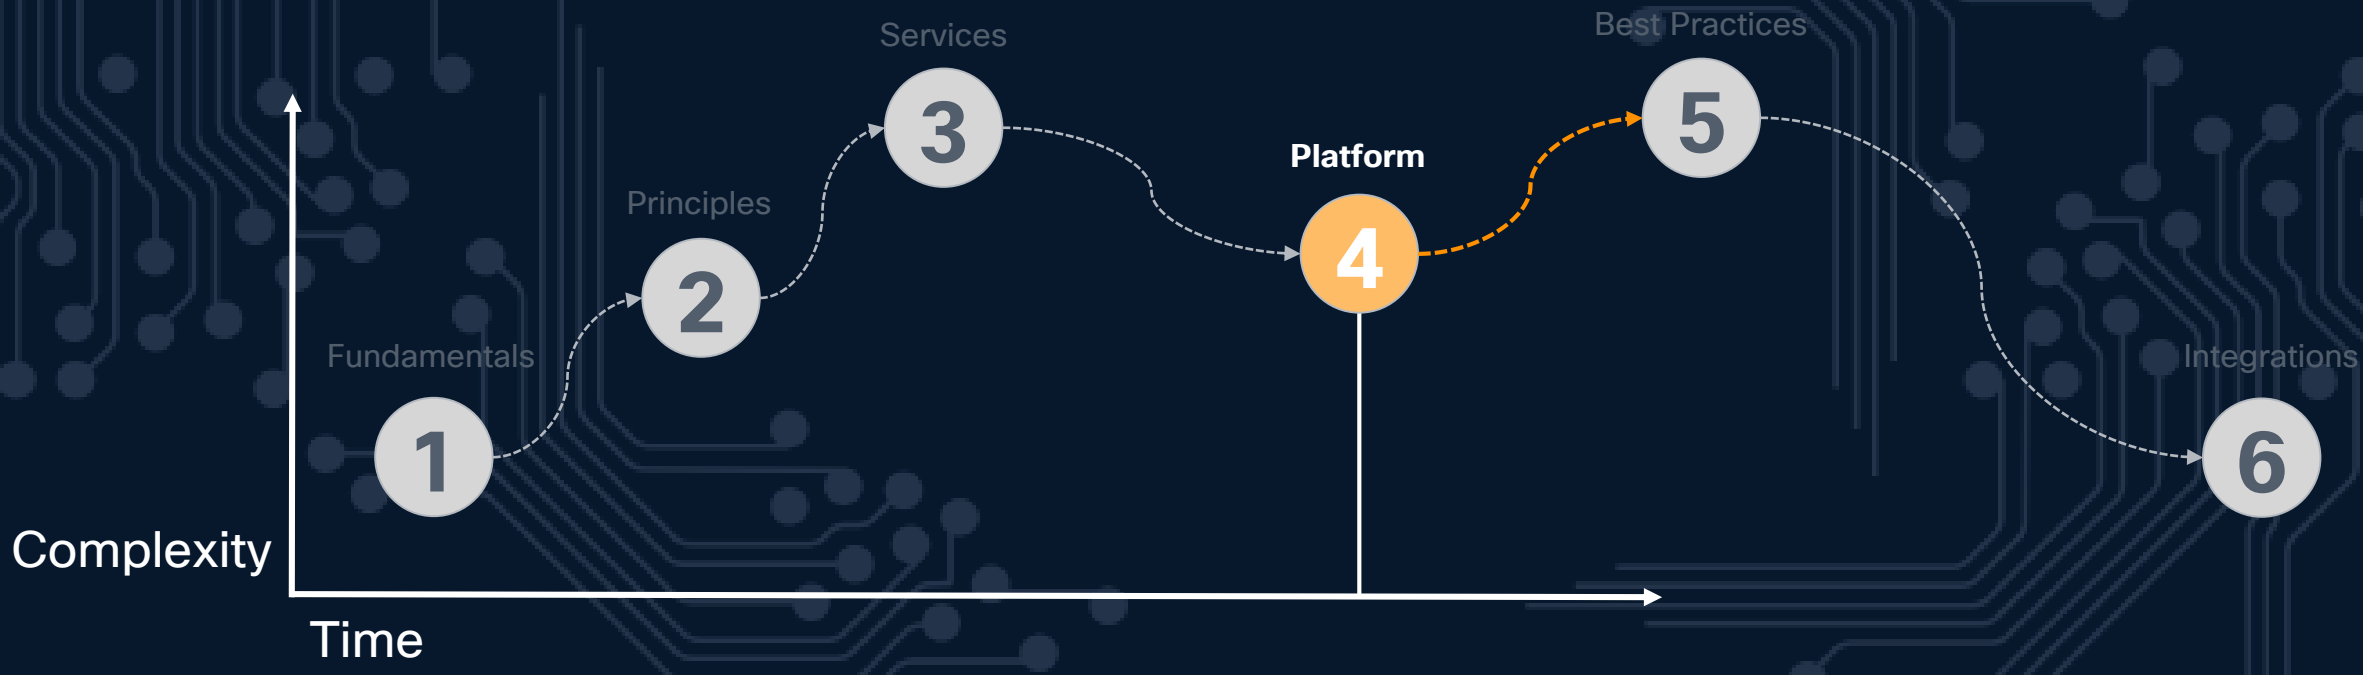

Platform Design

- **Chassis Considerations**
  - **Catalyst 9K** (Overview)
  - **Modular vs. Fixed**
- **Cabling Considerations**
- **Feature Considerations**

Cisco Catalyst 9000 Switching Portfolio

2025+ NEW

One Family from Access to Core – Common Hardware & Software

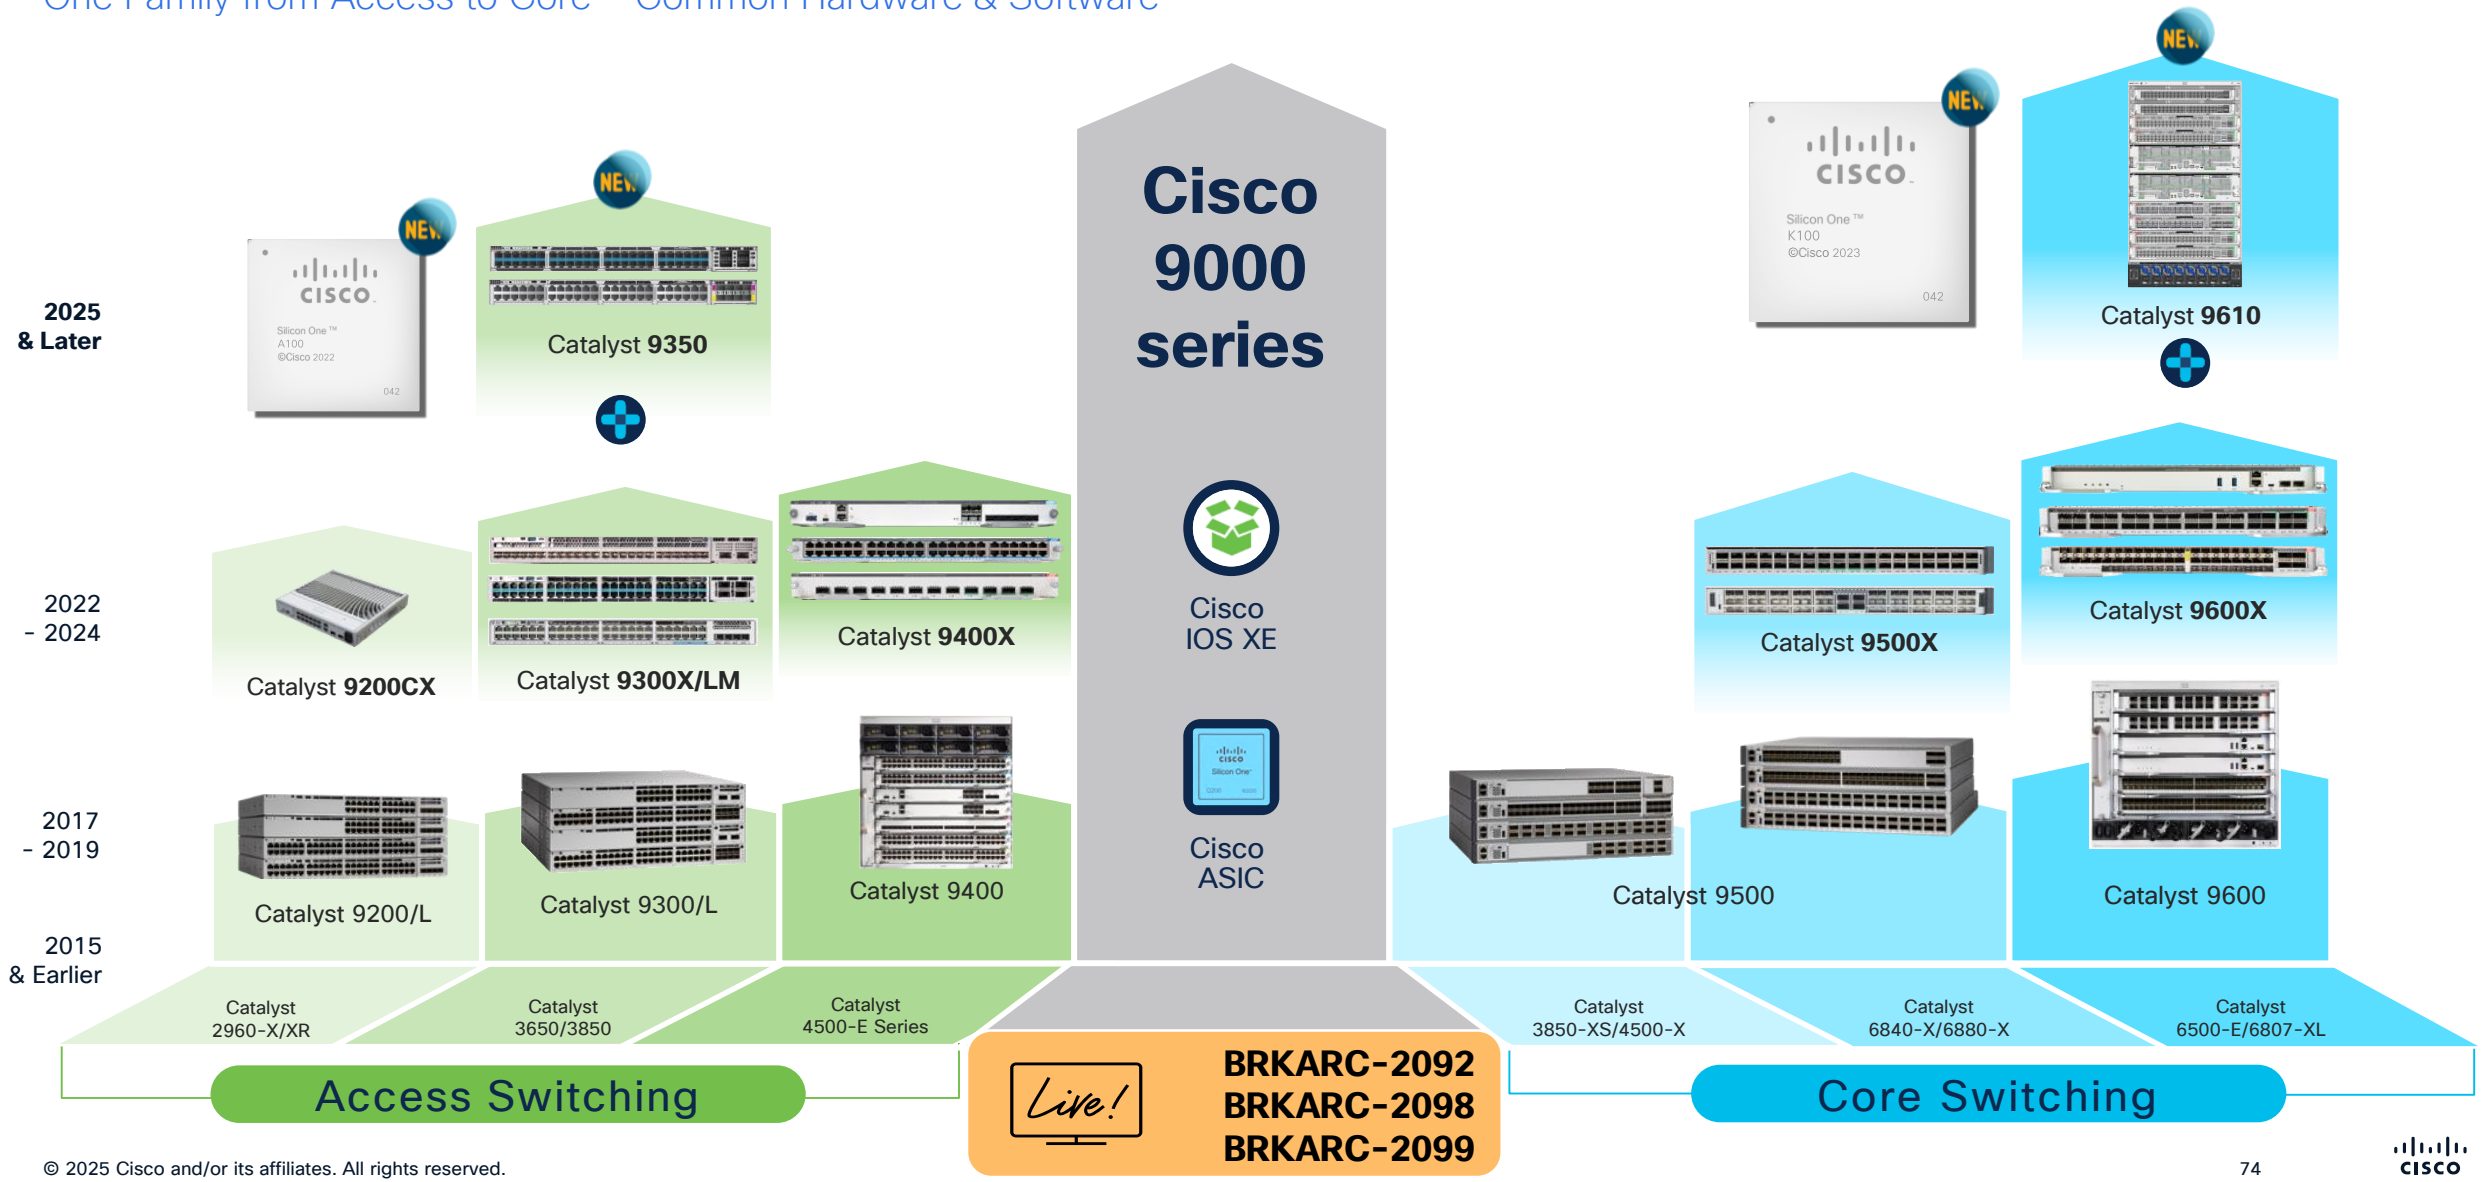

Software vs. Hardware

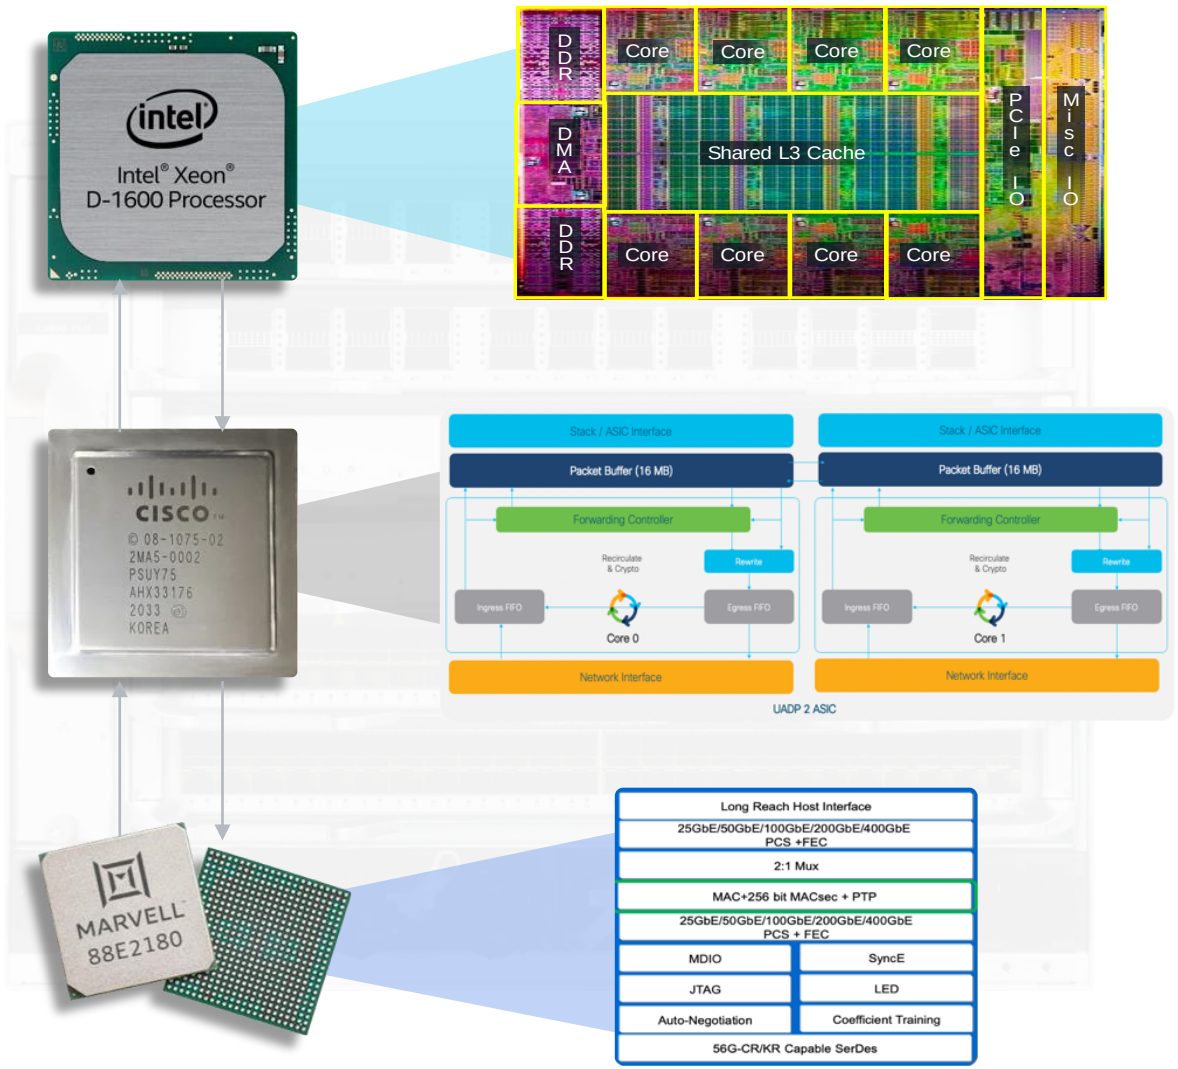

CPU/DRAM

Where the OS “software” runs. Includes control-plane, data-plane and system-management functions.

- **OS layer** – IOSXE (IOSd) and Features, etc.
- **System layer** – FMAN, CMAN, IOMD, FED, etc.

ASIC(s)

Where the “hardware” processing of traffic & services runs. Uses forwarding and state tables programmed by the software.

- **Forwarding** – L2, L3, ECMP, Encap, etc.
- **Services** – ACLs, QoS, Analytics, Encryption, etc.

Stub/PHY(s)

Transforms electrical and optical signals, splits or combines signals, and other various “physical” layer functions, such as encryption and timestamping.

Modular vs. Fixed Platforms

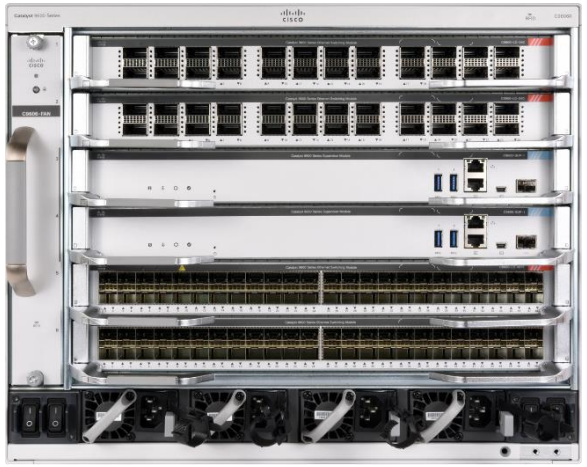

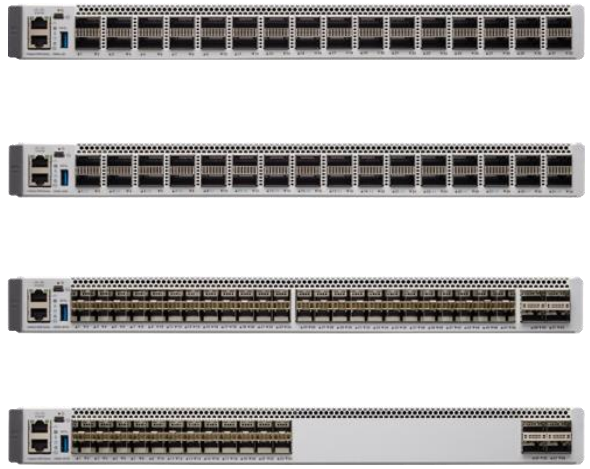

Modular

PROs	CONs
• **More Flexible**	• **More Complex**
• Longer Life-Cycle	• BW limit by Chassis
• Higher Port Density	• Slow(er) Dev & Test
• More Power/Cooling	• Lower MTBF
• Redundant Processors	• Higher COGs

Fixed

PROs	CONs
• **Less Complex**	• **Less Flexible**
• Swap Chassis for BW	• Shorter Life-Cycle
• Faster Dev & Test	• Lower Port Density
• Higher MTBF	• Less Power/Cooling
• Lower COGs	• Single Processor

Modular Platform Features & Benefits

Redundancy, Expansion, Efficiency & Flexibility

Highest Resiliency

Redundant Supervisors

StackWise® Virtual

Easy Upgrades with ISSU & GIR

Redundant Fans (Fan-Tray)

Redundant PSUs (1:1, N+1)

Highest Flexibility

SUP1 for Small Designs

SUP2/XL for Large Designs

Custom ASIC Scale Templates

Traditional Multi-Layer Designs

Fabric Overlay Designs

Highest Efficiency

Lowest Watts per Port

3000W Power Supplies

Titanium Rated (95%) PSUs

AC and/or DC Power

Configurable Power Priority

Longest Lifecycle

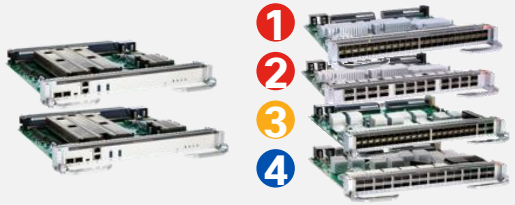

Start w/ SUP1 & few Gen1 LCs

Add Gen1 LCs as Access grows

Replace SUP1 with SUP2

Gen1 LCs get a 2X boost

Add new Gen2 LCs as Core grows

Most Port Options

Mixes of RJ45, SFP & QSFP

C9600-LC-40YL4CD
40x 50G SFP + 2x 100G + 2x 400G QSFP

C9600X-LC-32CD
32x 100G or 24x + 8x 400G QSFP

C9400-LC-48XS
48x 1/10G SFP

C9400-LC-48HX
48 x 10G mGig + UPOE®

Modular Design for Large Campus

BRKENS-1500

Architecture Perspective – Full Mesh vs. Hierarchical Design

Fixed System Design

- Static
- Costly
- Complex

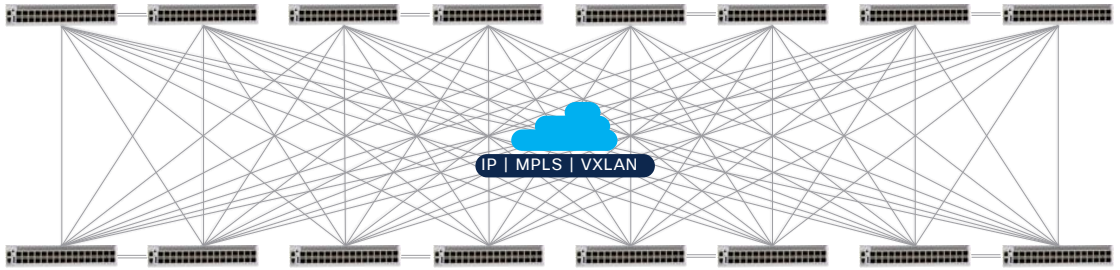

Modular System Design

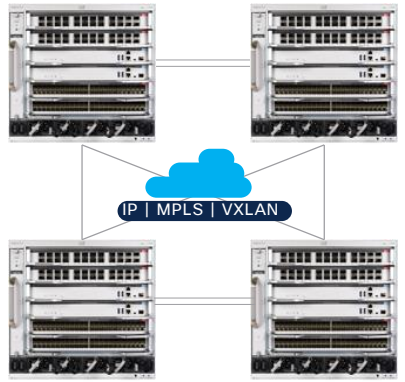

- Simple
- Scalable
- Sustainable

Modular System Benefits

Sustainable

Reduce Energy Demand
Reduce Carbon footprint
Environmental efficient

Cost

Reduce cost – CAPEX | OPEX
License & Service Management
Reduce product life-cycle TCO

Operation

Proven for large Enterprise
Day 0 – N scalable architecture
Simplified Tools and Management

Flexible

Pay-As-You-Grow model
Elastic Aggregation. Static Core.
Simple and large L2 boundaries

Resilient

Non-stop communication
Protected network performance
Reduced MTTR and MTBF

Platform Design

- Chassis Considerations
- Cabling Considerations
  - Link Scale
  - Why 2.5, 5 & 10G?
  - Why 25G & 50G?
  - Why 100G & 400G?
- Feature Considerations

Cisco Multi-Gigabit Ethernet

Providing greater Copper speeds with existing RJ45 cables

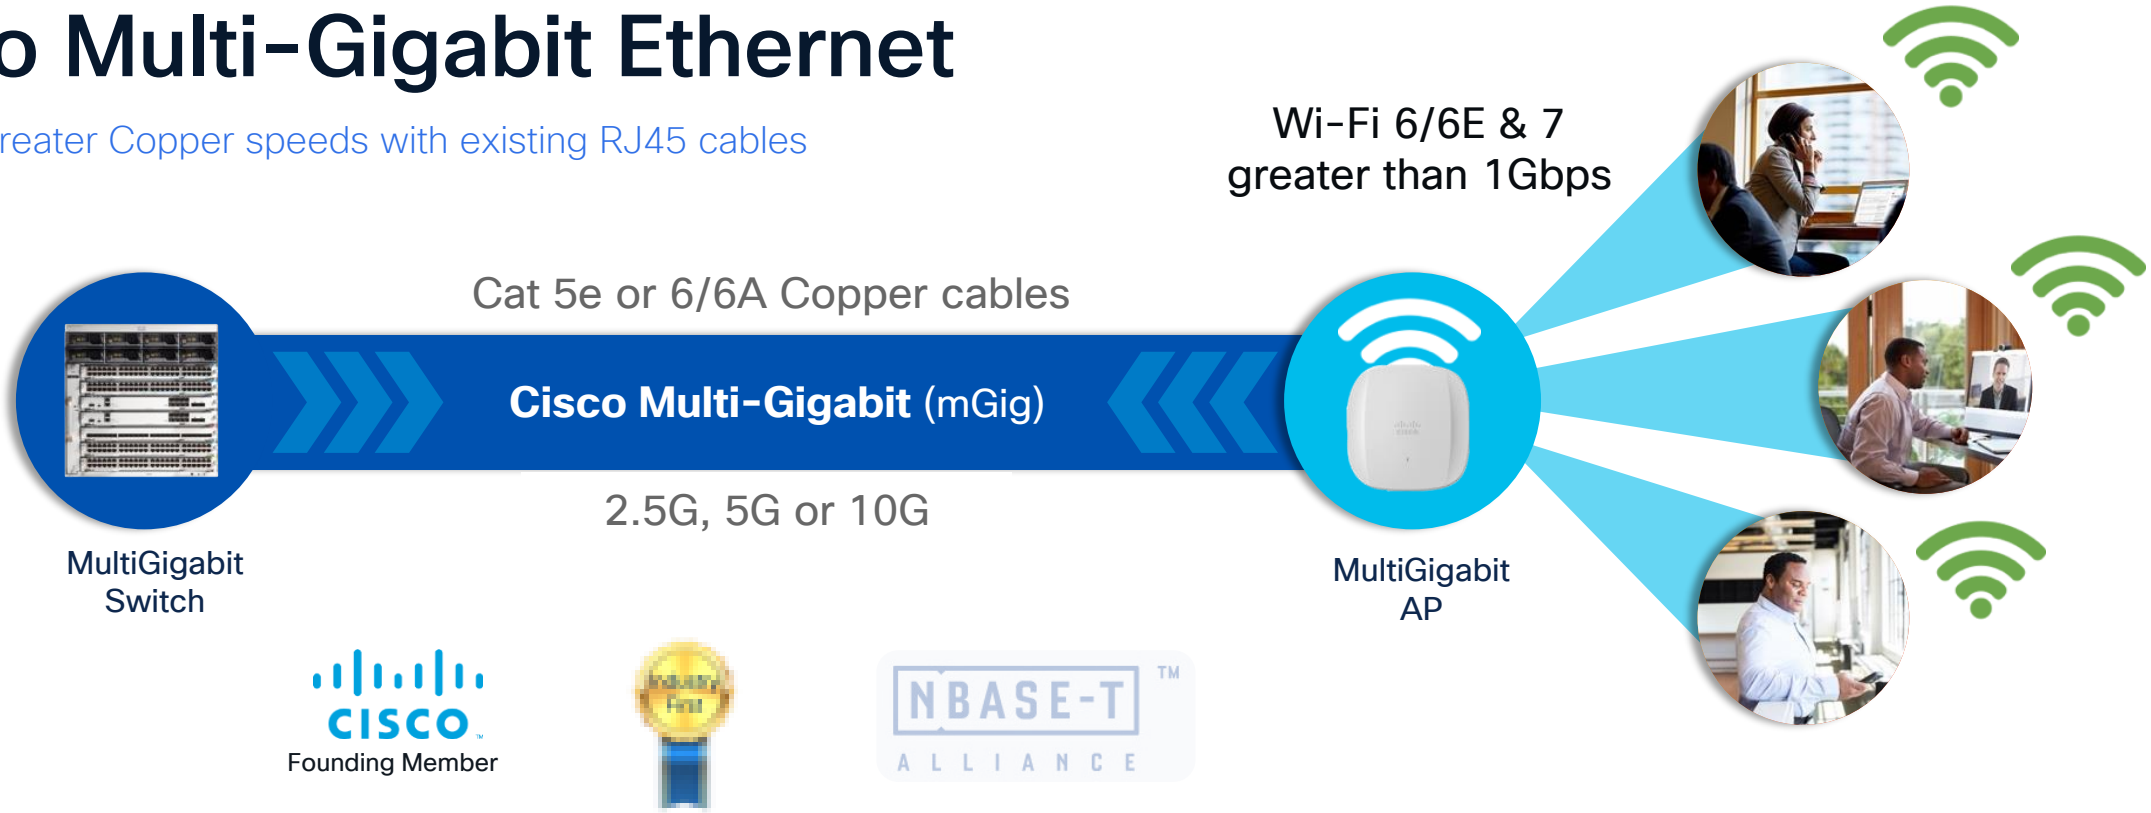

Founding Member

A game-changing innovation for Enterprise LAN to evolve beyond 1G

Enables **2.5Gbps** to **5Gbps** up to 100 meters on legacy **Category 5e** cables

Supports all PoE standards up to **90W**

Delivers up to **5x faster speeds** in enterprise **without replacing cabling** infrastructure

WiFi6E/7 driving Multi-Gigabit & UPOE

Future Proof Speeds and More Power Over Ethernet

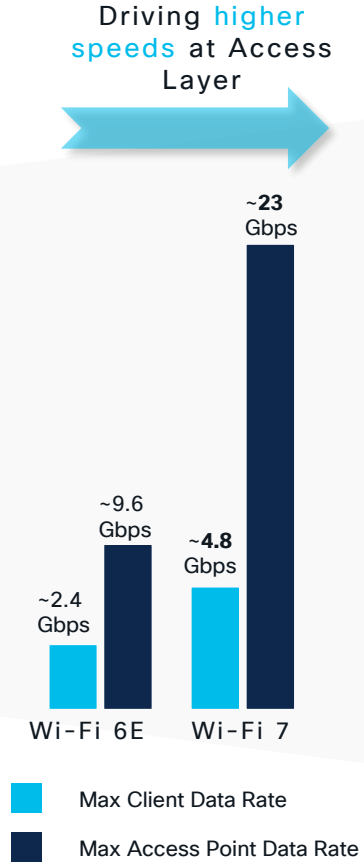

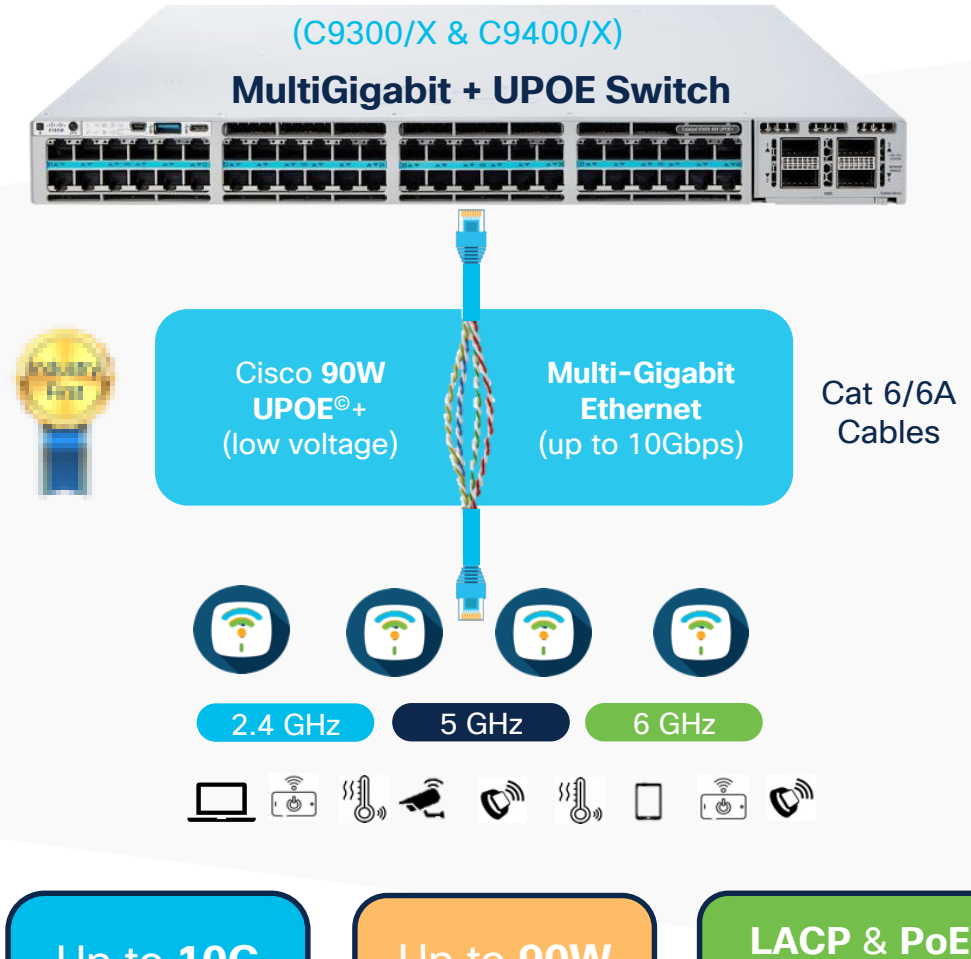

Up to **10G** Bandwidth

Up to **90W** UPOE

LACP & **PoE** Redundancy (Dual 10G ports)*

25GE & 50GE – A Better Alternative

Provide a seamless migration path from 1/10GE SFP

Designation	Speed
L	50GE
Y	25GE
X	10GE

Cisco C9600

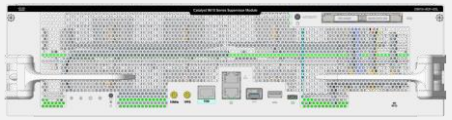

C9610-SUP-3

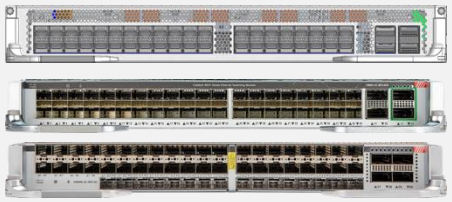

LC-48YL, 40YL4CD & 56YL4C

Cisco C9500

C9500X-60L4D

C9500-48Y4C

Cisco C9400

C9400X-SUP2XL

C9400-SUP1XL-Y

Cisco C9300

C9350-NM-8Y

C9300-NM-2Y & C9300X-NM-8Y

- Reduced CapEx through reuse of existing cabling
- Single-Lane Optics provide similar port densities of 10G
- Gradual migration options with support of Dual-Rate Optics
- Reduced OpEx through savings in power & cooling

Introducing 50GE – First in Campus

Provides seamless migration path from 10/25GE

Designation	Speed
D	400GE
C	100GE
L	50GE
Y	25GE
X	10GE

Supported on

SUP-2 & LC-48YL*

SUP-2 & LC-40YL4CD
LC-56YL4C

C9500X-60L4D

IOS-XE 17.12.1

Compatible 50G SFPs

50G SFP56

SFP-50G-CUxM (1/1.5/2/2.5/3/4/5M)
SFP56-25/50G-SL (7/10/20/30M)
SFP-50G-SR-S (70/100M)
SFP-50G-LR-S (10/15KM)

www.cisco.com/c/en/us/products/collateral/interfaces-modules/transceiver-modules/50gbase-sfp56-modules-ds.html

Reduced CAPEX

Reuse of existing cable infrastructure

Reduced OPEX

Optimized power and cooling

Single Lane Optics

Same port densities as 10/25G offerings

Dual Rate Optics

Staggered/Gradual upgrades and future proofing

Use Case

- Less oversubscription from Distro to Core
- Increase SFP-based SVL bandwidth

100GE & 400GE - A Better Alternative

Provide a seamless migration path from 40GE QSFP

Designation	Speed
D	400GE
C	100GE
Q	40GE

Cisco C9600

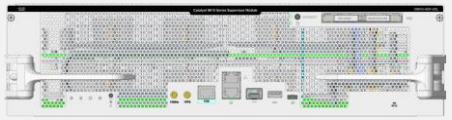

C9610-SUP-3

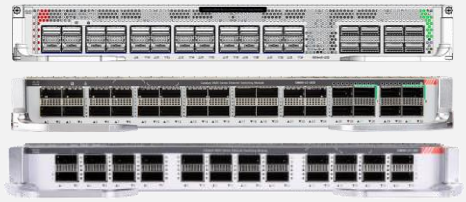

LC-24C & LC-32CD

Cisco C9500

C9500X-60L4D

C9500-32C

Cisco C9400

C9400X-SUP2XL

C9400-SUP1XL

Cisco C9300

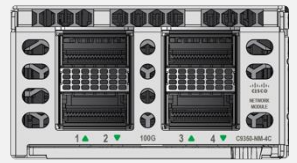

C9350-NM-2C & 4C

C9300X-NM-2C & 4C

- Reduced CapEx through reuse of existing cabling
- Single-Lane Optics provide similar port densities of 40G
- Gradual migration options with support of Dual-Rate Optics
- Reduced OpEx through savings in power & cooling

Introducing 400GE – First in Campus

Provides seamless migration path from 40/100GE

Designation	Speed
D	400GE
C	100GE
L	50GE
Y	25GE
X	10GE

Supported on

SUP-2 & LC-32CD

SUP-2 & LC-40YL4CD

C9500X-28C8D & 60L4D

Compatible 400G QSFPs

400G QSFPDD

QDD-400-CUxM (1/1.5/2/3M)
QDD-400-AOCxM (1/3/5/7/15/30M)
QDD-400G-SR8 (100M)
QDD-400G-DR4 (500M)
QDD-400G-FR4 (2KM)
QDD-400G-LR4/8 (10KM)

www.cisco.com/c/en/us/products/collateral/interfaces-modules/transceiver-modules/datasheet-c78-743172.html

Reduced CAPEX

Reuse of existing cable infrastructure

Reduced OPEX

Optimized power and cooling

Single Lane Optics

Same port densities as 40/100G offerings

Dual Rate Optics

Staggered/Gradual upgrades and future proofing

Use Case

- High-Speed from Campus to DC or SP
- Increase QSFP-based SVL bandwidth

* No MACsec when using LC-48YL with Supervisor 2

Catalyst 9200/CX, 9300/X & 9400/X

BRKARC-2098

Catalyst 9000 Series Switching Family – Access

Manish Sharma - Technical Marketing, Cisco

This session will cover the platform overview of Cisco Catalyst 9000 Series access switches.

It will share the details of the Catalyst 9000 product portfolio, which will include new additions in fixed and modular access series – Catalyst 9200/CX, Catalyst 9300/X, and Catalyst 9400/X.

The session will talk about the component at the heart of these switches, which is the ASIC. It will also cover common attributes, technologies, and features in the Catalyst 9000 Series switches.

Introducing Cisco C9350 Series Smart Switches NEW

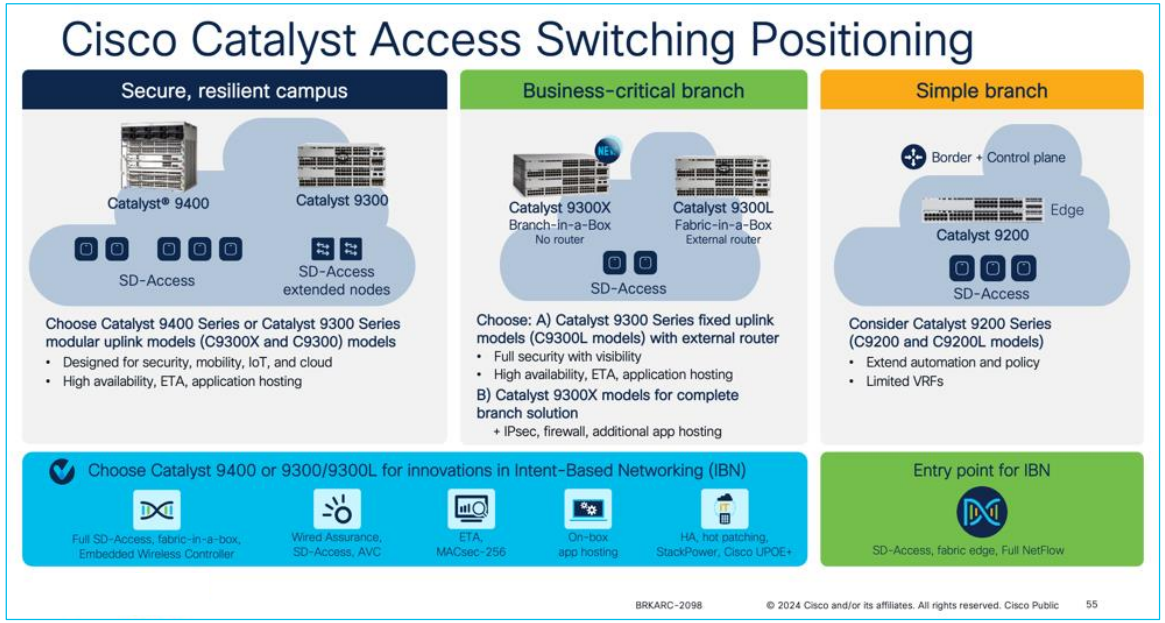

Catalyst 9500/X & 9600/X

BRKARC-2099

Catalyst 9000 Series Switching Family – Core & Distribution

Kenny Lei - Leader Technical Marketing, Cisco

This session will cover the platform overview of Catalyst 9000 Series core and distribution switches.

It will share the details of the Catalyst 9000 Series product portfolio, which will include new additions in fixed and modular core and distribution switching series - Catalyst 9500/X and Catalyst 9600/X.

The session will discuss the component at the heart of these switches, which is the ASIC, and it will also cover common attributes, technologies, and features in Catalyst 9000 switches.

Introducing Cisco C9610 Series Smart Switches NEW

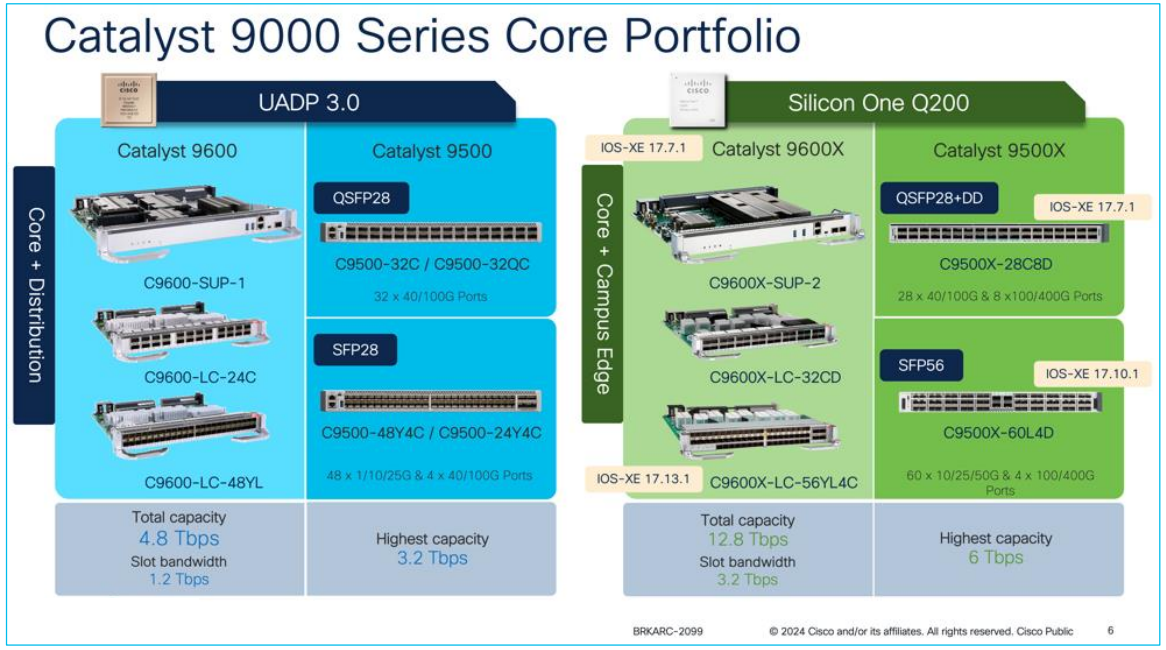

Also check out BRKARC-2668

Current 3-Tier Campus Design

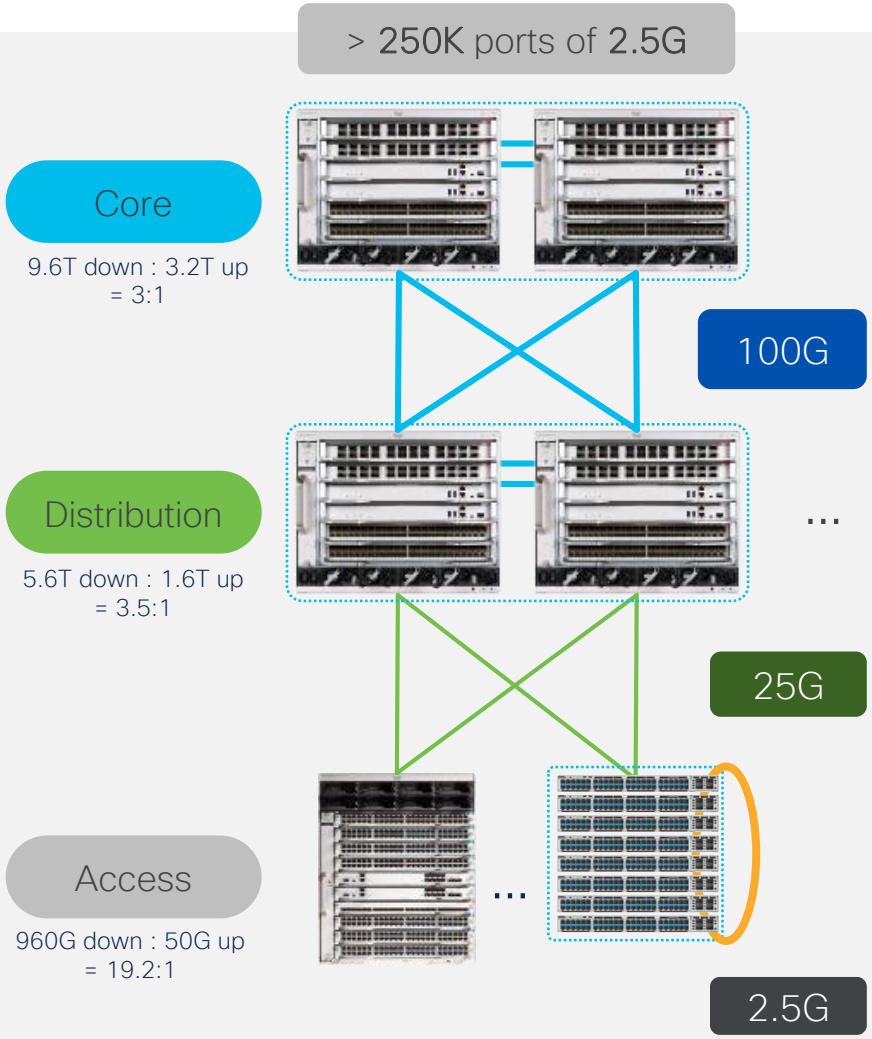

- **C9606R** Switch with **4x C9600-LC-32CD** for 100G connection to the Distribution layer, WAN, and between Core devices
- Soft oversubscription ratio of ~**4:1 from Distribution to Core**

Core

- **C9606R** with **4x C9606-LC-56YL4C**: 100G as the uplink connection to the Core and SVL links, and 10G/25G as the connections to the Access layer
- Soft oversubscription ratio of ~**20:1 from Access to Distribution**

Distribution

- **C9410 with SUP2/2XL** and 1/2.5 mGig downlink modules and 10/25G uplink (sup or modules)
- **8x C9300 StackWise-480** and 1/2.5 mGig fixed downlinks and 10G/25G uplink modules

Access

384 ports x **112** Access Pods x **6** Distribution Blocks = **258K** ports of 2.5G

Simplified 3-Tier Campus Design

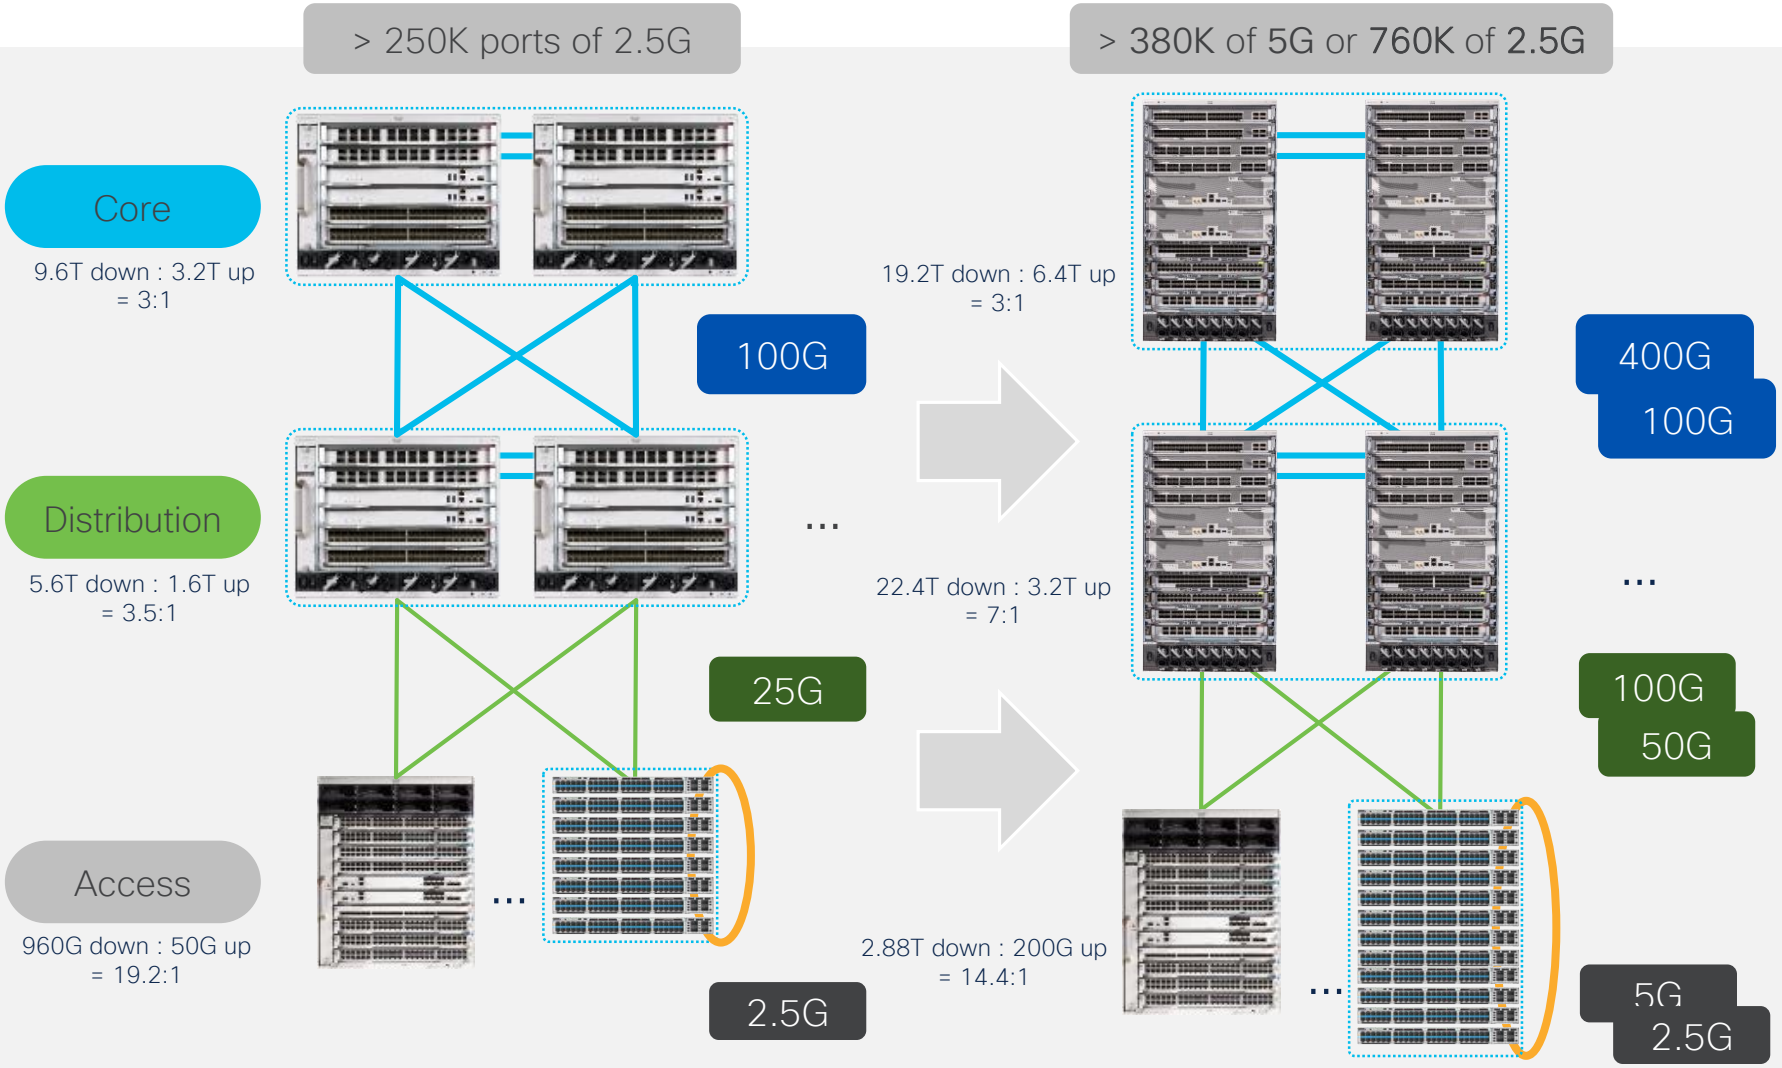

Core

- Double Bandwidth
- Double the number of Distribution blocks
- C9610R with SUP-3/3XL and 8x LC-32CD

Distribution

- Double Bandwidth
- Double the number of Access Pods
- C9610R with SUP-3/3XL and 8x LC-56YL4C or 32CD

Access

- 12x NG-Stackwise
- C9350 Access PODs with 2.5G or 5G downlinks and 50G or 100G uplinks

576 ports x **112** Access Pods x **6** Distribution Blocks = **387K** ports of 5G

Platform Design

- **Chassis Considerations**
- **Cabling Considerations**
- **Feature Considerations**
  - **L2** (Unicast & Multicast)
  - **L3** (Unicast & Multicast)
  - **Security** (AAA & ACL)
  - **Quality of Service** (QoS)
  - **NetFlow** (AVC & XDR)

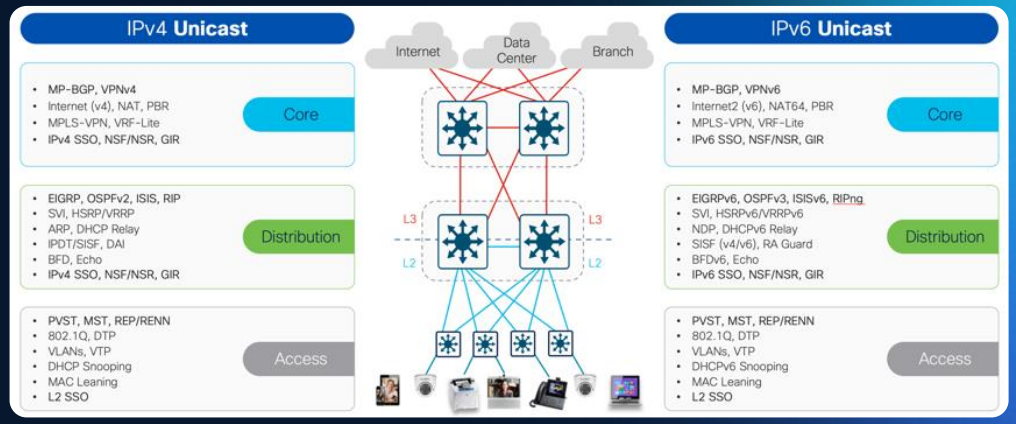

Campus Networks

BRKENS-1500

L2/L3 Unicast Technologies

IPv4 Unicast

- MP-BGP, VPNv4
- Internet (v4), NAT, PBR
- MPLS-VPN, VRF-Lite
- IPv4 SSO, NSF/NSR, GIR

Core

- EIGRP, OSPFv2, ISIS, RIP
- SVI, HSRP/VRRP
- ARP, DHCP Relay
- IPDT/SISF, DAI
- BFD, Echo
- IPv4 SSO, NSF/NSR, GIR

Distribution

- PVST, MST, REP/RENN
- 802.1Q, DTP
- VLANs, VTP
- DHCP Snooping
- MAC Leaning
- L2 SSO

Access

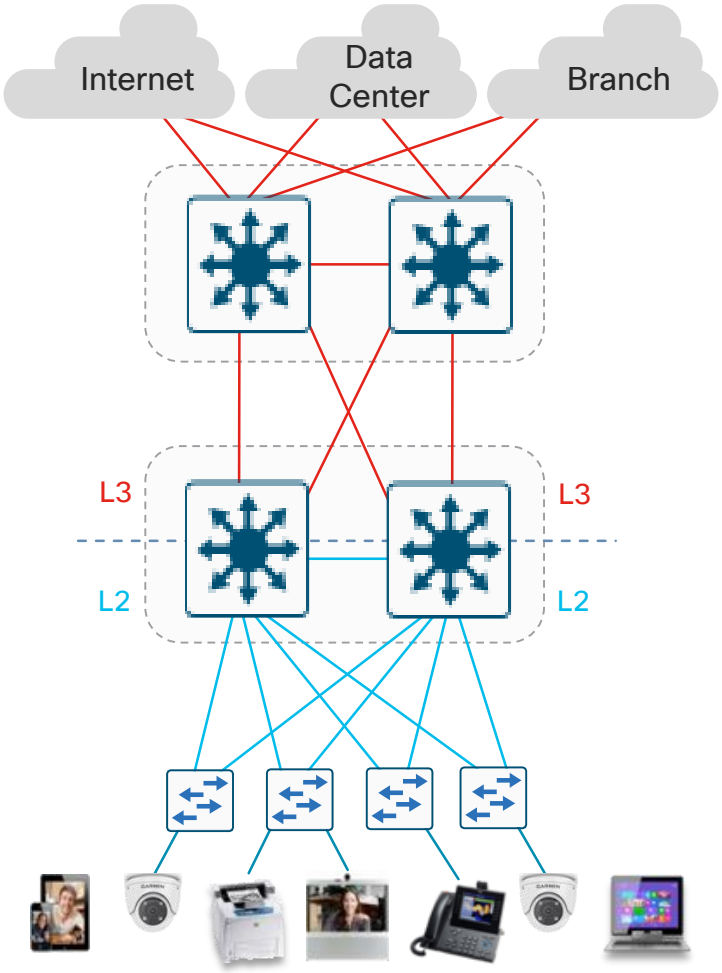

IPv6 Unicast

- MP-BGP, VPNv6
- Internet2 (v6), NAT64, PBR
- MPLS-VPN, VRF-Lite
- IPv6 SSO, NSF/NSR, GIR

Core

- EIGRPv6, OSPFv3, ISISv6, RIPng
- SVI, HSRPv6/VRRPv6
- NDP, DHCPv6 Relay
- SISF (v4/v6), RA Guard
- BFDv6, Echo
- IPv6 SSO, NSF/NSR, GIR

Distribution

- PVST, MST, REP/RENN
- 802.1Q, DTP
- VLANs, VTP
- DHCPv6 Snooping
- MAC Leaning
- L2 SSO

Access

Understanding L2 Scale

MAC Address Scale

Soft recommendation for Access to Distribution ≤ 20:1

- Each unique Endpoint (Host) will have 1x MAC address
  - Access: # **Hosts** = # **MAC**
- All MACs are also learned on Distribution (STP Root)
  - Distro: Sum of # Access

1000-2000 x 20

SUM: **20-40K MACs**

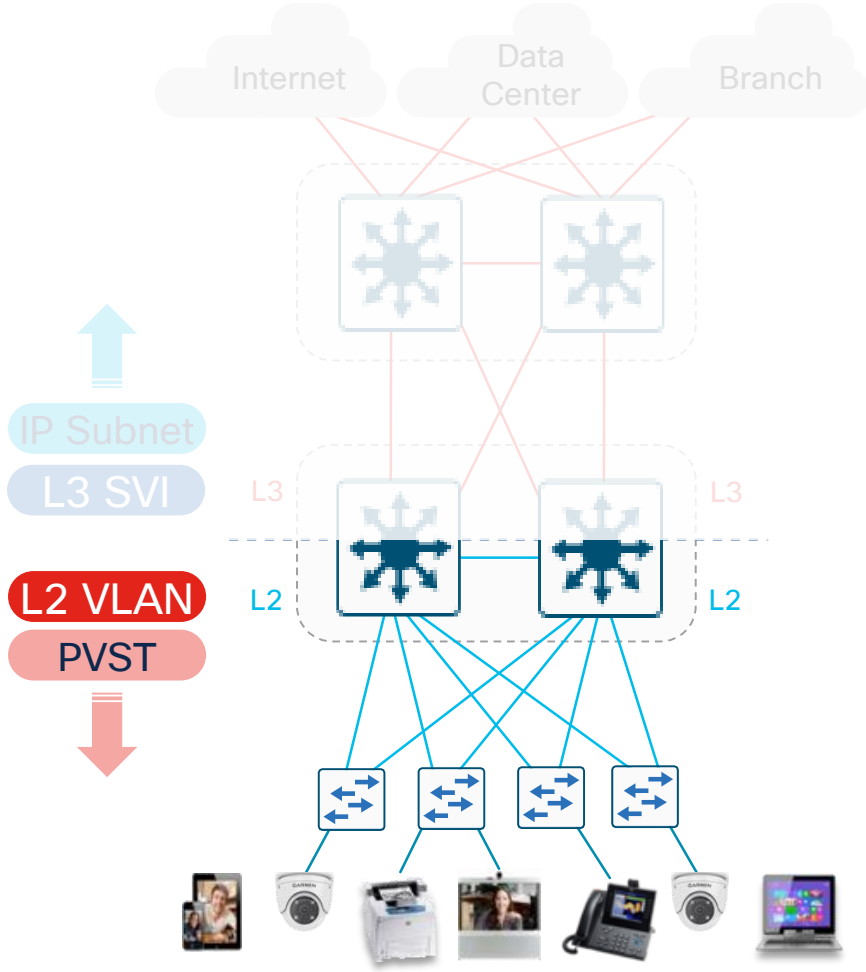

Campus Networks

BRKENS-1500

L2/L3 Multicast Technologies

IPv4 Multicast

Core

- PIM-SM, SSM and Bidir
- AutoRP, BSR RP, MSDP
- MVPN, Multicast VRF-Lite
- Multicast load splitting
- IPv4 multicast HA

Distribution

- Dual-stack IPv4 / IPv6
- PIM-SM, SSM and Bidir
- IGMPv2,v3 snooping
- Stub multicast routing
- PIM BFD
- IPv4 multicast HA

Access

- IGMP v1,v2,v3 snooping
- IPv4 multicast QoS & ACL
- IGMP v1,v2 filtering

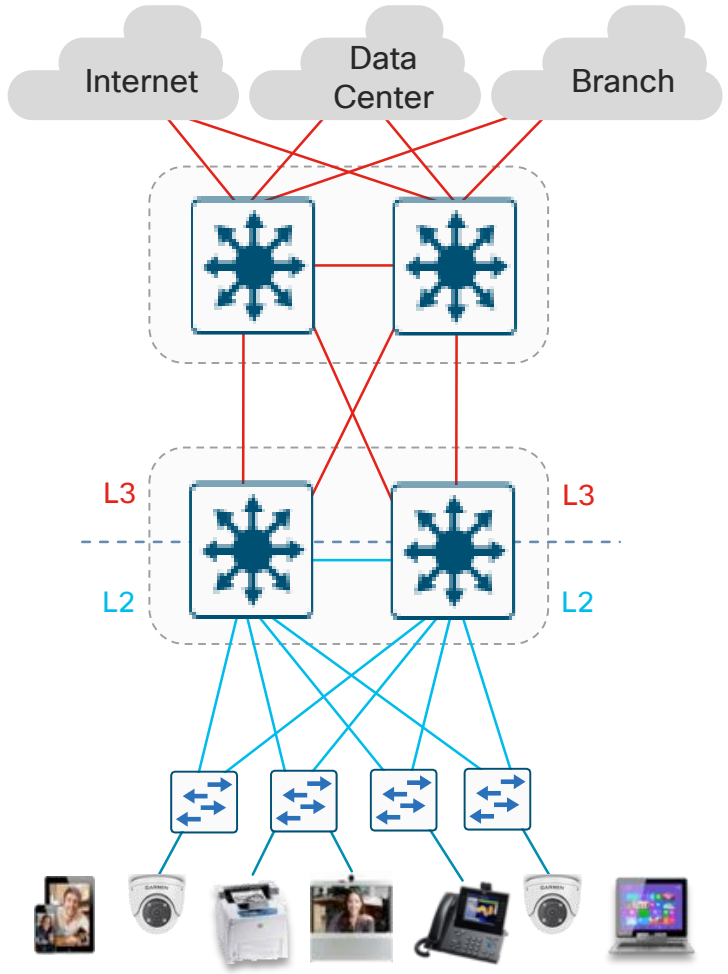

IPv6 Multicast

Core

- PIM-SM and SSM
- IPv6 BSR RP
- IPv6 embedded RP
- IPv6 multicast HA

Distribution

- Dual-stack IPv4 / IPv6
- PIM-SM and SSM
- MLDv1,v2 snooping
- HW register and RPF
- HSRP-aware PIM
- IPv6 multicast HA

Access

- MLD v1,v2 snooping
- IPv6 multicast QoS & ACL
- MLD v1,v2 filtering

Understanding L3 Scale

BRKENS-1500

IP Route Scale

- Each unique Endpoint (Host) will have 1x ARP (and/or 3+ NDP)
- All ARP/NDP resolve on Distribution (L3 SVI)
  - Distro: Sum of # Access = **1-3K x 20 = 20-60K**
  - VLANs: 5-10 per Access = **4-5 x 20 = 100-200**
- All L3 + WAN/DC (x VRF) Subnets on Core
  - Core(Site): Sum of # Distro = **10-20 x 200 = 2K-4K**
  - WAN/DC: Sum of # Sites = **10-20 x 2K = 20K-40K**
  - Internet: Q1 2025 = **~990K IPv4, ~220K IPv6**

Soft recommendation for
Access to Distribution ≤ 20:1
Distribution to Core ≤ 4:1

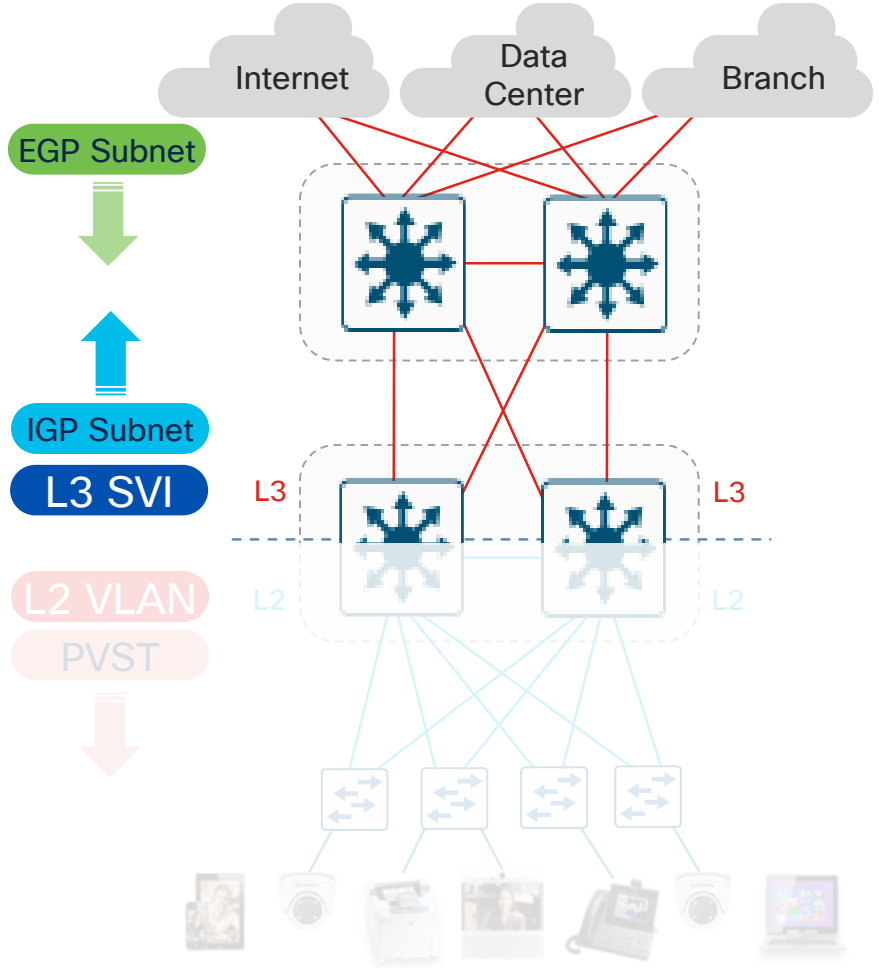

End User Authentication

802.1x Authentication

Live! BRKENS-2062

- Security at the access layer was originally done with techniques like port-security
  - Port security is vulnerable to things like MAC address spoofing, and not flexible or easily traceable.
- IEEE developed the 802.1x standard to address the concerns and create a framework to secure *wired & wireless ports*
  - **Extensible Authentication Protocol** (EAP)
  - **MAC Address Bypass** (MAB)
  - **Web Authentication** (WebAuth)
- Generally - each of these AAA methods serves a specific purpose (depending on the use-case and AAA server).

	802.1x	MAB	WebAuth
802.1x enabled **Employee**	✓	✗	✓
Non-802.1x managed **Device**	✗	✓	✗
802.1x or non-802.1x **Contractor**	✗	✗	✓
802.1x or non-802.1x **IP Phone**	✓	✓	✗

802.1X Host modes

802.1x on switches provides multiple host modes that change the behavior of the number of clients authenticated on the port.

Mode	Behavior
Single Host	Only a single host can connect/authenticate to the port.
Multi Host	A single devices authenticates - all further devices are allowed without authentication.
Multi Domain	Allows one client in the voice VLAN -and- one authenticated client on the data VLAN.
Muti Auth	Allows one client in the voice VLAN and multiple authenticated clients on the data VLAN.

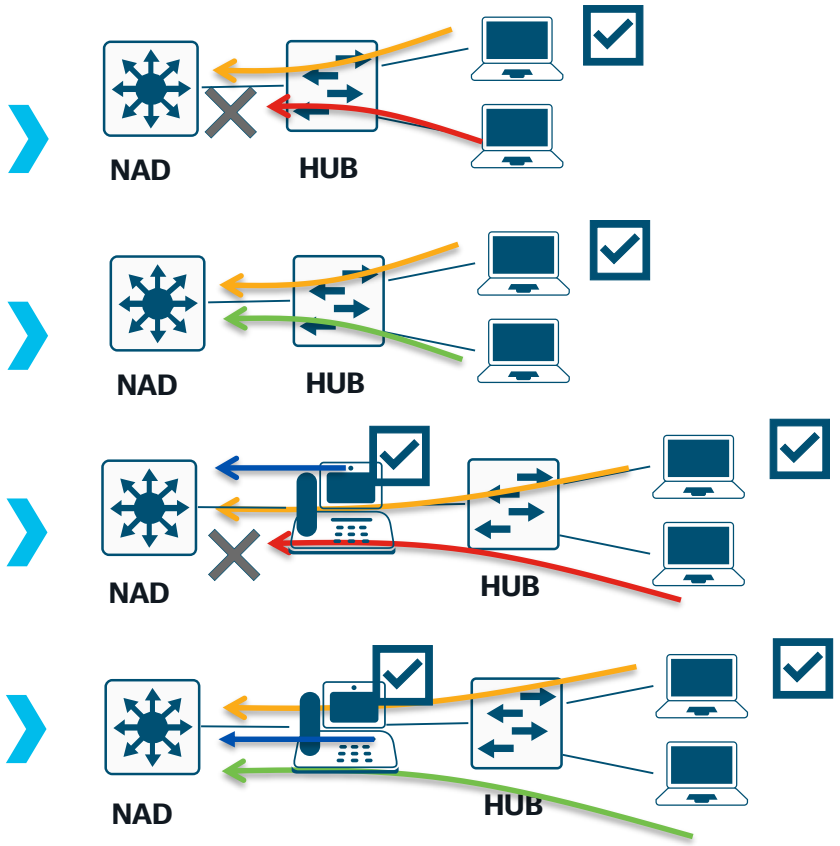

Understanding AAA Scale

Authentication, Authorization & Accounting Scale

Remember - 1K Ports & 20:1 Rules

- Each unique Endpoint (Host) will have 1+ AAA session
  - For IP Phones = 1x **Data** + 1x **Voice** (2x)
- All AAA sessions are learned on Access or Distro (Authenticator)
  - Access: # **Hosts** = # **AAA**

    1-2K AAA
  - Distro: Sum of # Access

    1-2K x 20 = 20-40K AAA

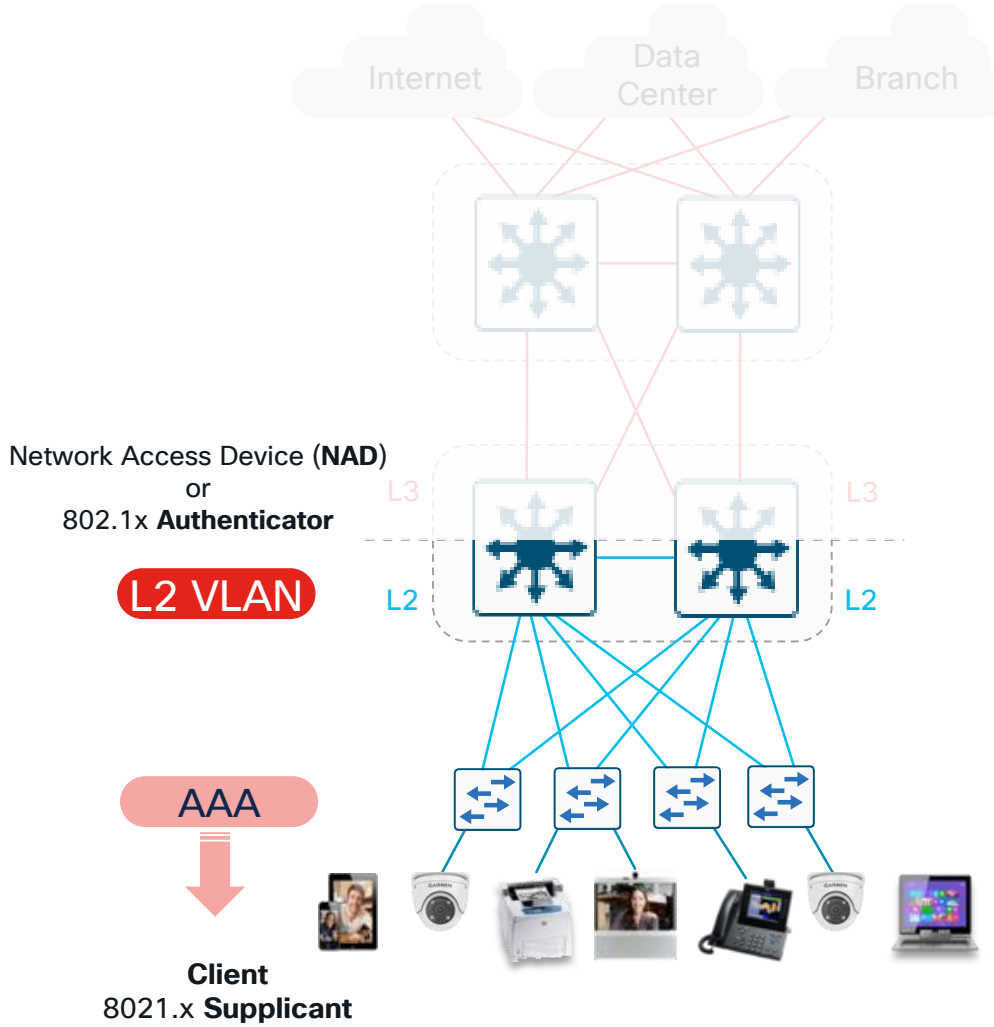

Security ACL Types

Understanding ACLs in Software & Hardware

Live! BRKENS-2062

PACL
Port Based ACL

Applied to Physical Ports


```
access-list extended PACL permit 36.48.0.3
access-list extended PACL deny 36.48.0.0 0.0.255.255
access-list extended PACL permit 36.0.0.0 0.255.255.255
!
interface gigabitethernet 2/0/1
  switchport
  ip access-group PACL in
```

VACL
VLAN Based ACL

Applied to L2 VLAN domains

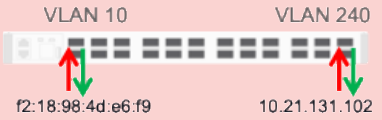


```
ip access-list extended SERVER_ACL
  permit ip 10.1.2.0 0.0.0.255 host 10.1.1.100
  permit ip host 10.1.1.4 host 10.1.1.100
vlan access-map VACL
  match ip address SERVER_ACL
  action drop
!
vlan filter VACL vlan-list 10
```

RACL
Router (IP) ACL

Applied to L3 Ports & SVIs

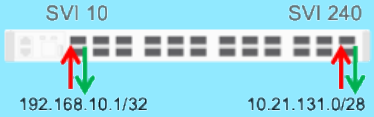


```
access-list extended RACL permit 36.48.0.3
access-list extended RACL deny 36.48.0.0 0.0.255.255
access-list extended RACL permit 36.0.0.0 0.255.255.255
!
interface vlan 36
  ip address 36.64.0.1 255.255.255.0
  ip access-group RACL in
```

OGACL
Object-Group ACL

Maps IP to Object

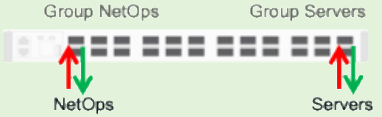


```
object-group network NetOps
 11.96.179.0
 11.96.180.0
object-group network Servers
 10.10.172.12
 10.10.179.210
ip access-list extended OGACL
 permit object-group NetOps Servers any
```

Why OGACL & SGACL in Campus?

Object-Group ACLs for IP

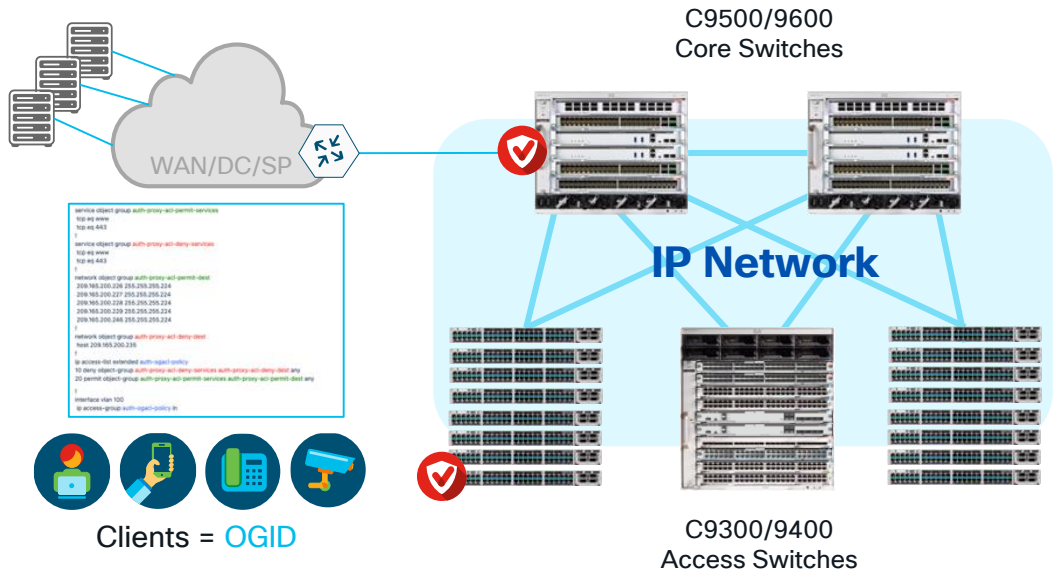

Object-Groups map IP/mask to Labels in CEM

- User defines IP/masks to simple OG name
- OGID labels are stored in Exact Match table

OGACL ACEs take minimal space in ACL TCAM

- Only the Permit/Deny ACEs in TCAM
- OGACLs with same ACEs can reuse entries

Source-Groups ACLs for SDA

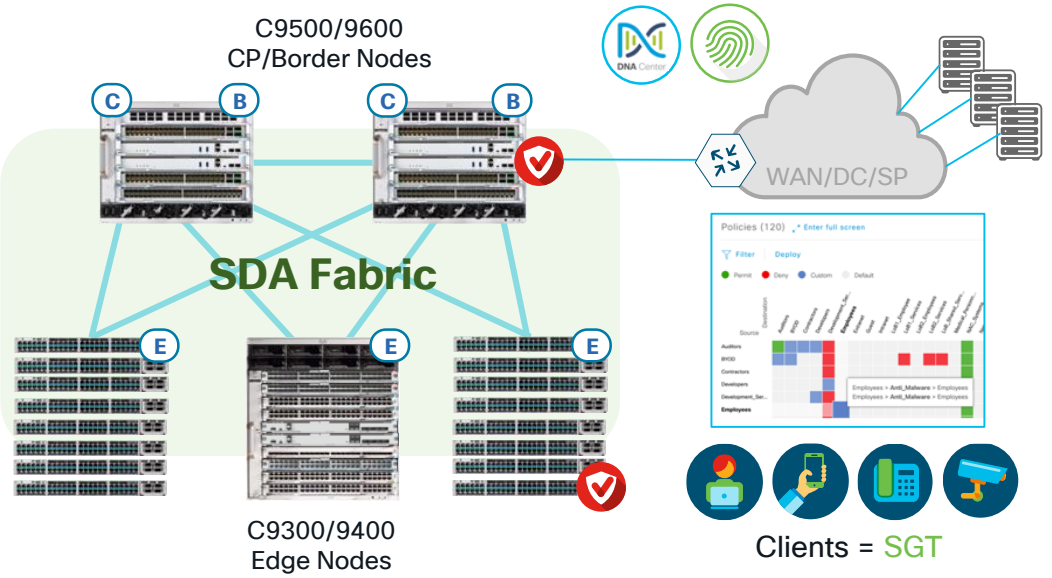

Source-Groups map IP/mask to Labels in CEM

- ISE/DNAC defines IP/masks to simple SG name
- SGT labels are stored in Exact Match table

SGACL ACEs take minimal space in ACL TCAM

- Only the Permit/Deny ACEs in TCAM
- SGACLs with same ACEs can reuse entries

Cisco Flexible NetFlow (FNF)

Collecting & Exporting Full or Sampled Data Analytics

Live! BRKOPS-2038

```
flow record SAMPLE-RECORD
 match ipv4 source address
 match ipv4 destination address
 match transport source-port
 match transport destination-port
 match flow direction
 collect counter bytes
 collect counter packets
 collect timestamp sys-uptime first
 collect timestamp sys-uptime last
```

Flow Record(s)

```
flow exporter SAMPLE-EXPORT
 description FNF v9 Exporter
 destination 10.1.30.230
 source Loopback0
 transport udp 999
```

Flow Exporter(s)

```
flow monitor SAMPLE-MONITOR
 description FNF v9 Monitor
 cache timeout active 15
 record SAMPLE-RECORD
 exporter SAMPLE-EXPORT
```

Flow Monitor

```
interface GigabitEthernet1/0/36
 ip flow monitor SAMPLE-MONITOR input
 ip flow monitor SAMPLE-MONITOR output
```

Apply to Interface(s)

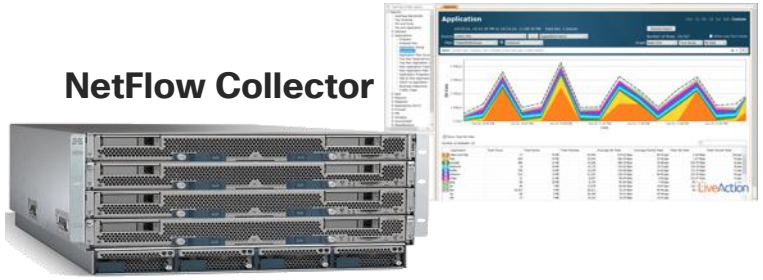

Flexible NetFlow
data exporting

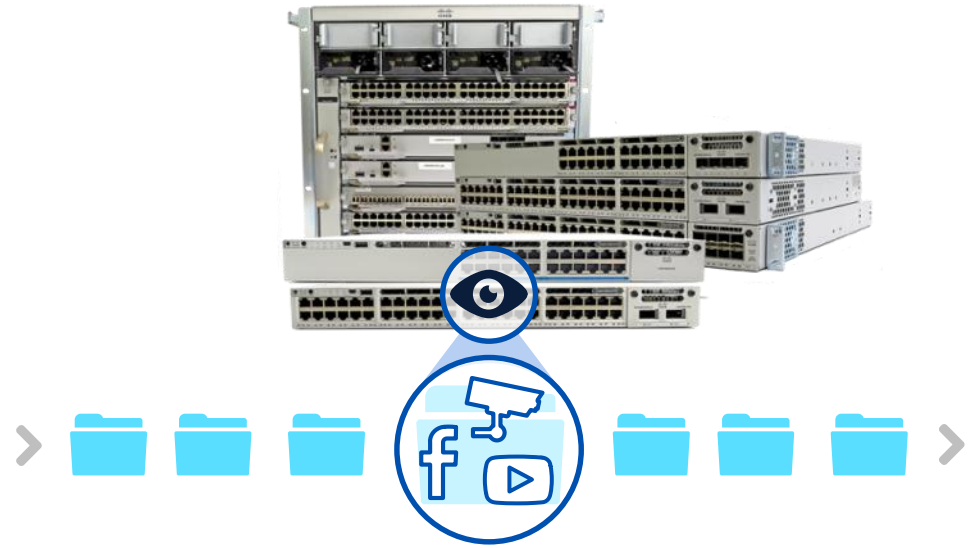

Distributed FNF design for Campus

ID @ Access + Monitor @ Core

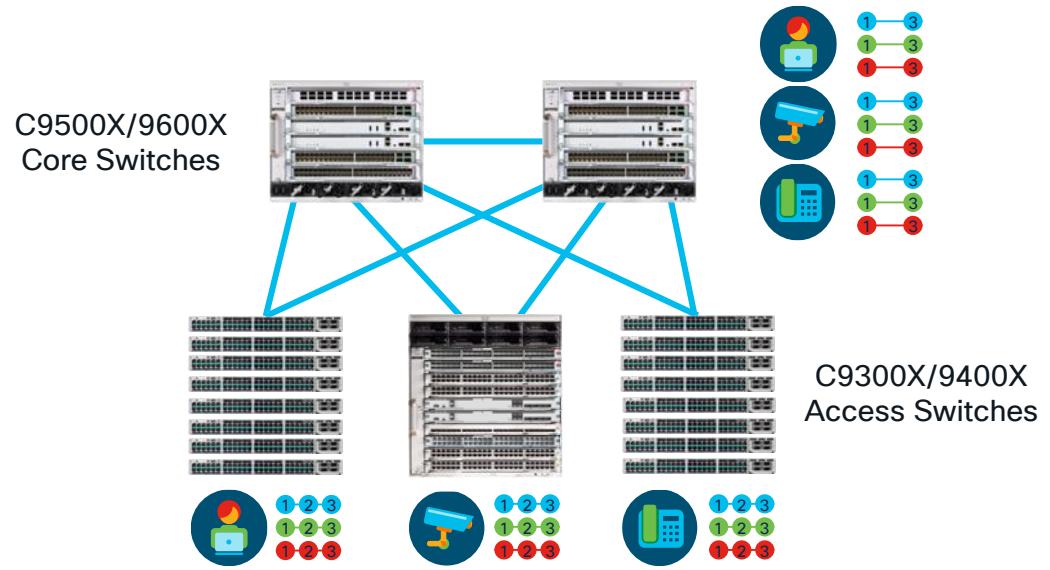

Detailed (1:1) flow identification at the Access

- Better to ID flows as they enter the network
- Full accounting of every client/flow (ETA/AVC)

Aggregate (1:1K) monitoring of flows at the Core

- Just need to monitor the overall network usage
- Adjust sample rates to balance scale & load

Campus-wide FNF Scale

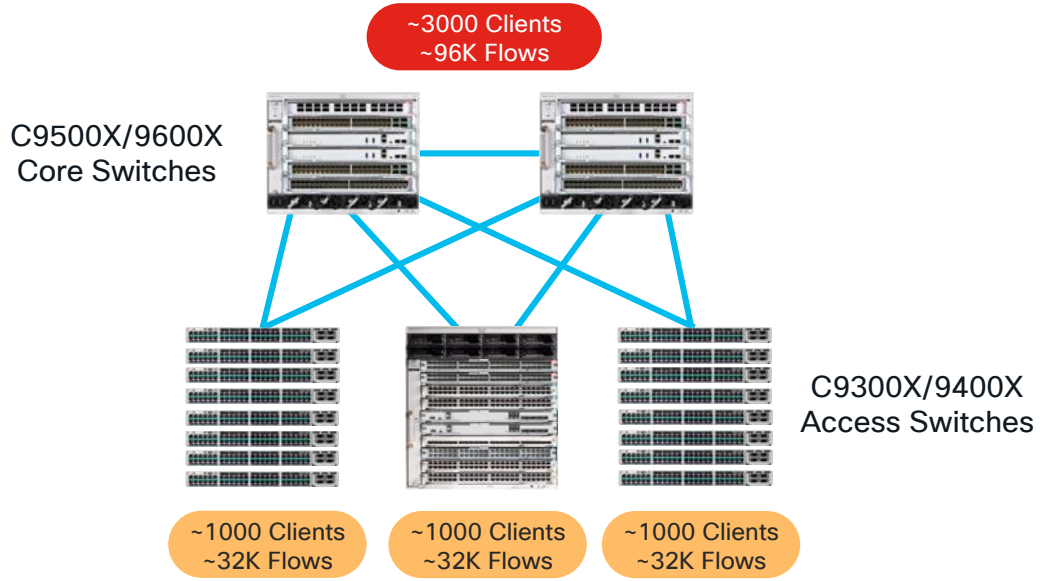

Low-Moderate FNF scale at the Access

- Fewer number of connected clients/flows
- Average ~1K clients x ~32K flows per Access

Medium-High FNF scale at the Core

- Need to aggregate all clients/flows (# Access x 32K)
- Adjust cache aging to increase overall scale

Enhanced Threat Detection & Response (XDR)

Live! BRKOPS-1263

Cisco XDR & NDR - enabled by Flexible NetFlow + Secure Network Analytics

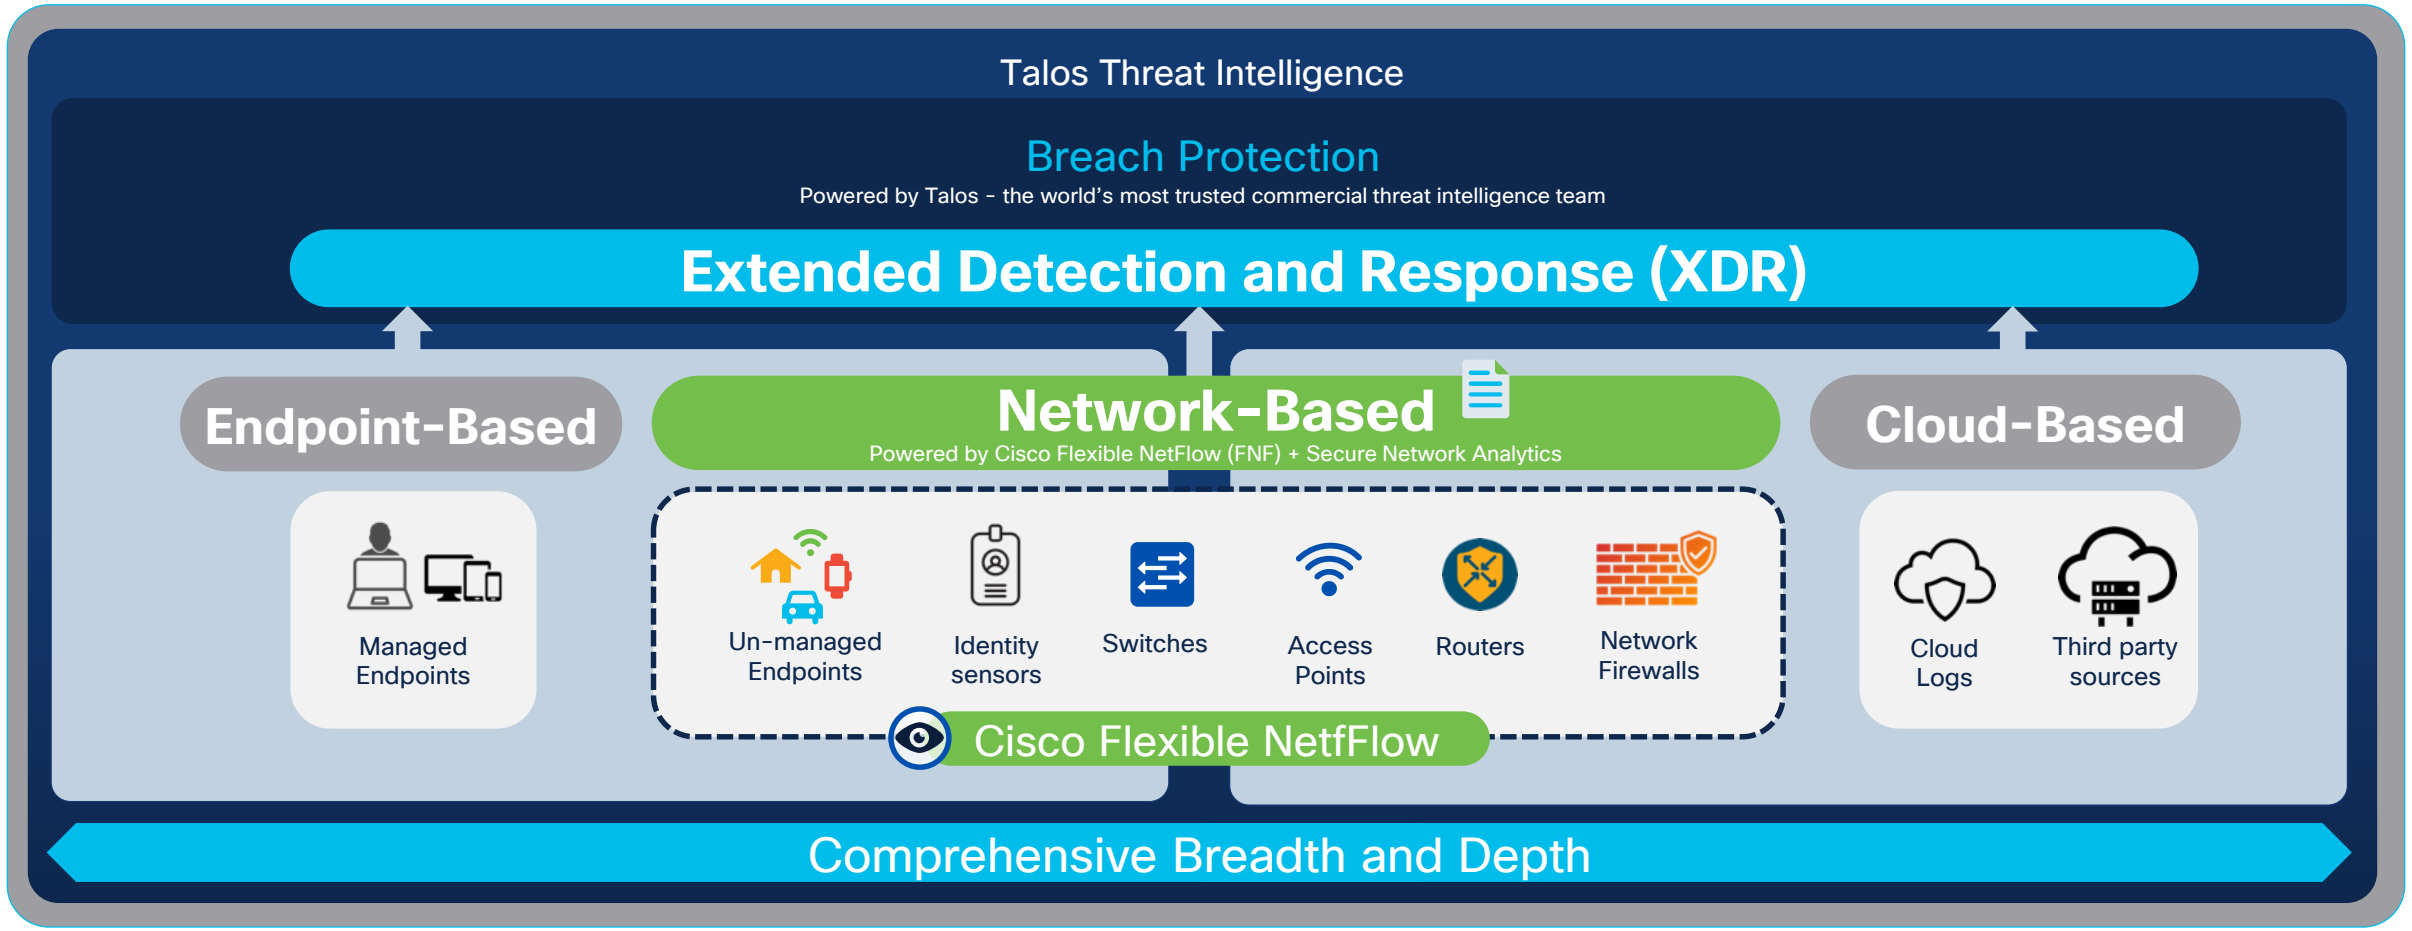

⚠ Network-based detection is critical for unmanaged endpoint threat response!

LAN QOS – Transmit Queue Congestion

The [main] reason for Campus LAN QoS

- The primary role of QoS in Campus networks is to manage packet loss
- Rich media applications (audio/video) are extremely sensitive to packet drops
- In Campus networks – it takes only a few *milliseconds* of congestion to cause drops

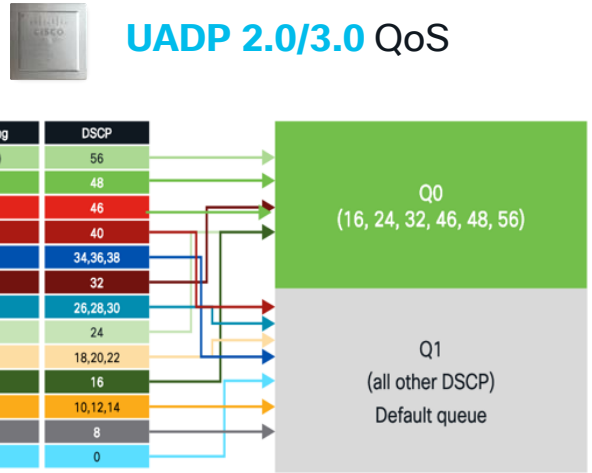

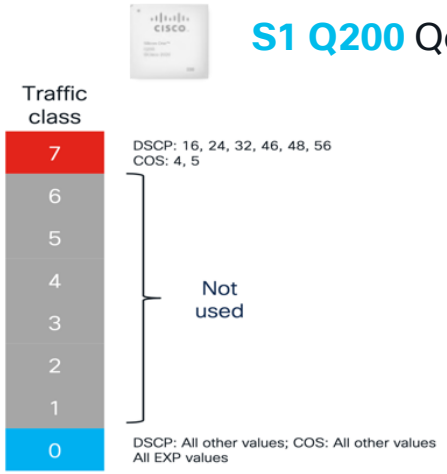

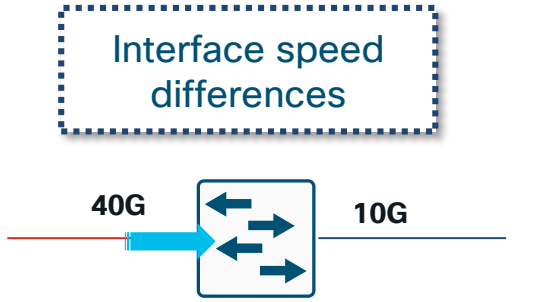

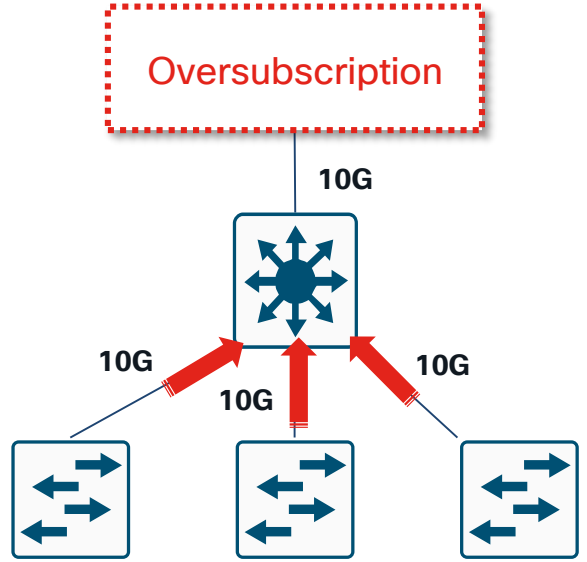

LAN QoS End-to-End

- Prepare your QoS Strategy – what are the Critical/Business Relevant/Default applications?
- Understand QoS capabilities of chosen platforms
- Match the strategy against the platform capabilities
- Always build bi-directional and End-to-End policy

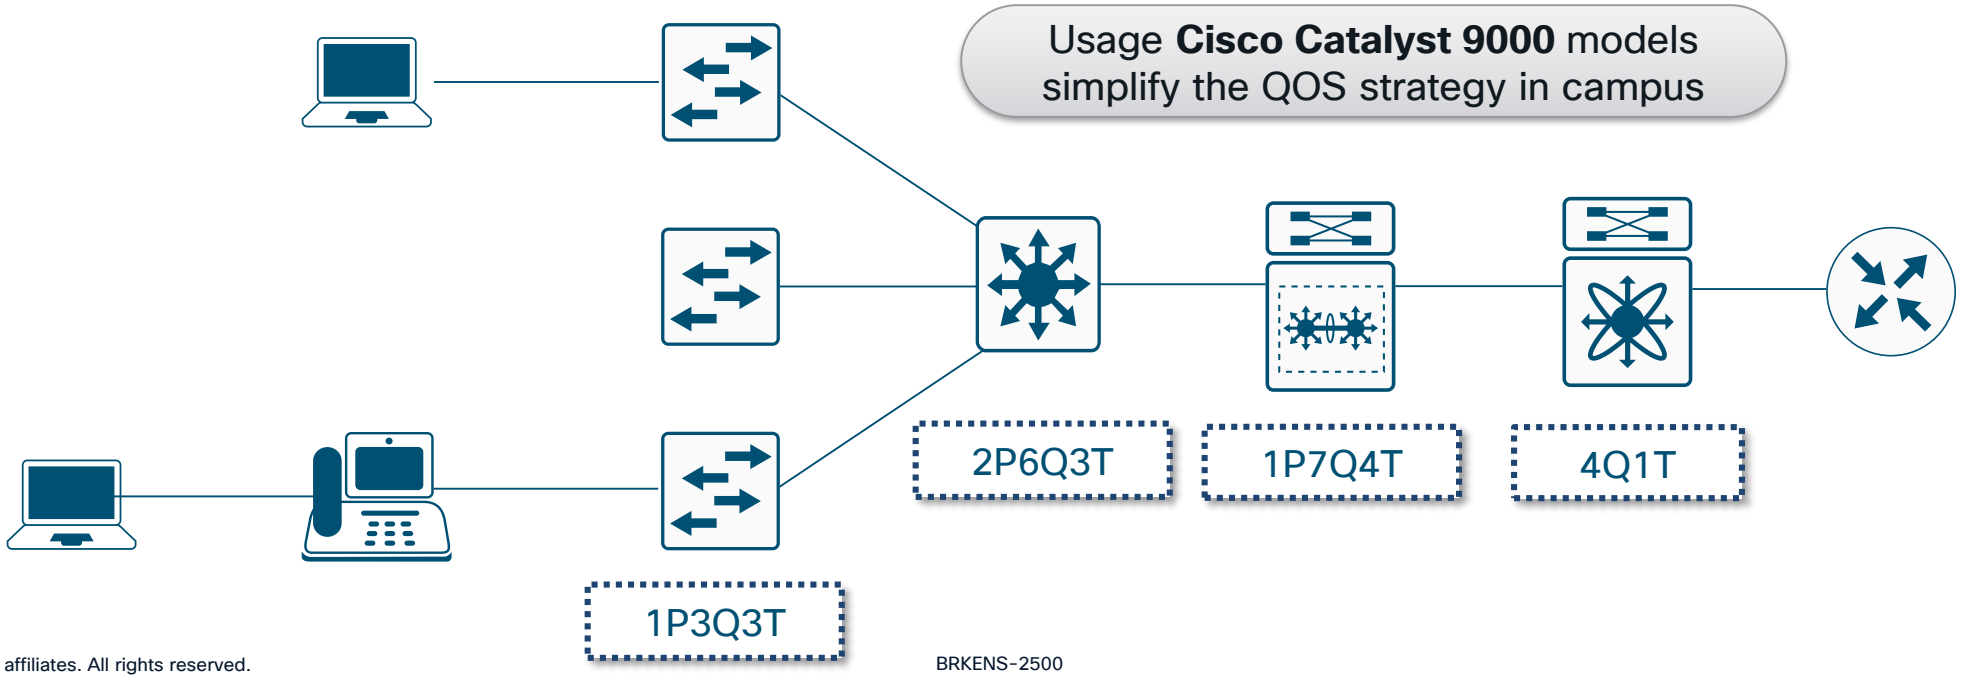

Campus LAN Port QoS Roles

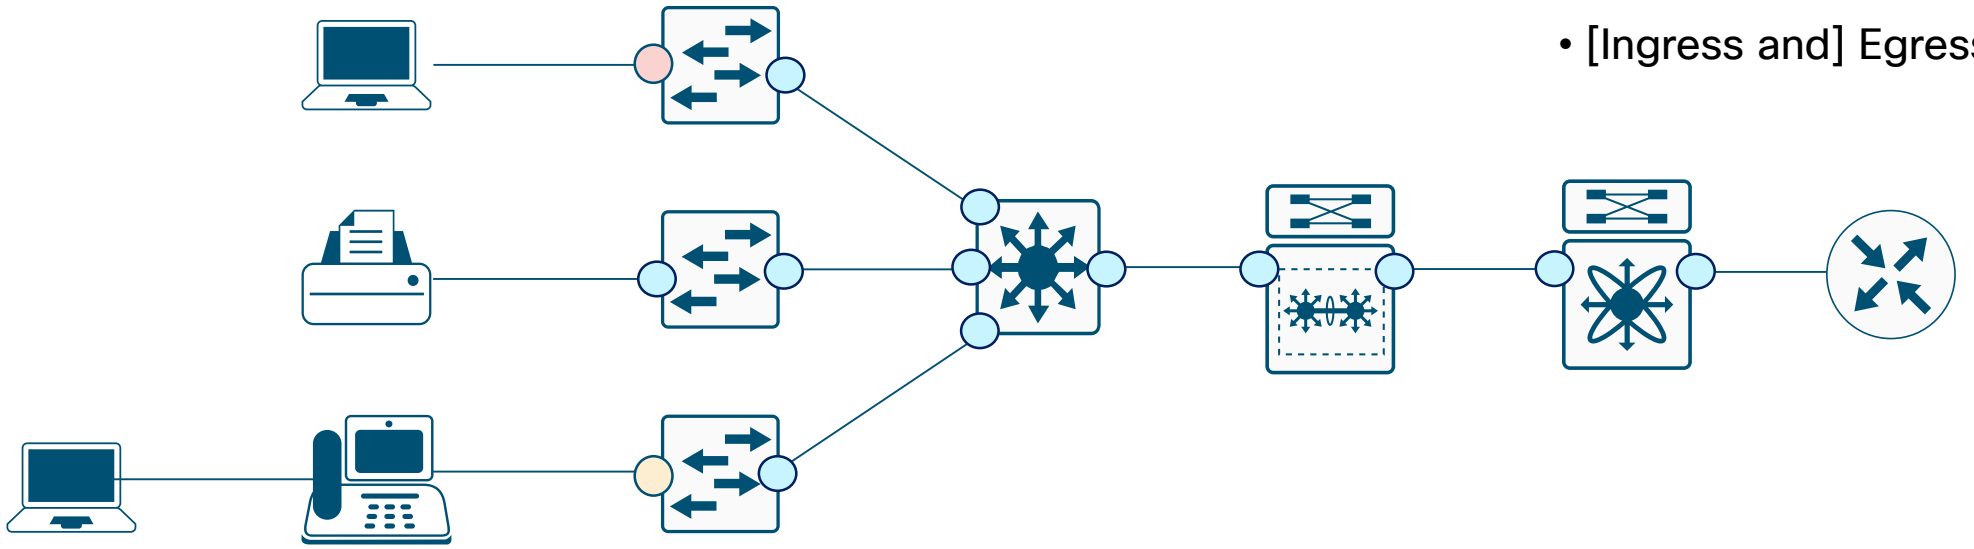

Untrusted Endpoint:

- Port Set to Untrusted State
  (or Explicit Policy to Mark to DSCP 0)
- [Optional Ingress Marking and/or Policing]
- [Ingress and] Egress Queuing

Conditionally-Trusted Endpoint

- Conditional-Trust with Trust-CoS or DSCP
- [Optional Ingress Marking and/or Policing]
- [Ingress and] Egress Queuing

Trusted Port

- Trust DSCP
- [Ingress and] Egress Queuing

Catalyst 9000 Switching QoS

BRKENS-2096

Cisco Catalyst 9000 Switching QoS Deep Dive

Ninad Diwakar - Technical Marketing, Cisco

This session will deep dive into the QoS model used in the Cisco Catalyst 9000 Series of switches powered by the Cisco UADP and Cisco Silicon One Q200 ASICs.

The session will cover platform-specific designs for classification, policing, and ingress and egress queueing policies which are applicable to the Catalyst 9200, 9300, 9400, 9500 and the 9600 switches. To close things off, the session will cover thought processes to be followed for migration configurations from Catalyst 6500 Series switches over to the Catalyst 9500/9600 Series switches.

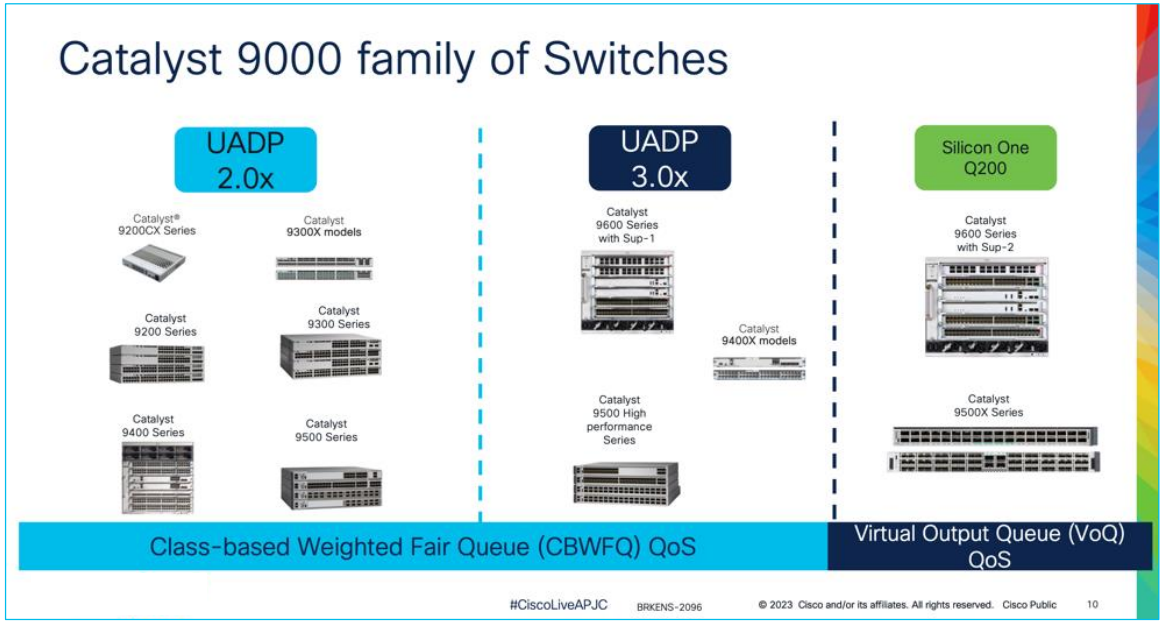

Cisco IOS XE – App Hosting

Enabled by IOS-XE and CAF

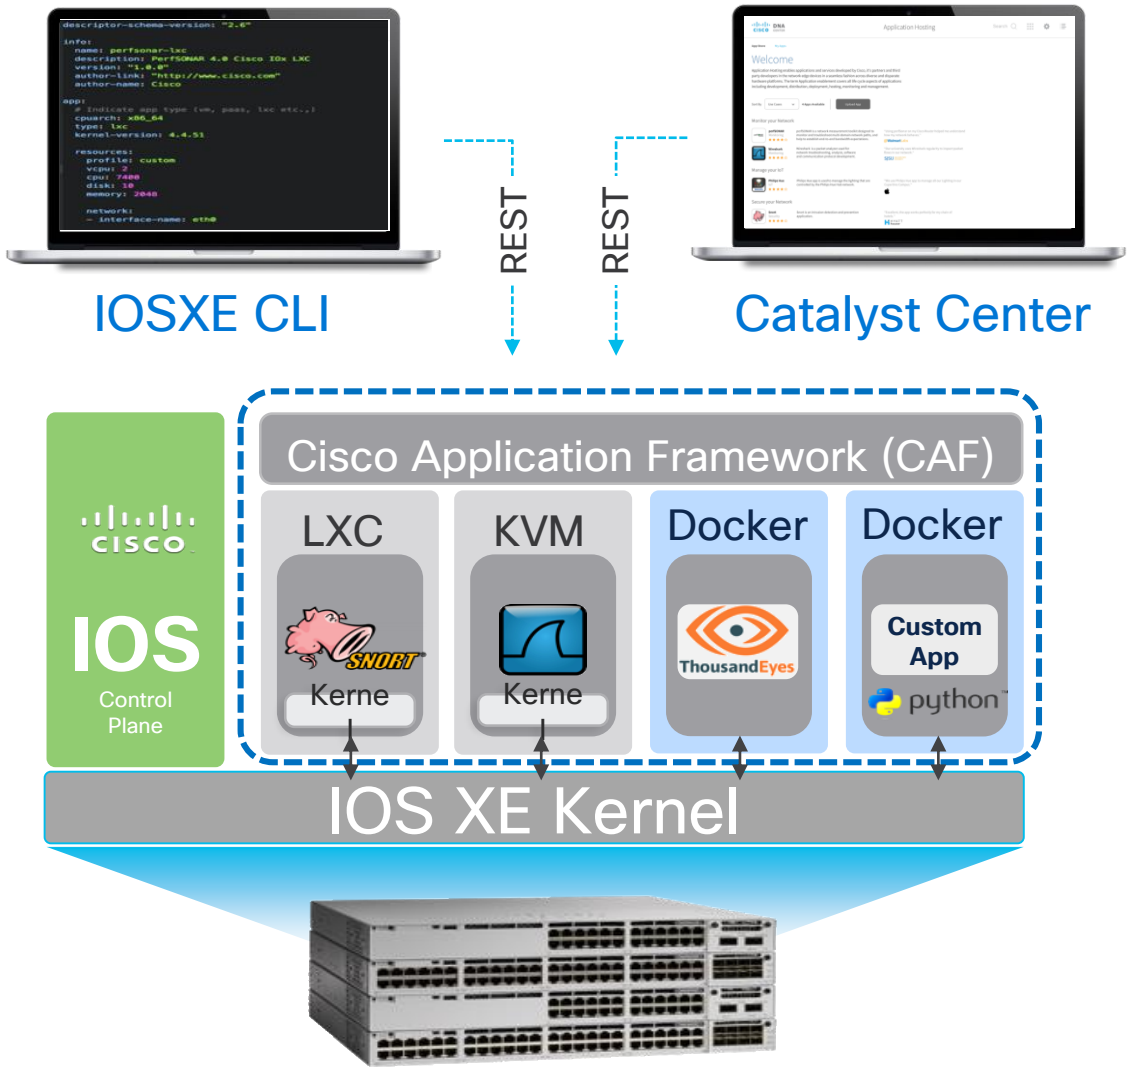

Application Ecosystem

TRex Realistic Traffic Generator · Attivo Networks · Hawkeye

TRAFip · ThousandEyes · NETBEEZ

perfSONAR · Snort · AWS Greengrass

ISC DHCP Server · BIND 9 · iPerf

Nagios · Alleantia · More...

- Cisco will not support third-party apps or open-source apps, unless specifically called out
- Such apps, however, will be validated for compatibility on Catalyst 9000 switches
- DevNet ecosystem will indicate the partners who have worked on Catalyst 9000 switches

Catalyst 9000 & ThousandEyes

Catalyst 9300/L

Catalyst 9400

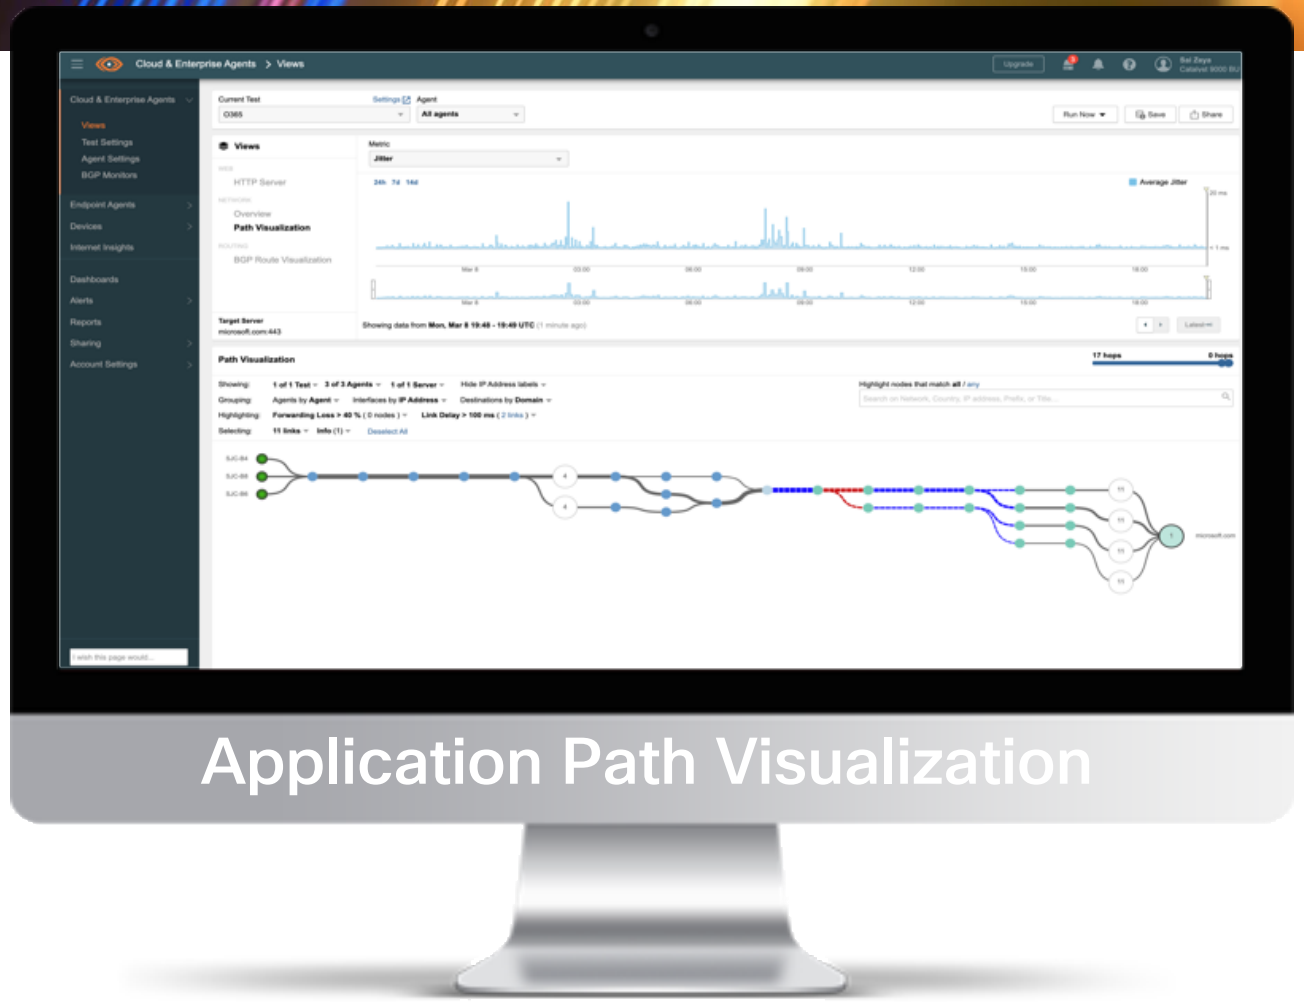

Application Path Visualization

Visualize Applications path to the DC/Cloud and identify Anomalies

Catalyst 9000 Series – App Hosting

Container Networking

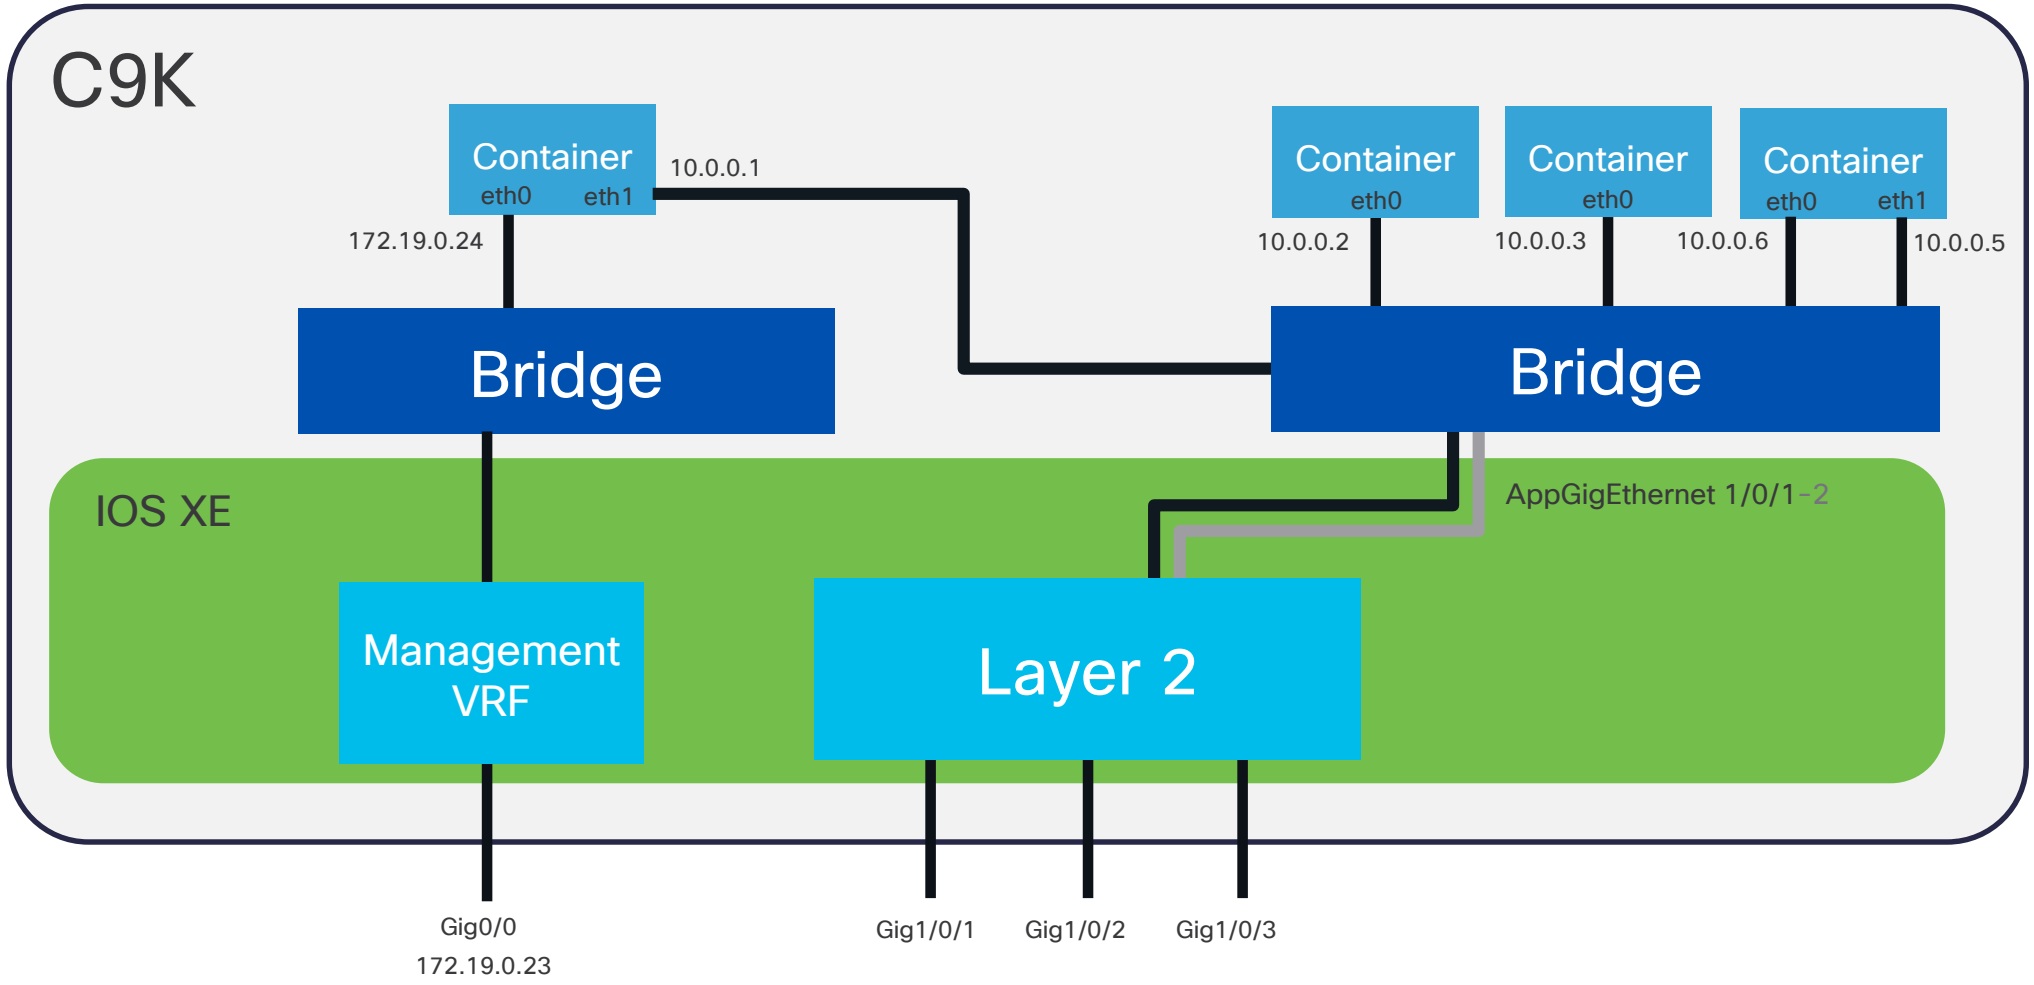

App Hosting - Hardware Resources

Product Selection is critical for App Hosting

* Using loopback with external ports

	Resource type	Catalyst 9300	Catalyst 9300-X	Catalyst 9400	Catalyst 9400-X	Catalyst 9500	Catalyst 9500-X	Catalyst 9600	Catalyst 9600-X
Networking	AppGig Port	1x1G	**2x10G**	1x1G	**2x10G**	Mgmt Port*	**2x10G**	Mgmt Port*	Mgmt Port* (2x10G CPU ports)
Resources	Memory	2GB	**8GB**	8GB	8GB	8GB	8GB	8GB	8GB
	CPU	1 core	**2 core**	1 core	1 core	1 core	1 core	1 core	1 core
	Storage	240GB (USB3.0/SSD)	240GB (USB3.0/SSD)	480-960GB (SATA)	480-960GB (SATA)	480-960GB (SATA)	480-960GB (SATA)	480-960GB (SATA)	480-960GB (SATA)

Catalyst 9300-X

USB 3.0
240GB

Back panel

Catalyst 9400-X

M2 SATA
480/960GB

Plug into removable SUP

Catalyst 9500-X

M2 SATA
480/960GB

Back panel

Catalyst 9600-X

M2 SATA
480/960GB

Plug into removable SUP

Session Agenda

Design Fundamentals | **Design Considerations**

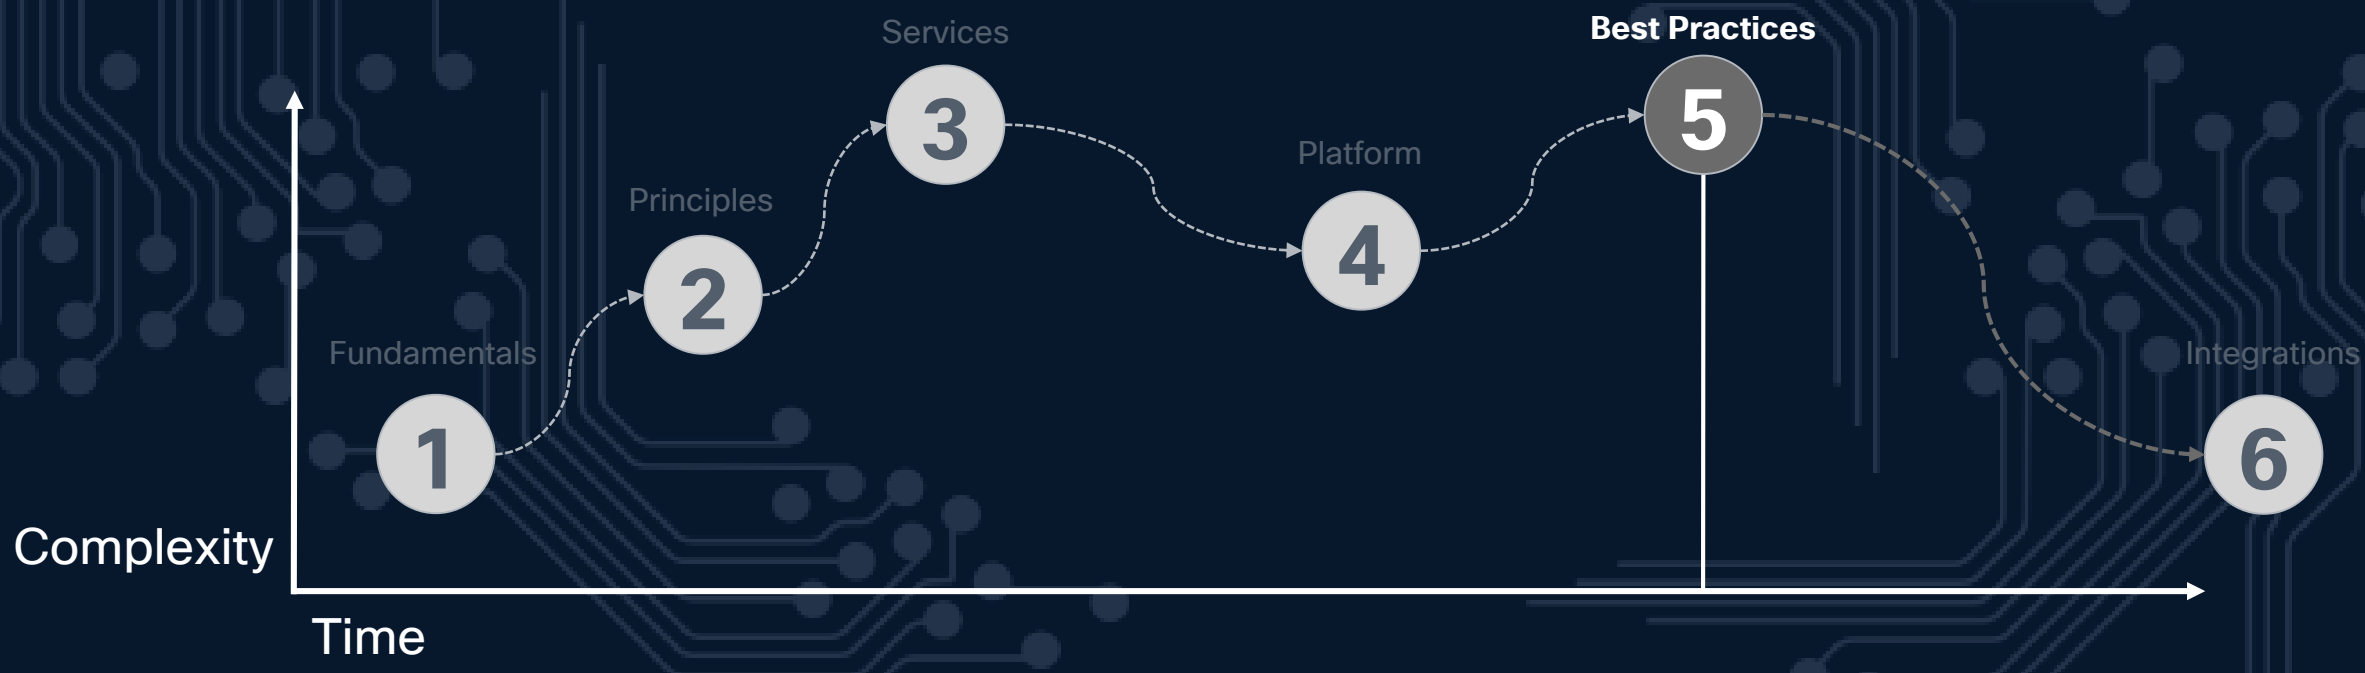

Platform Design

- **LAN High Availability**
  - Basics of LAN HA
  - SSO, Stacking & SVL
  - SMU, ISSU & xFSU
  - ECMP, MEC & mLAG
- **LAN Security**
- **Virtual Networking**

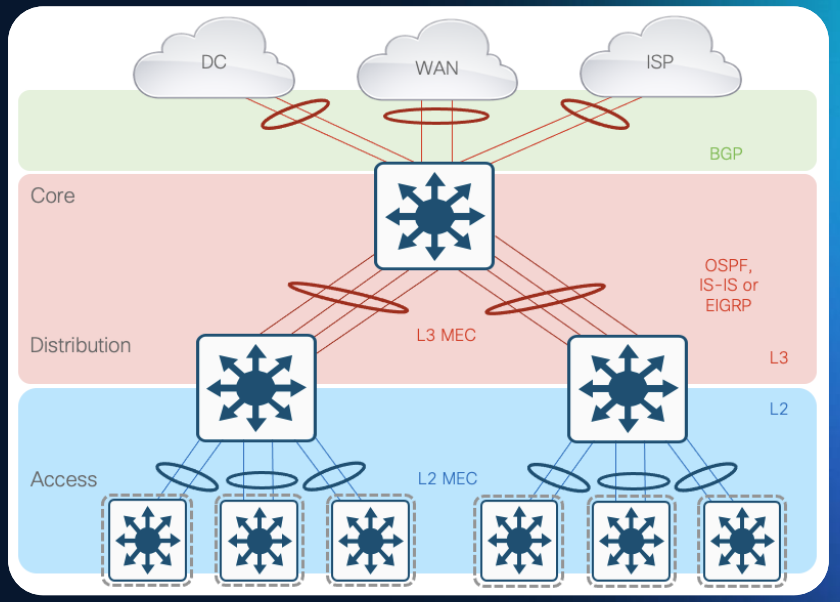

Mission-Critical Resiliency

BRKENS-1500

Your business stops if the network is down

Cost of only **one hour** of downtime to an average enterprise **> $300,000****

** Based on industry reports from Gartner and ITIC

Catalyst 9600
(Dual chassis w/ StackWise Virtual)

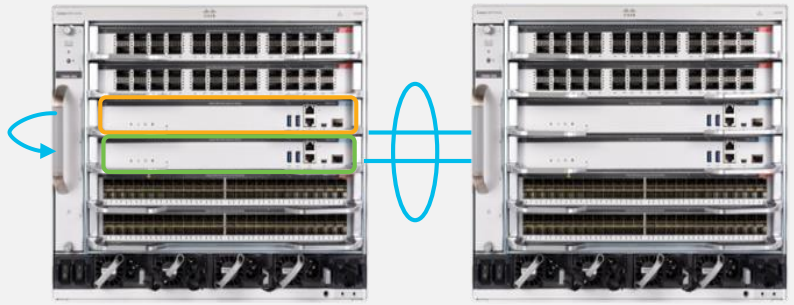

Catalyst 9300, 9400 & 9500

Architecture

StackWise® & StackWise Virtual

- Virtualized redundant systems for simplified configuration & protocols

Graceful Insertion/Removal (GIR)

- No downtime when device in maintenance mode

Operating System

Software Maintenance Upgrade (SMU)

- Minimal or no downtime patches

In-Service Software Upgrade (ISSU)

- Minimal or no traffic loss upgrade

Extended Fast Software Upgrade (xFSU)

- < 5 sec downtime - Stack upgrade

Platform

Redundant Control & Data-Plane

- Dual Sup or Stack SSO/NSF
- SVL with Quad-SUP RPR

Redundant Power & Fans

- N+1 or Combined mode

StackPower for StackWise

Eliminate downtime with **High Availability** designed at every level

High-Availability - SSO & NSF

BRKENS-1500

Stateful Switchover (SSO) and Non-Stop Forwarding (NSF)

- **Stateful Switchover** (SSO) synchronizes active process state and running-config, between Active & Standby supervisors or Active & Standby switches in a stack
  - Traffic loss minimized for Active supervisor or Active switch failure
- **Non-Stop Forwarding** (NSF) allows for graceful restart of L3 routing protocols

Modular Switch with **Redundant Supervisors**

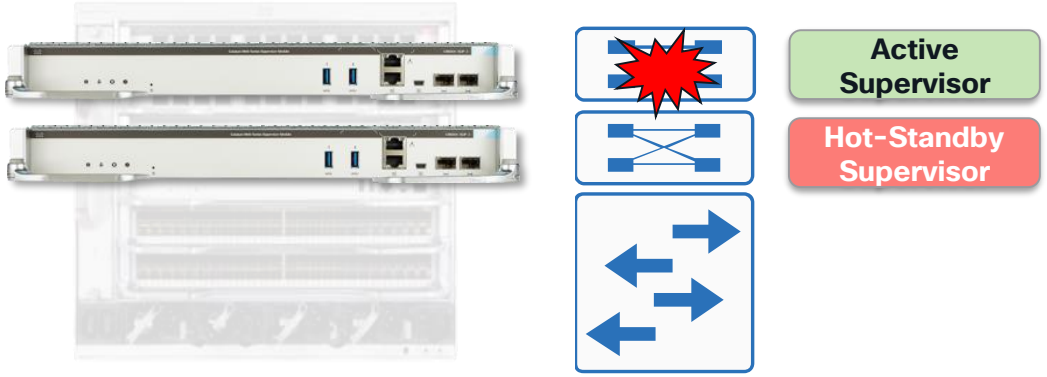

StackWise Stack or **StackWise Virtual Pair**

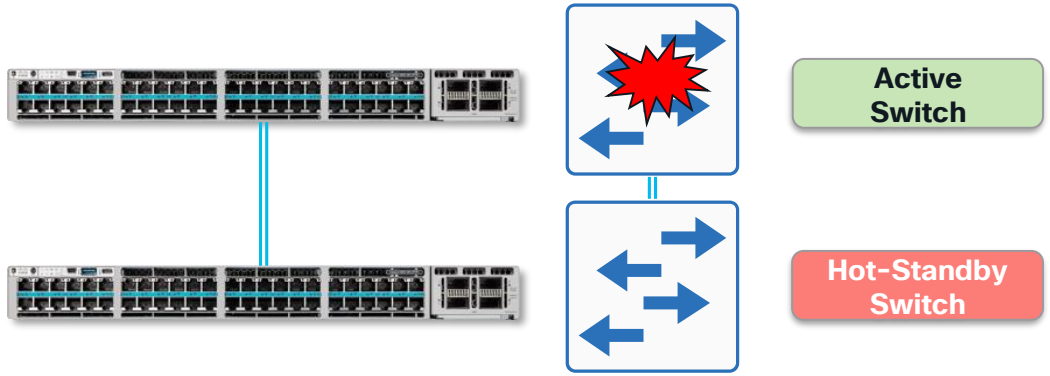

Cisco IOS XE High Availability

Control-Plane to Data-Plane Programming

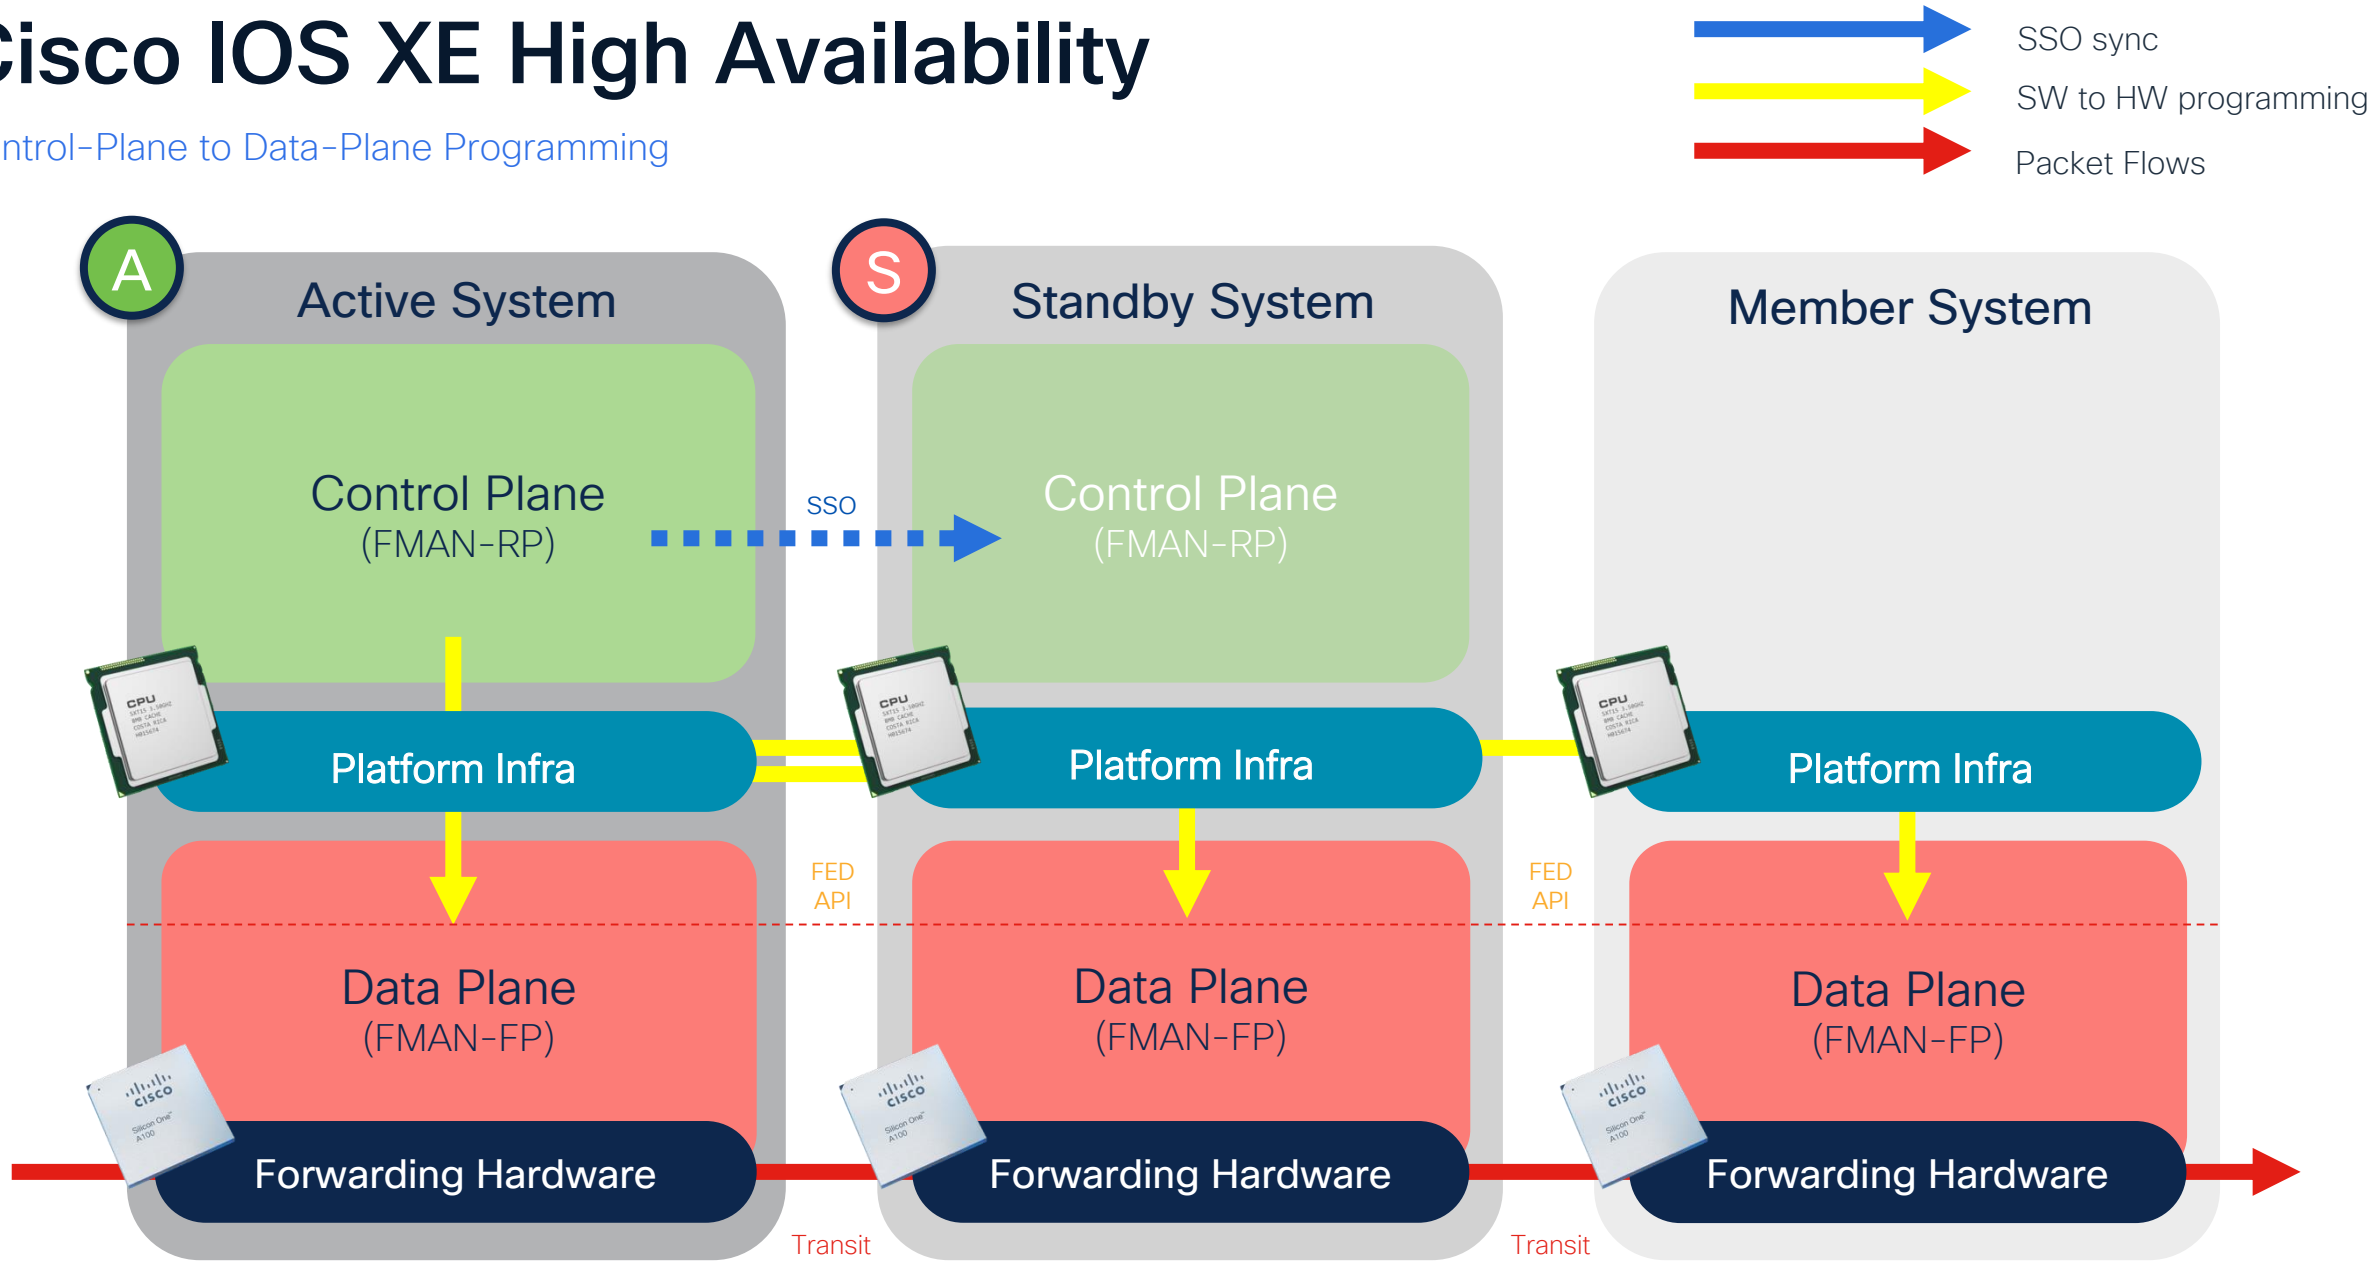

Modular High Availability

Maximum Network Up-Time for Physical, Control-Plane & Data-Plane

Catalyst 9600/X

Catalyst 9400/X

SSO Active responsible for:
- Management
- L2 protocols
- L3 protocols

Highest Resiliency

Redundant Supervisors

StackWise® Virtual

Easy Upgrades with ISSU & GIR

Redundant Fans (Fan-Tray)

Redundant PSUs (1:1, N+1)

Best for **Single Attached** Clients

Redundant Supervisors

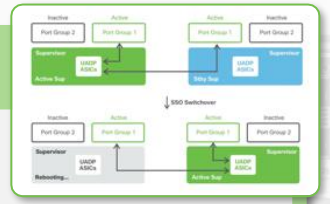

Active + Standby CPU & ASIC

1. **SSO**: Stateful Synch Infra & L2
   - Protects hardware Data-Plane
2. **NSF**: L3 Graceful Restart
   - Protects software Control-Plane
3. **LCs** connected to both Sups
   - ≤ 200ms traffic switchover

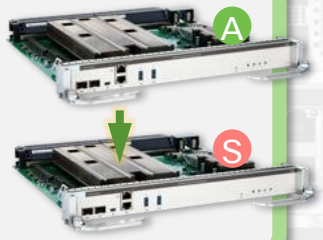

StackWise Virtual

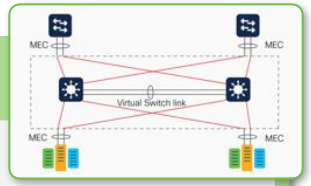

Active + Standby Chassis

- **Same SSO/NSF** as Redundant Sups
  - Stateful Synch vs. ECMP
- **Multi-chassis EtherChannel** (MEC)
  - Best Load-Balance & Convergence
- **Simplify Topology** and Less Peers
  - Eliminate HSRP/VRRP and ECMP

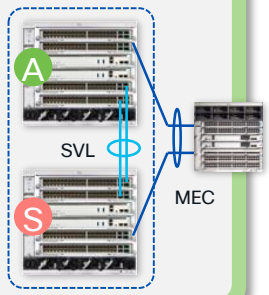

Redundant PSUs & Fans

Backup Power & Cooling

- PSUs can be **1:1 Redundant**
  - 100% backup (A+B Feeds)
- PSUs can be **N+1 Combined**
  - Shared Power + Backup
- Tray works with **2 faulty Fans**
  - Remaining increase RPM

SMU, ISSU & GIR

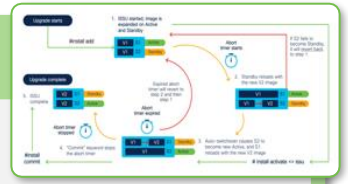

Redundancy during Upgrades

- **SMU** patching for Sev1/PSIRTs
  - Target patch, w/o recertification
- **ISSU** between/within EMR releases
  - SSO upgrades for Sups & SVL
- **GIR** for Traditional (ECMP) Network
  - "Maintenance" mode + ECMP

```
Switch# install
```

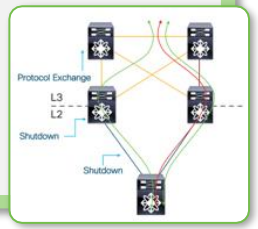

Fixed Access High Availability

Cisco StackWise® – Access Switch Stacking

Catalyst 9300X/LM

Catalyst 9300/L

Catalyst 9200/L

SSO Active responsible for:
- Management
- L2 protocols
- L3 protocols

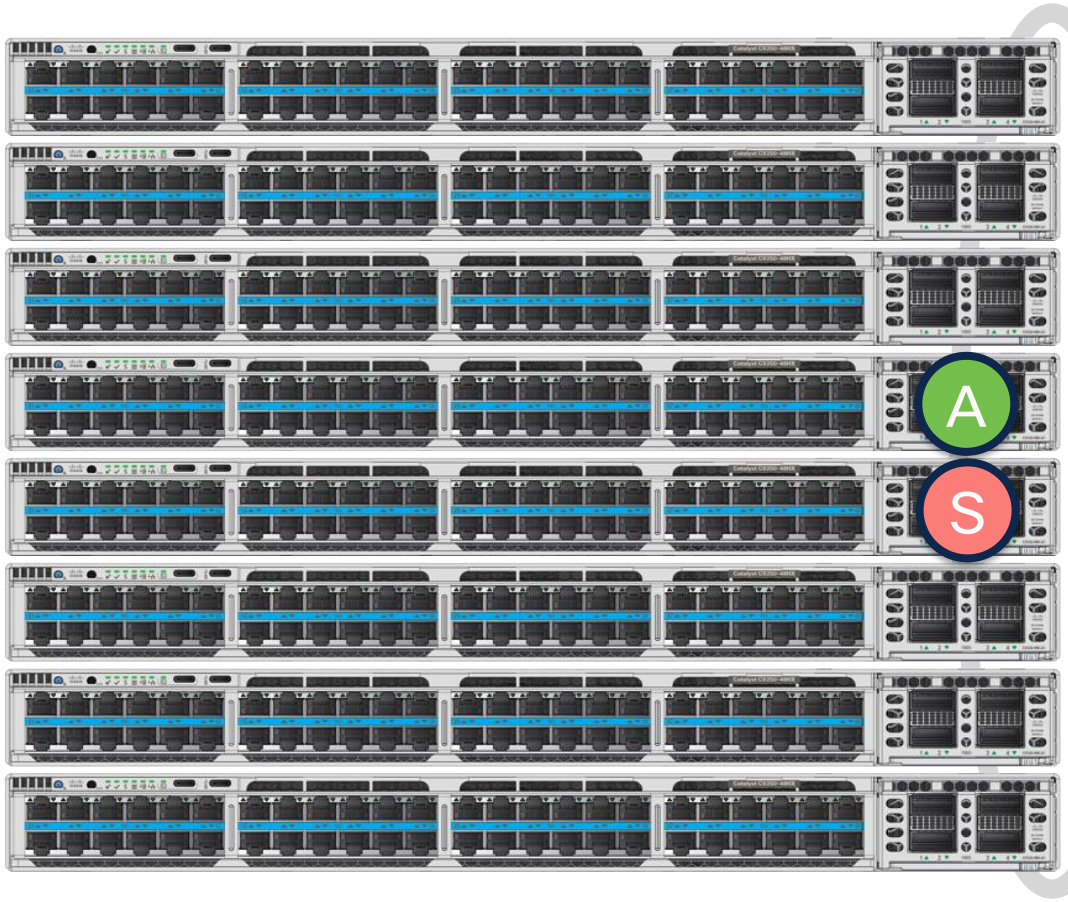

Up to 8 Members

StackWise – 80/160/360/480/1T*

BRKENS-1500

Centralized Control Plane

Distributed Data Plane

1+1 Stateful Redundancy with Active & Standby

Stateful Switchover SSO/NSF

*StackWise speeds vary depending on platform

LAN High Availability

Switch Stacking - Catalyst 9300

Modular Uplink
Catalyst 9300X models (10/25G Fiber)

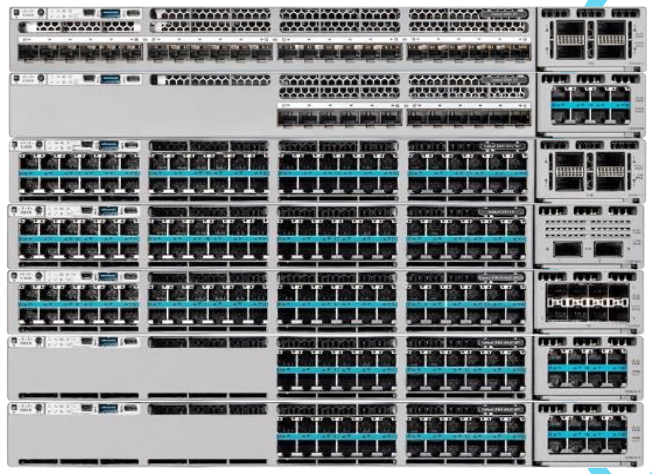

8 switches

1T

Modular Uplink
Catalyst 9300 (non –B) models
or mix with Catalyst 9300X

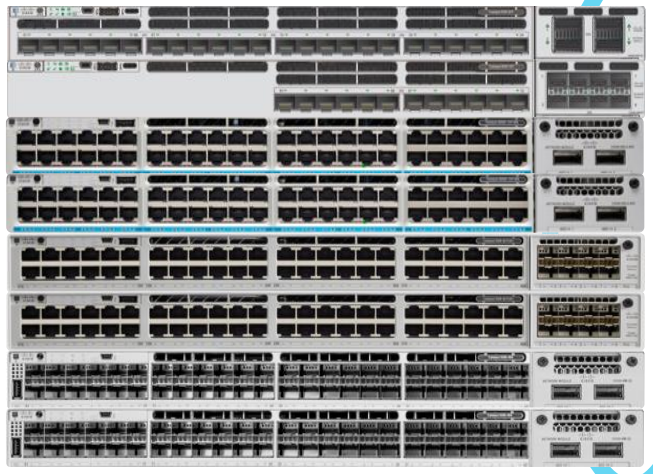

8 switches

480G

Fixed Uplink
Catalyst 9300L models

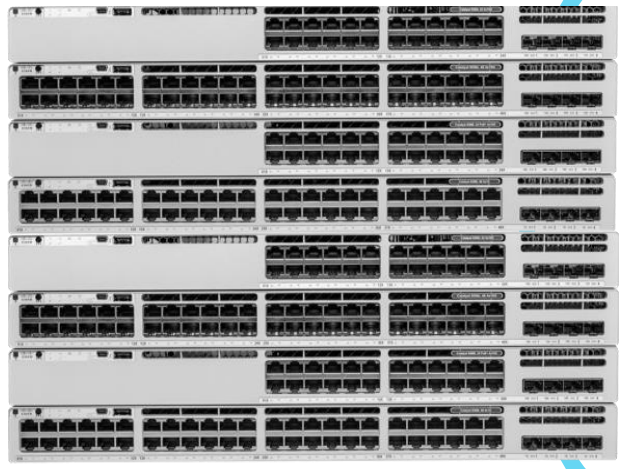

8 switches

320G

Stacking supported among Catalyst 9300X models
and mixed stacking between Catalyst 9300 and Catalyst 9300X models

Stacking supported among
Catalyst 9300L models only

9200 stacks with 9200 and 9200L stacks with 9200L

LAN High Availability

StackPower - Power HA - How it works?

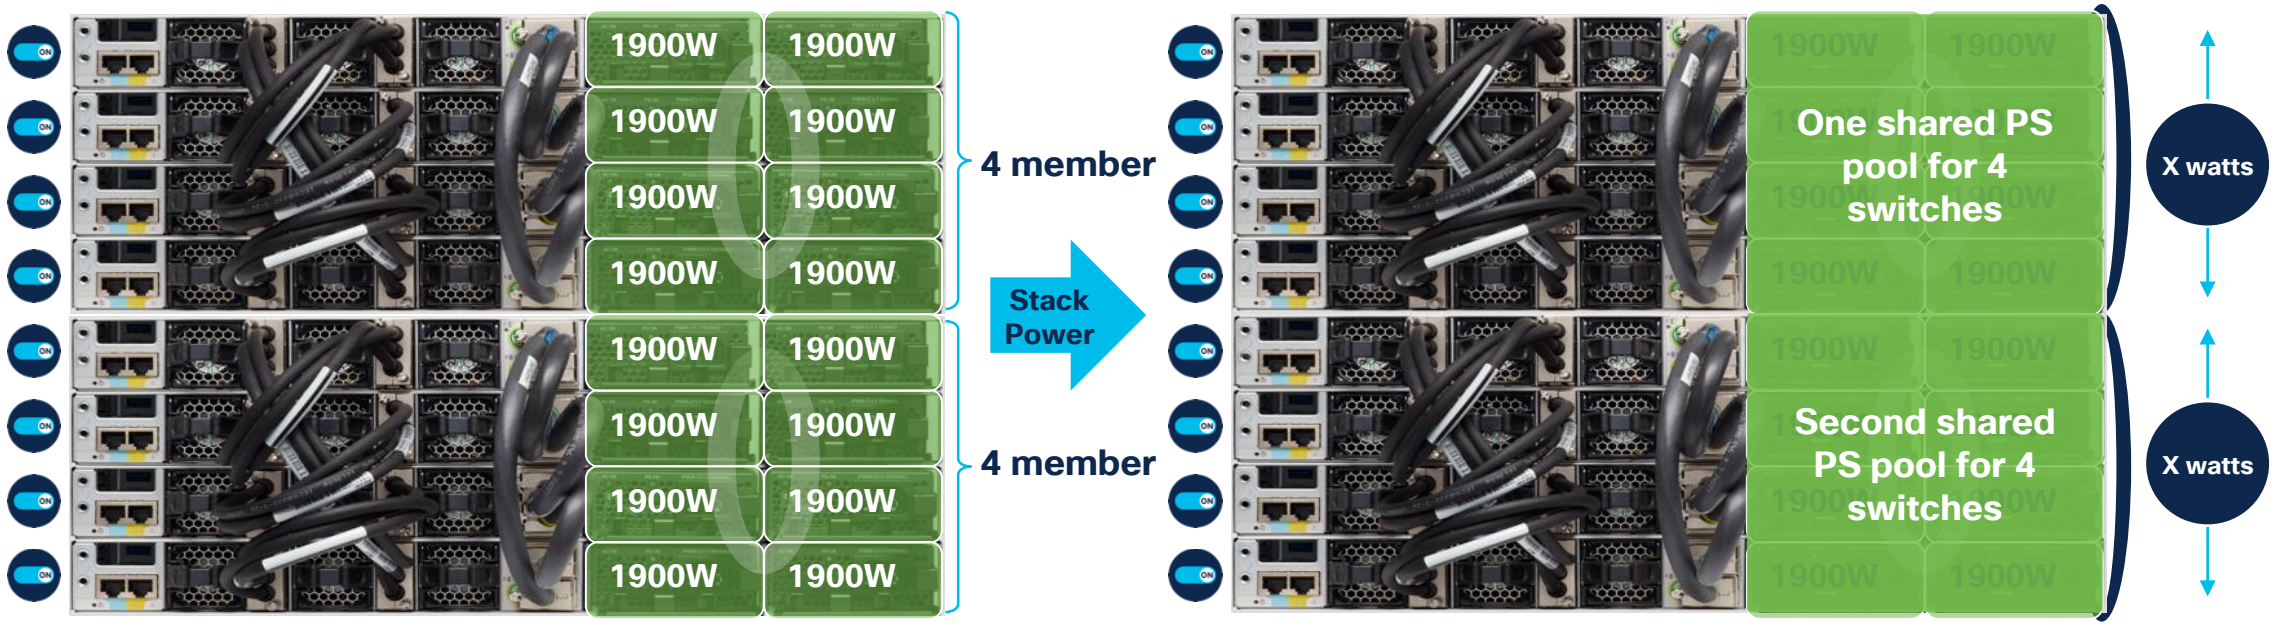

- Pools power from all Power Supplies (PS)
- All switches in StackPower share the available power in the pool
- Each switch is given its minimum power budget

- 1+N Redundancy with inline power
- Up to 4 switches in one StackPower Ring
- Multiple Power stacks possible in one data stack

LAN High Availability

In-Service Software Upgrade (ISSU)


Catalyst 9600/X


Catalyst 9500/X


Catalyst 9400/X

BRKENS-1500

1. ISSU started - image is expanded on active and standby supervisors

Upgrade start

 # install add

V1 S1 Active

V1 S2 Standby

Abort timer starts

2. Standby reloads with the new V2 image

V1 S1 Active

V1 → V2 S2 Standby

If S2 fails to become the active, it will revert back to Step 1

3. Auto-switchover causes S2 to become the new active and S1 reloads with the new V2 image

 # install activate <> issu

V1 → V2 S1 Standby

V2 S2 Active

Expired abort timer will revert to Step 2 and then Step 1

Abort timer expired

4. 'Commit' keyword stops the abort timer

 # install commit

Abort timer stopped

V2 S1 Standby

V2 S2 Active

Upgrade complete

5. ISSU complete

LAN High Availability

Extended Fast Software Upgrade (xFSU)

Catalyst 9300X/LM

Catalyst 9300/L

NEW. ≤ 5 Seconds - 17.15.2

C9300/X (≤ 30s) - 17.7.1

BRKENS-1500

Command to trigger xFSU

C9300# ***install add file*** *<image>* ***activate xfsu commit***

Control Plane Upgrade
V1 ⟶ V2

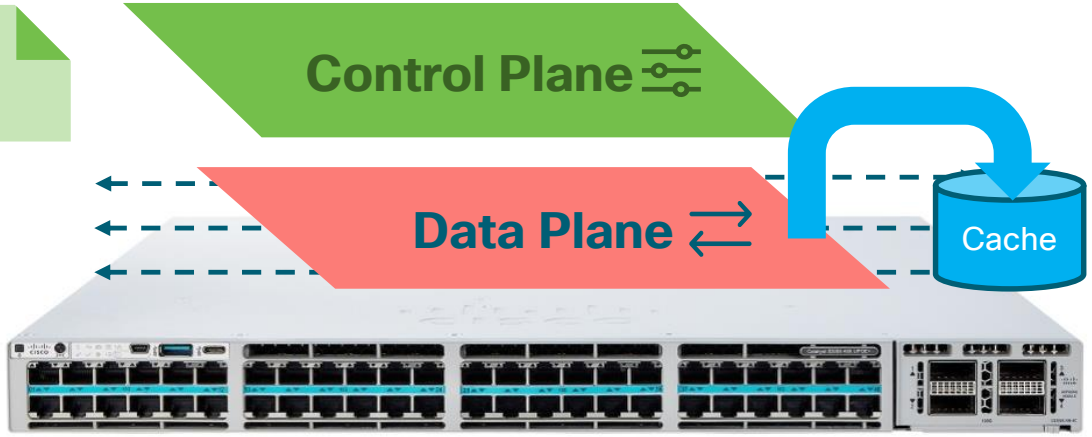

Data Plane upgrade
V1 ⟶ V2

C9300 | C9300L | C9300X

Cisco xFSU minimizes downtime to less than 5 seconds (standalone or stack)

IOS XE Install – SMU patches

Patch a previously 'certified' image, without full upgrade

DNAC Version: Guardian 2.3.3
License Level : DNA Advantage

Patching can also be performed using Cisco Catalyst Center

Software Maintenance Update (SMU) is a small **software package** that can be installed for a **fix** or **security resolution** on a released version.

- **Hot Patching** - does not require explicit reload after installation to get activated. (not traffic-affecting)
- **Cold Patching** - requires a system reload after installation to get activated. (traffic-affecting). The 9200 family supports only this.

NEW

- **Bundle SMU** - One single SMU file containing and applying multiple SMU fixes at once.
- **Independent SMU** - One single SMU file tailored to address a specific bug or vulnerability.

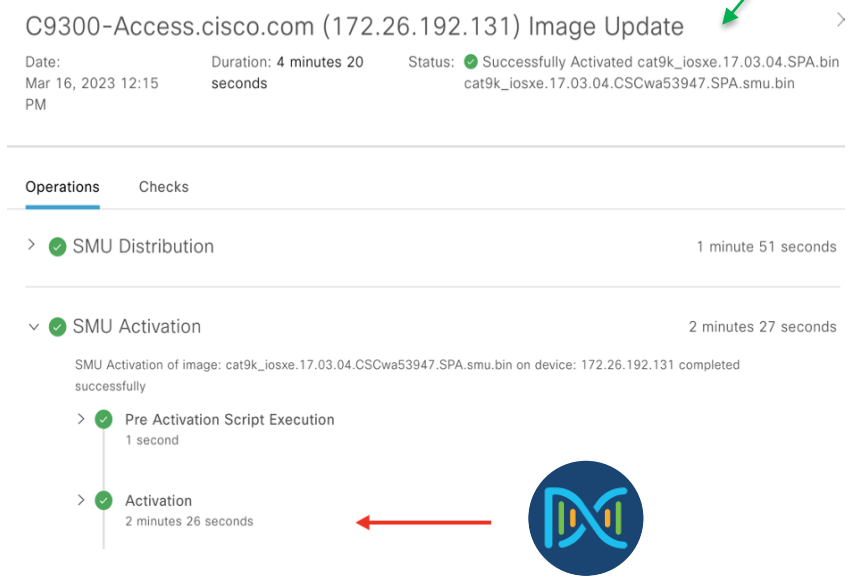

Reduces both time and scope for [re]certification

Self-sufficient and reliant bug fix solution for PSIRT, CFD & IFD

Hitless Zero-Downtime with Hot Patching

EMR Releases with 371 SMU files available

LAN HA - Campus Architecture

Control-Plane & Data-Plane Redundancy

1 ECMP (L2/L3 Paths)

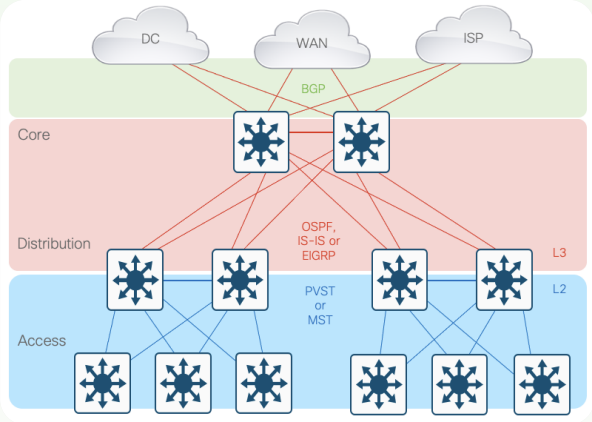

- Complex Topology
- More Nodes, Less Cables
- More Neighbors (+ Tuning)
- Protocol Load-Balancing (ECMP)
- Node-level Redundancy

L1 : Single Connections
L2: STP, MST, REP + ECMP (Port Cost)
L3: FHRP, IGP, BGP + ECMP(Port Cost)
More Neighbors = Requires Protocol Tuning

2 EtherChannel (L2/L3 LAG)

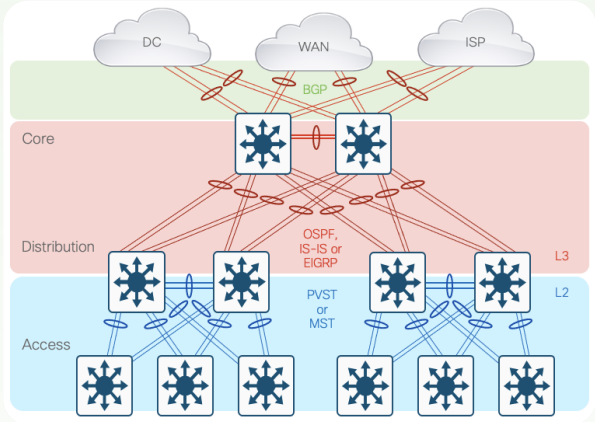

- Complex Topology
- Same Nodes, More Cables (2-8)
- Same Neighbors (+ Tuning)
- EtherChannel Load-Balancing
- Node & Link-level Redundancy

L1 : Multiple Connections
L2: STP, MST, REP + ECMP (Portchannel Cost)
L3: FHRP, IGP, BGP + ECMP (Portchannel Cost)
More Neighbors = Requires Protocol Tuning

3 StackWise (L2/L3 mLAG)

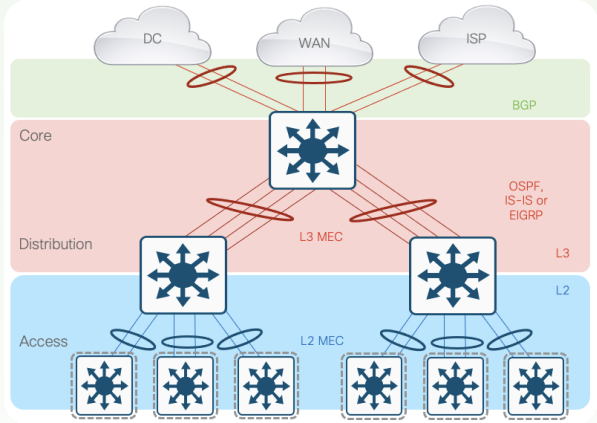

- Simple Topology
- Same Cables, Less Nodes
- Less Neighbors (No Tuning)
- Multi-chassis EtherChannel (MEC)
- Layer-level Redundancy

L1 : Multiple Connections
L2: L2 MEC (No STP or REP)
L3: IGP, BGP + L3 MEC (No FHRP)
Fewer Neighbors & No Protocol Tuning

Design Options

Traditional Multi-Layer Campus

Logical topology– L3: core/dist. L2: dist./acc.

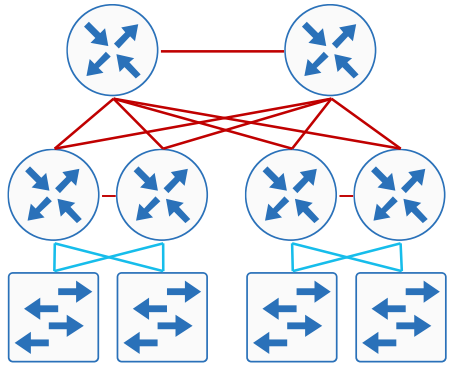

Physical topology: 2 core 2 dist./acc.

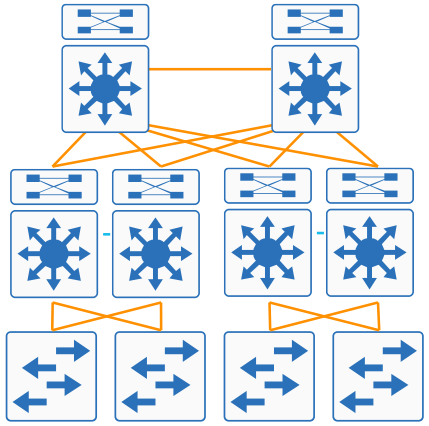

- Common design since the 1990's
- Complex configurations (prone to human error) related to spanning-tree, load balancing, unicast and multicast routing
- Requires heavy performance tuning resulting from reliance on FHRPs (HSRP, VRRP, GLBP)

Survives device and link failures	✔
Easy mitigation of Layer 2 looping concerns	
Rapid detection/recovery from failures	
Layer 2 across all access blocks within distribution	✔
Device-level CLI configuration simplicity	
Automated network and policy provisioning included	

Design Options

Layer 3 Routed Access

Logical topology–

L3: everywhere

L2: edge only

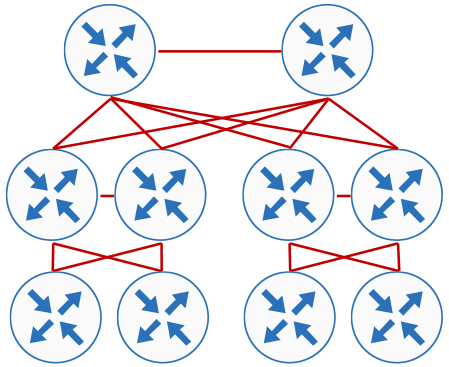

- Complexity reduced for Layer 2 (STP, trunks, etc.)
- Elimination of FHRP and associated timer tuning
- Requires more Layer 3 subnet planning; might not support Layer 2 adjacency requirements

Physical topology:

2 core

2 dist./acc.

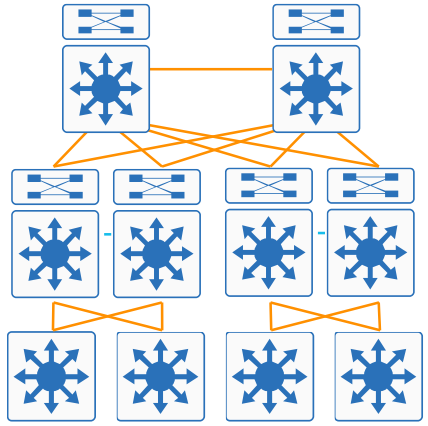

Survives device and link failures	✔
Easy mitigation of Layer 2 looping concerns	✔
Rapid detection/recovery from failures	✔
Layer 2 across all access blocks within distribution	
Device-level CLI configuration simplicity	✔
Automated network and policy provisioning included	

HA Design Options

Simplification with routed access design: Layer 3 distribution with Layer 3 access

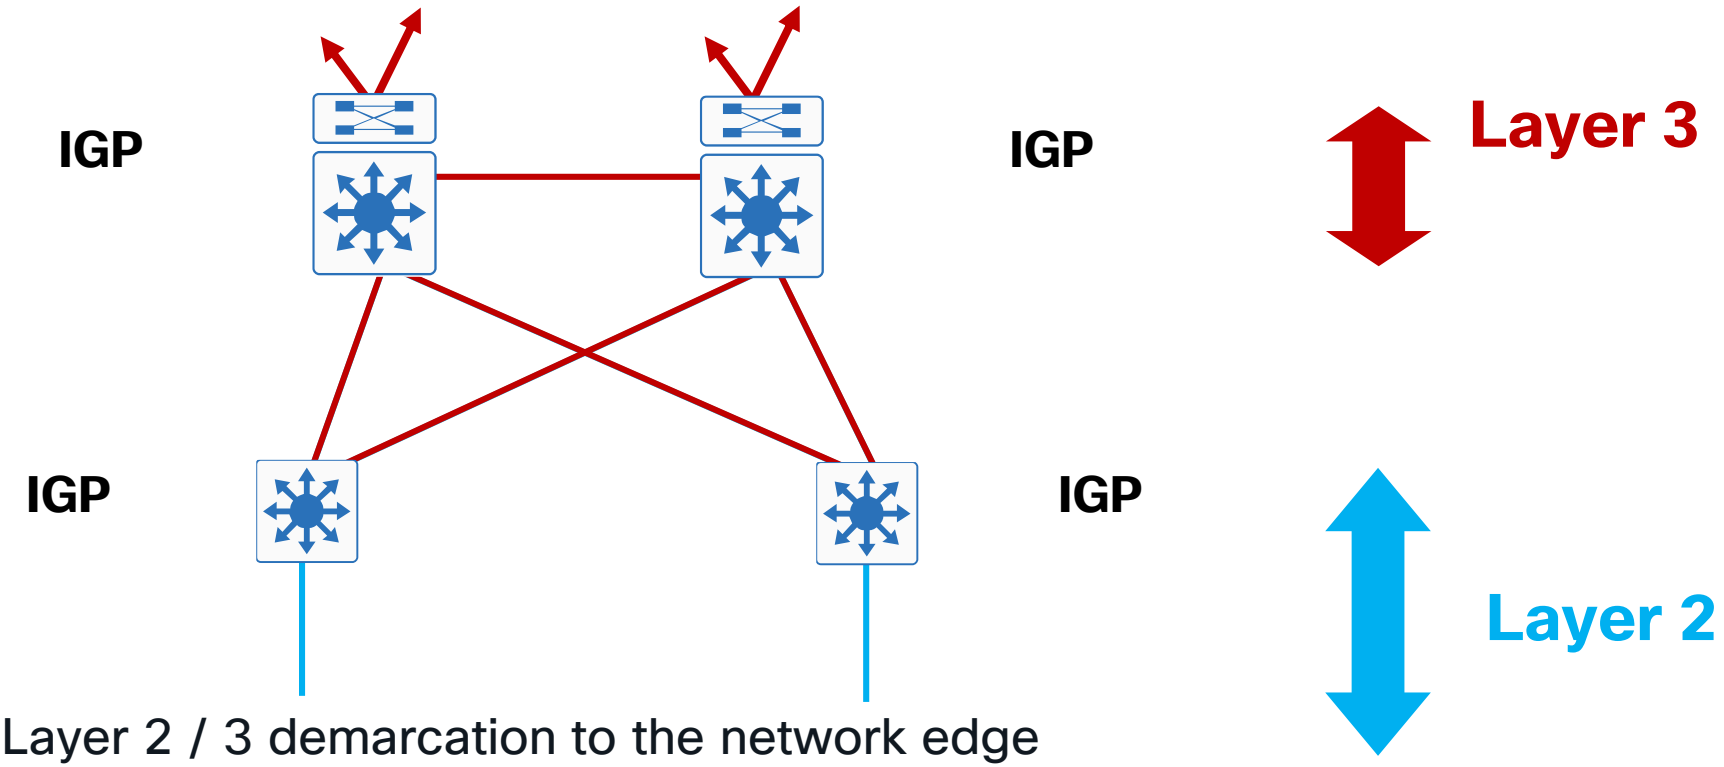

- Move the Layer 2 / 3 demarcation to the network edge
- Leverages Layer 2 only on the access ports, but builds a Layer 2 loop-free network
- **Design Motivations** – Simplified control plane, ease of troubleshooting, highest availability

Design Options

Layer 2 Access with “Simplified” Distribution

Logical topology–
L3: core/dist.
L2: dist./acc.

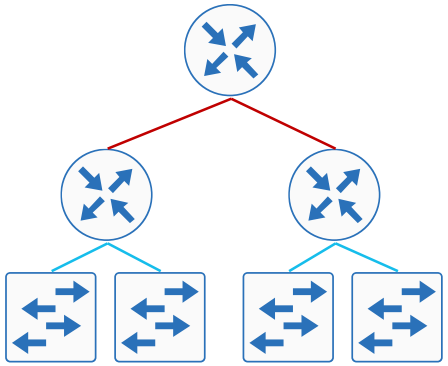

Physical topology:
2 core
2 dist./acc.

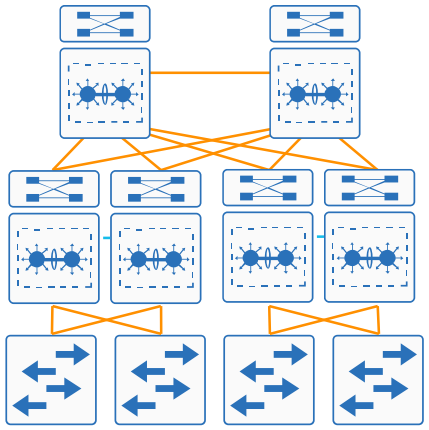

- Leading campus design for easy configuration and operation when using stacking or similar technology (VSS, StackWise Virtual)
- Flexibility to support Layer 2 services within distribution blocks, without FHRPs.
- Easy to scale and manage

Survives device and link failures	✔
Easy mitigation of Layer 2 looping concerns	✔
Rapid detection/recovery from failures	✔
Layer 2 across all access blocks within distribution	✔
Device-level CLI configuration simplicity	✔
Automated network and policy provisioning included	

Cisco StackWise technology

Manage one logical switch - with up to 8 physical switches

Catalyst 9200/L Series StackWise-160/80

- Catalyst 9200 Series stacking of up to 8 switches and 416 ports
- StackWise-160 is supported on Catalyst 9200 switch models
- StackWise-80 is supported on Catalyst 9200L switch models

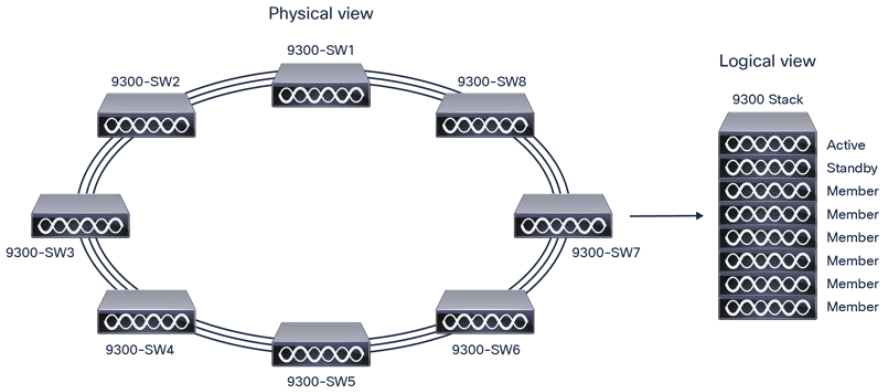

Catalyst 9300/L Series StackWise-480/360

- Catalyst 9300 Series stacking of up to 8 switches and 448 ports
- StackWise-480 is supported on Catalyst 9300 switch models
- StackWise-360 is supported on Catalyst 9300L switch models

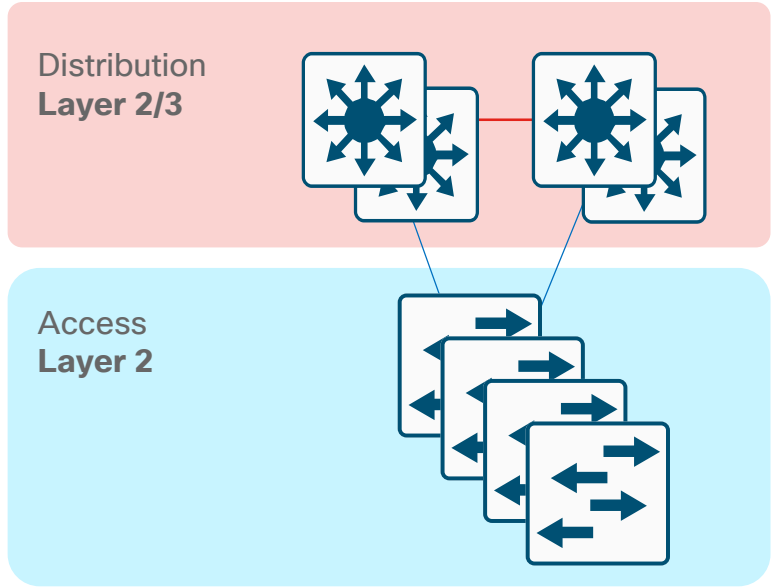

Catalyst 9300X/LM Series StackWise-1T

- Catalyst 9300X Series stacking of up to 8 switches and 448 ports

Cisco StackWise technology

StackWise with Multichassis EtherChannel (MEC)

- Allows up to a of 8 switches to be stacked together physically in a ring topology to form a single, unified, virtual stack system.
- Unified control and management plane by electing one switch in the stack as the Active switch and another switch as the Hot-Standby.
  - Remaining switches become **stack** members
- Cross-stack Multichassis EtherChannel (MEC) extends traditional EtherChannel by allowing member ports to be aggregated on different physical chassis

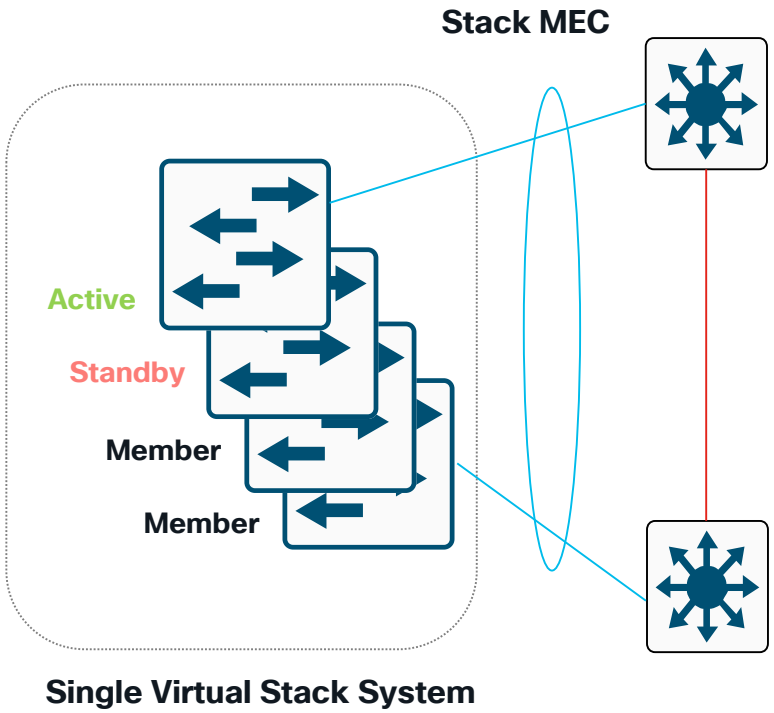

StackWise Access

The StackWise Access PIN focuses on combining multiple Access switches into a single virtual switch to increase access-layer port density.

- Typically, the same layer as Access (Tier 1)
- The same 'physical' topology as a multi-layer network

Main goal is to expand port density of Access layer

Same L2 protocols & features as Access

- North: VLAN, 802.1Q, STP, MAC, IGMP Snooping
- South: AAA, STP, Portfast, Storm-Control

Leverages **Stateful Switchover (SSO)**

- Active/Standby Control-Plane (synchronized)
- Works with NSF/NSR for L3 protocols

Leverages **Multi-chassis EtherChannel (MEC)**

- Active/Active Data-Plane (both switches forwarding)
- L2 Portchannel (neighbor sees single neighbor)

Tends to require **med-high L2 + feature scale**

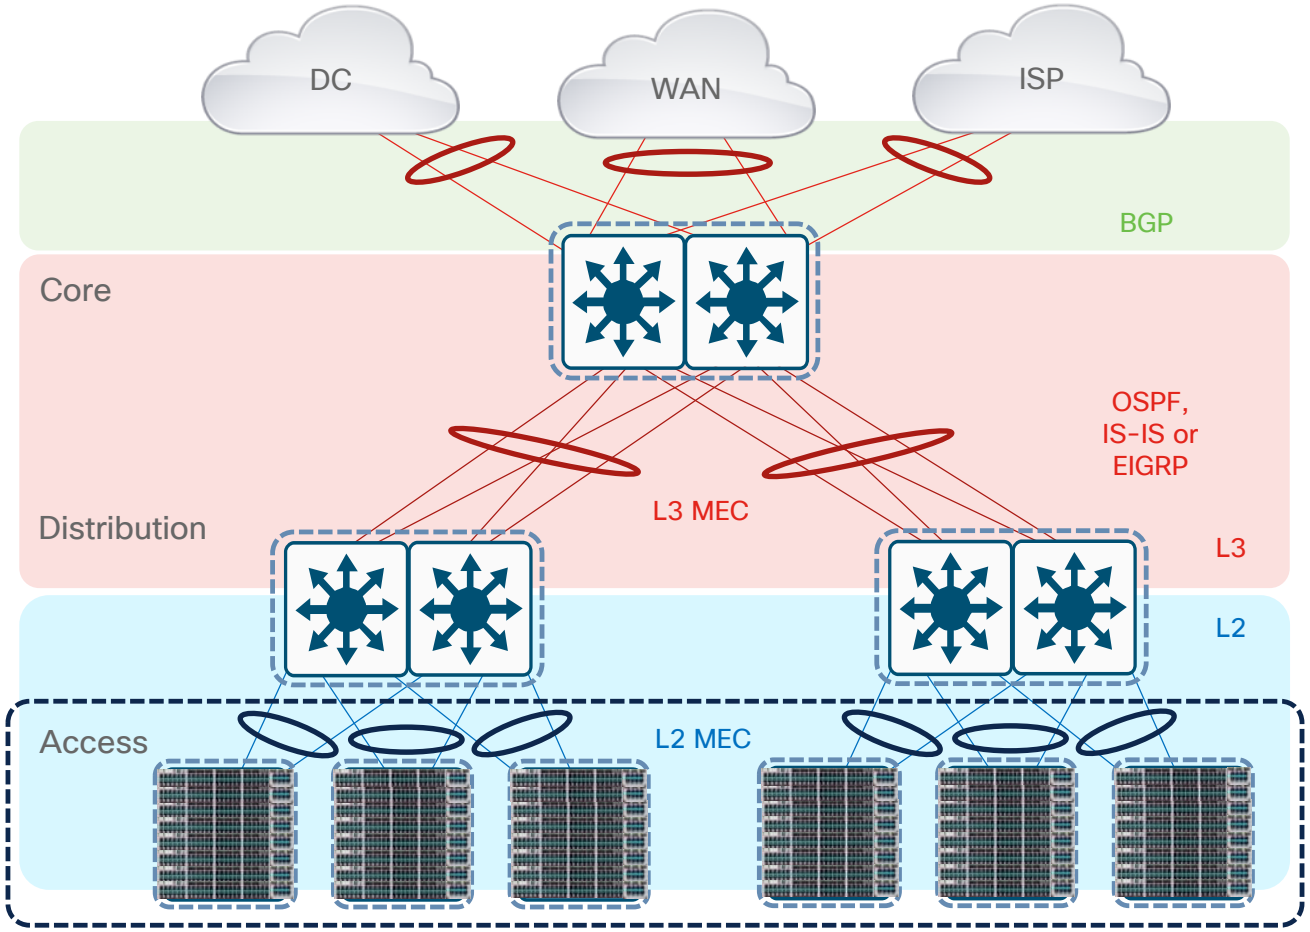

StackWise Virtual Technology


Catalyst 9600/X


Catalyst 9500/X


Catalyst 9400/X

- Intended for **Distribution** and/or **Core** layer
- Available on **C9400**, **C9500** and **C9600**
- Formed using **Front Panel** ethernet ports

up to 8x
10/25/50G SFP

up to 8x
40/100/400G QSFP

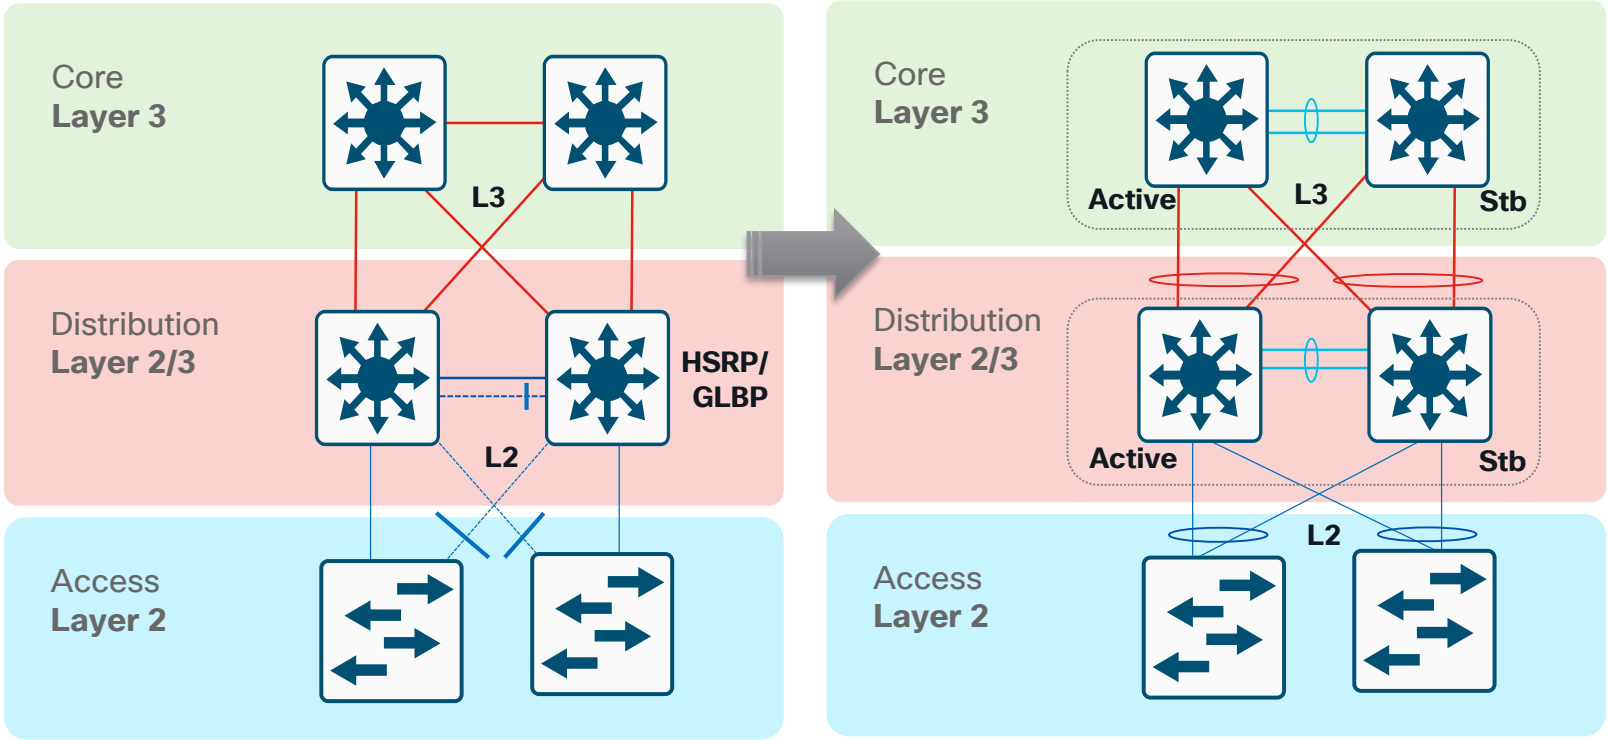

- **Simplify Operations** by Eliminating STP, FHRP and Multiple Touch-Points
- **Double Bandwidth** and Reduce Latency with Active-Active Multi-chassis EtherChannel (MEC)
- **Minimize Convergence** with Sub-second Stateful and Graceful Restart (SSO/NSF)

StackWise Virtual Core/Distro

The **StackWise Virtual (SVL)** Core PIN focuses on combining Core and/or Distribution into a single virtual switch to connect to outside areas.

- Typically, the same layer as Distribution or Core (Tier 2-3)
- The same 'physical' topology as a multi-layer network

Main goal is to simplify and expand the Distribution and/or Core layer

Same L2/L3 protocols & features as Distro/Core

- North: SVI, ARP/ND, IGP/BGP, PIM
- South: VLAN, 802.1Q, MAC, IGMP (No STP)

Leverages **Stateful Switchover (SSO)**

- Active/Standby Control-Plane (synchronized)
- Works with NSF/NSR for L3 protocols

Leverages **Multi-chassis EtherChannel (MEC)**

- Active/Active Data-Plane (both switches forwarding)
- L2 & L3 Portchannel (neighbor sees single neighbor)

Tends to require **med-high L2, L3 & feature scale**

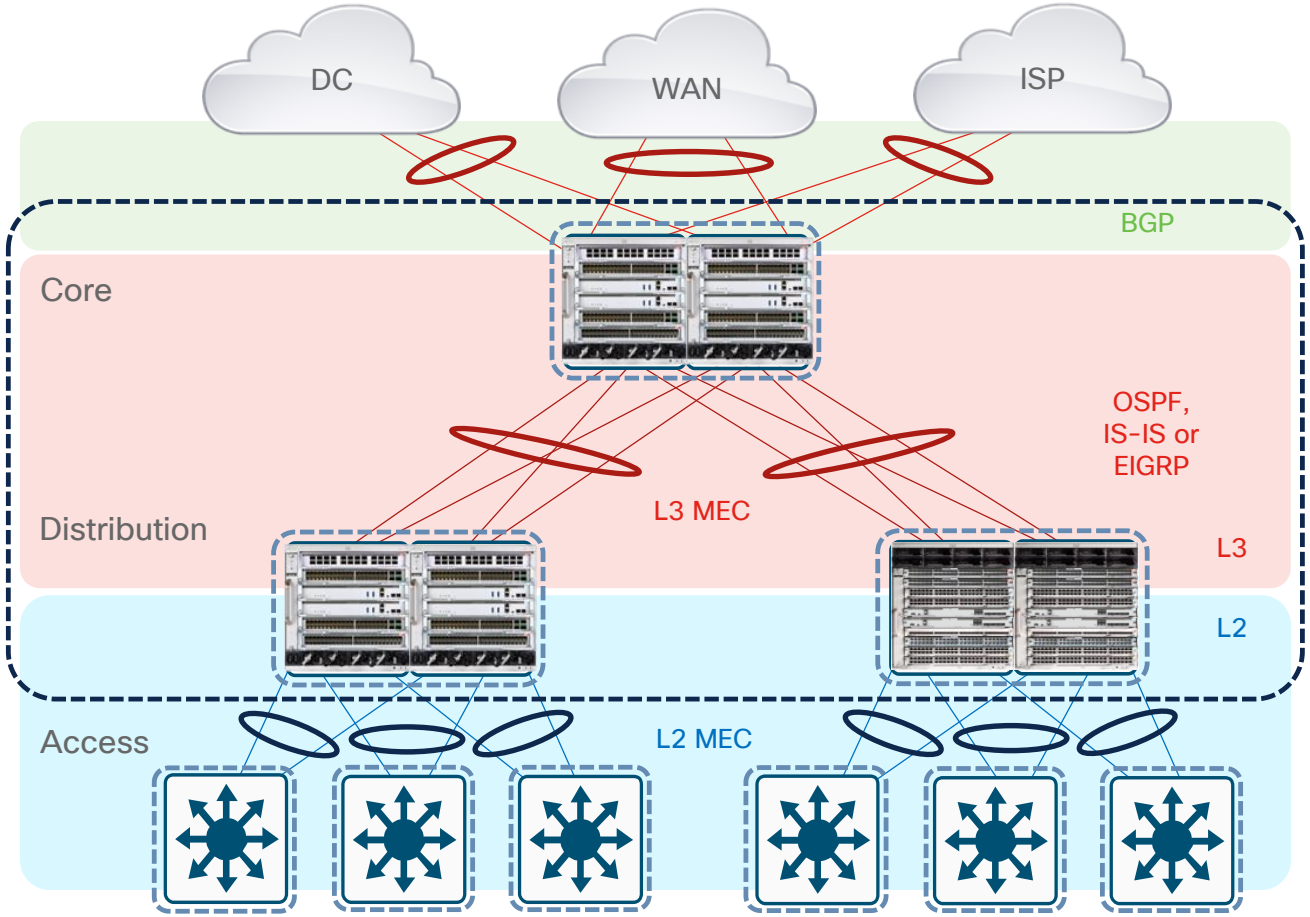

LAN High Availability

StackWise Virtual – Distro/Core Switch Stacking

Catalyst 9600/X

Catalyst 9500/X

Catalyst 9400/X

SSO Active responsible for:
- Management
- L2 protocols
- L3 protocols

BRKENS-1500

Traditional L2/L3

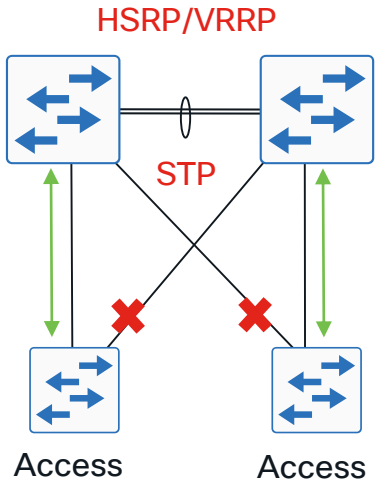

StackWise Virtual – Physical

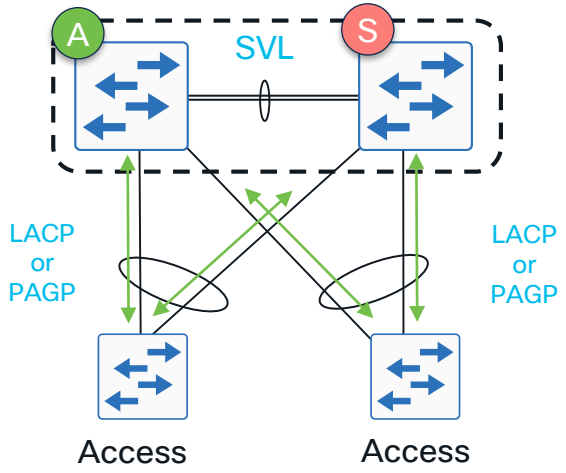

StackWise Virtual – Logical

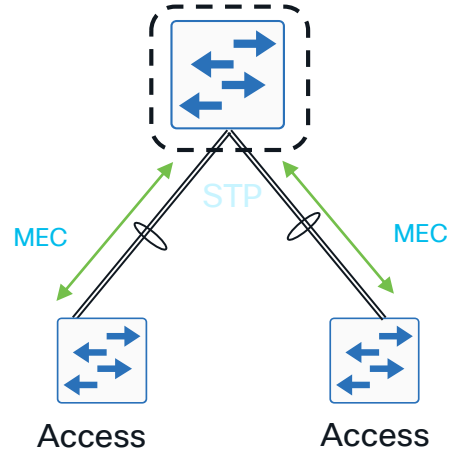

Both Active & Standby switches have Active data plane and make forwarding decisions

LAN High Availability

SWV/VSS: connecting Distribution to Access layer

- General Best-Practice - always use EtherChannel for link resiliency and load sharing
- For SWV/VSS - use **Multi-chassis EtherChannel (MEC)** and dual-home to each switch

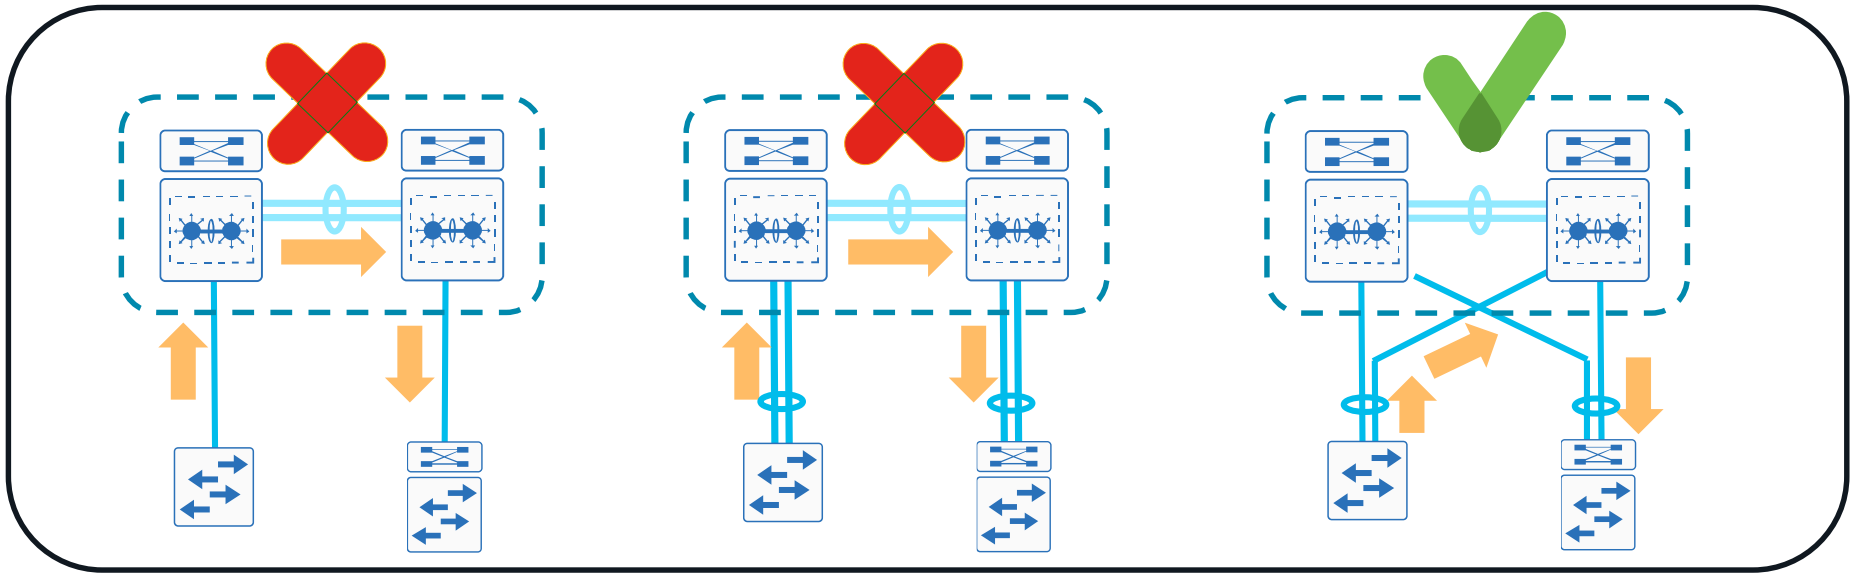

Routed Access vs. Virtual Switching

Evolutions of and Improvements to Existing Designs

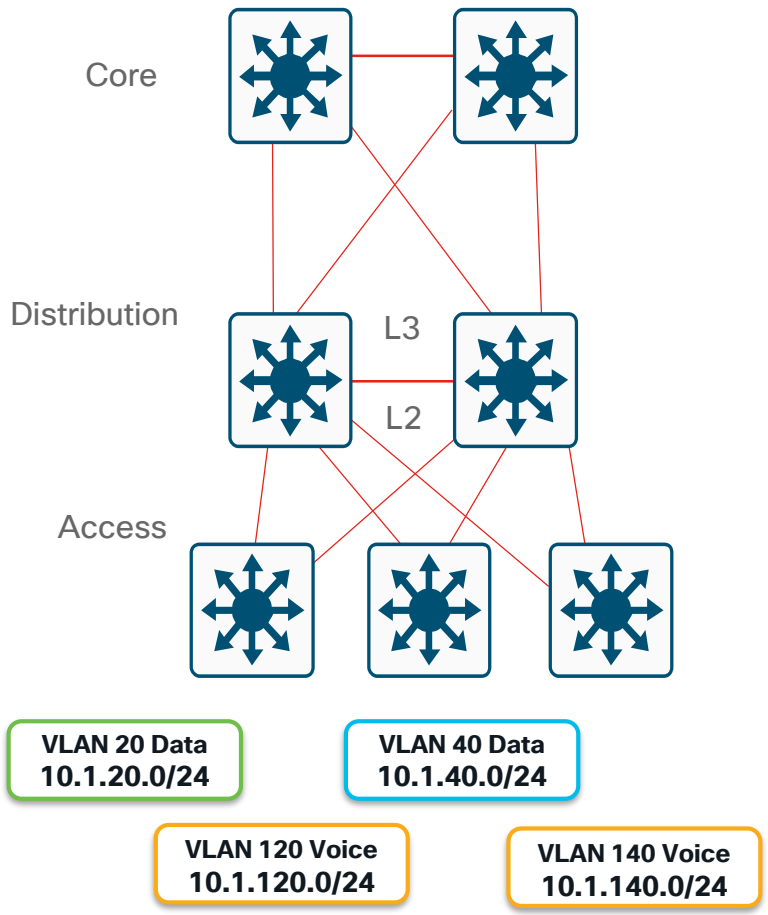

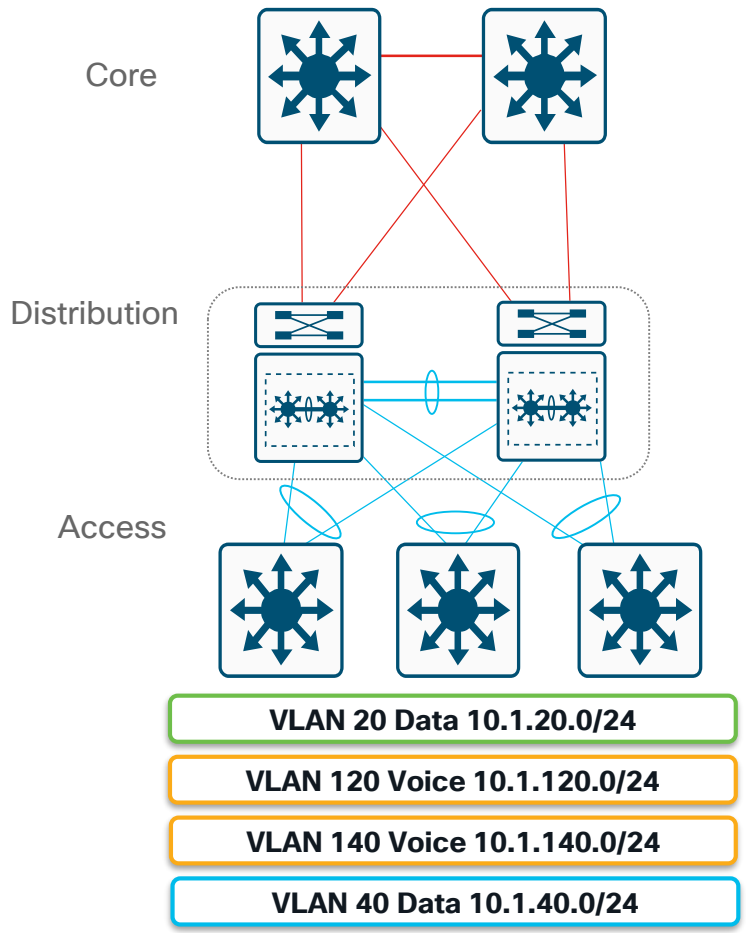

Introducing Next Generation Stackwise

A new era of micro-services device clustering

A new Stacking Architecture
based on standard Ethernet protocols!

Cloud-Native IOSXE includes a new Stack architecture based on standard Ethernet protocols and encapsulation. Using either back-side or front-side ethernet links - for greater scale, stability, better SSO convergence, and ISSU/XFSU & SMU for hitless software upgrades.

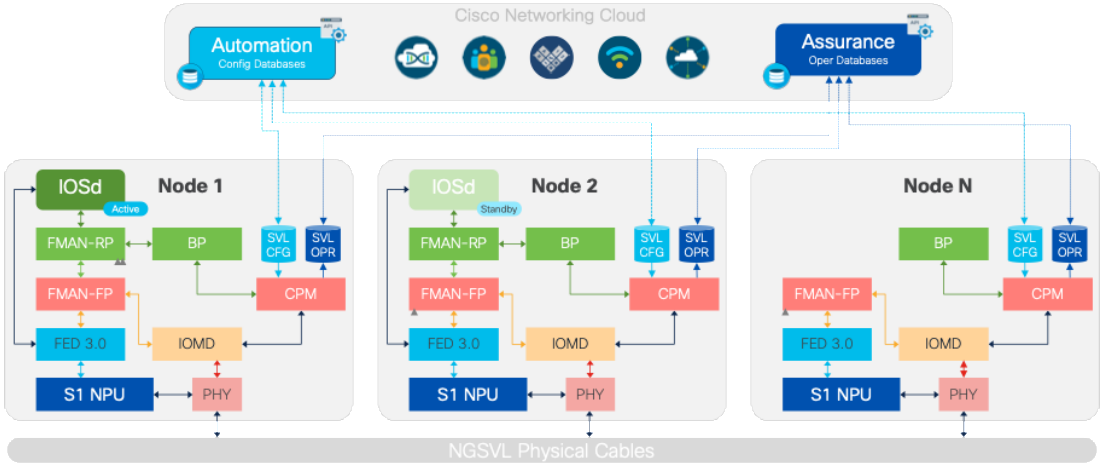

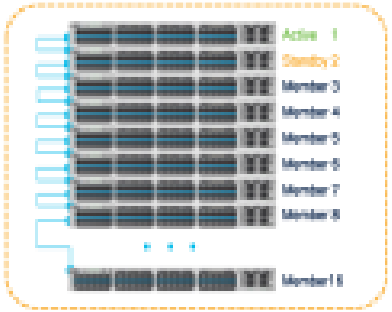

How is the Architecture new?

- A new Bootstrap Process (BP) & Cluster Process Manager (CPM), as a separate Linux thread.
- Runs IP SPF + VXLAN-GPO over standard Ethernet

Is it front panel or back panel?

- Both! Dedicated cables or Transceivers + cables.
- Stack Interface (SIF) is treated as any ethernet port

Can I still manage it the same?

- Familiar (same) config and show commands
- Uses same '*stackwise*' & '*stackwise-virtual*' CLI

How does it benefit me?

- Common to both C9350 & C9610
- Simplified cabling & dynamic link add/change
- Improved SSO and ISSU/SMU convergence

... and much more, coming soon

What is Next Gen Stackwise?

Introducing an entirely new Stacking Architecture

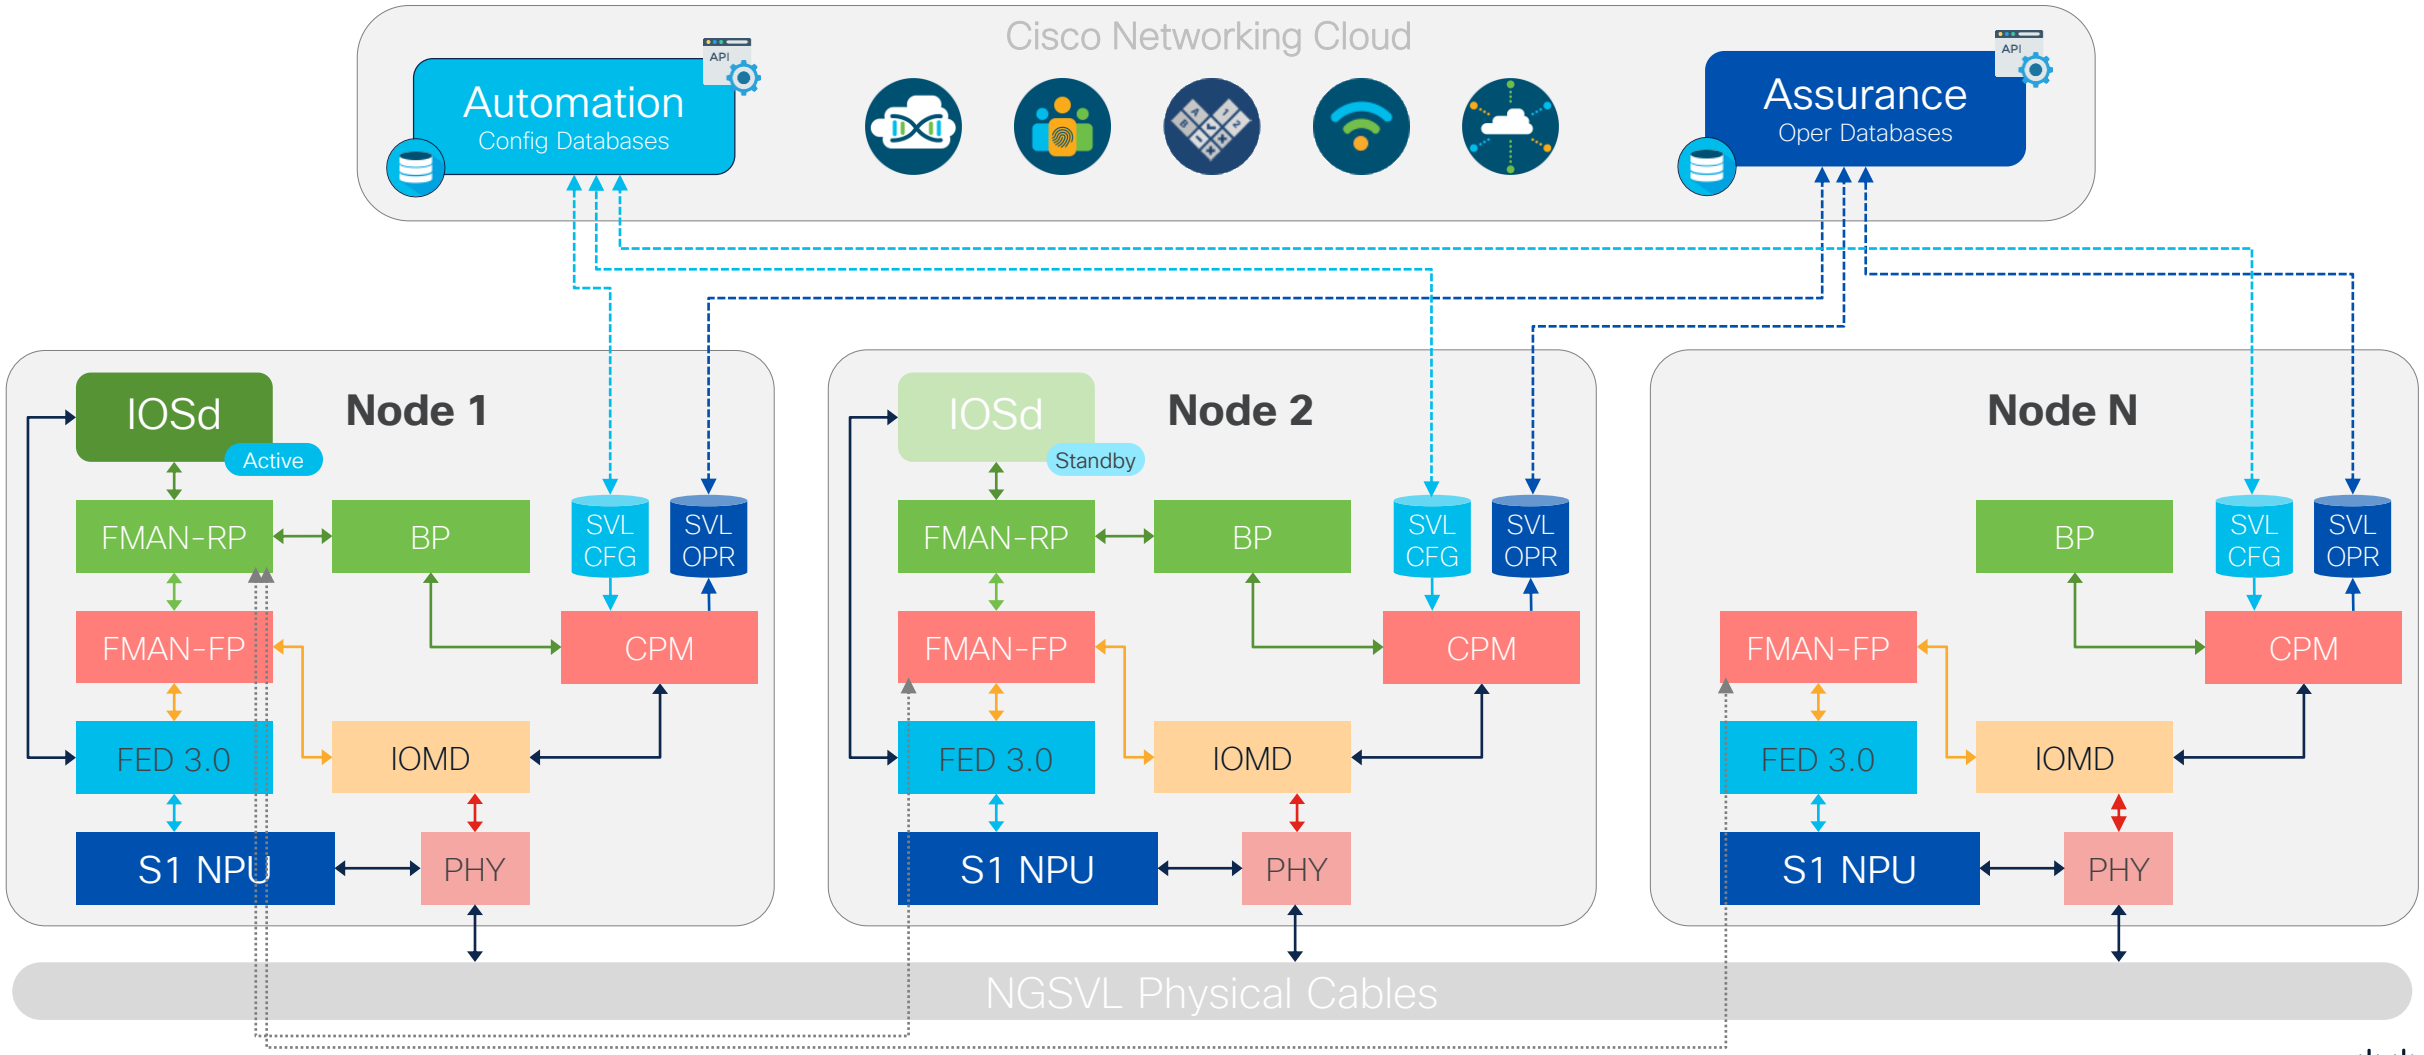

NOTE: Simplified Diagram

Campus with Next Generation Stackwise

Reducing the number of Nodes & Peers, with superior L2 & L3 convergence

L2/L3 Equal Cost Multi-Path

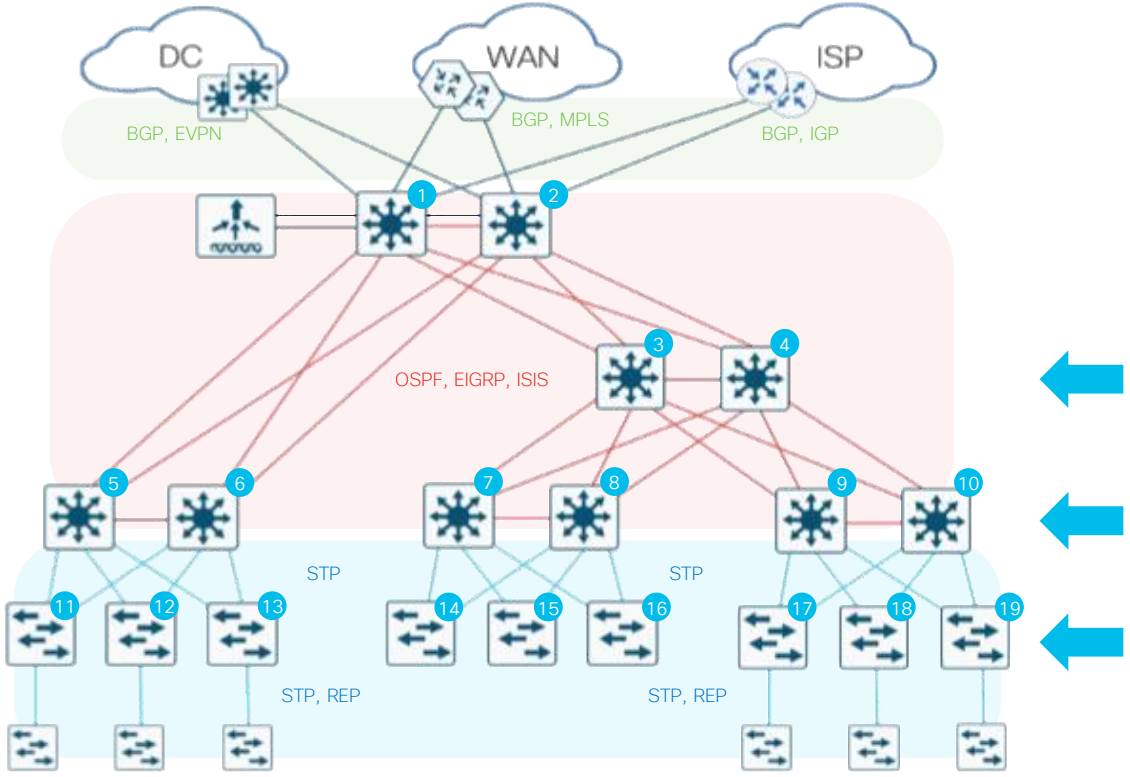

- ✕ 19 Nodes to manage (IPs, Syslog, FNF, SNMP, YANG, etc.)
- ✕ 7 Core IGP links & peers (10 IGP speakers)
- ✕ 4 Access STP links & peers (8 STP speakers)
- ✕ Failover based on Protocol + CEF (seconds)

L2/L3 with Stacking & MEC

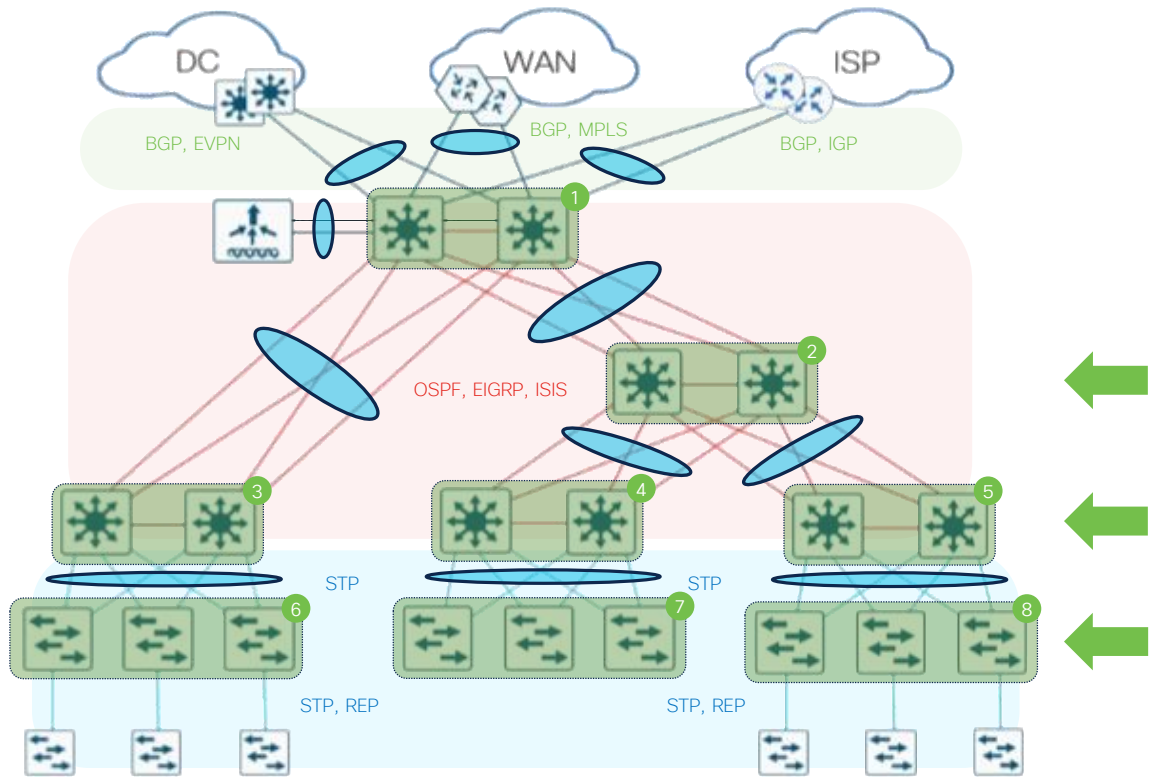

- ✓ 8 Nodes to manage (IPs, Syslog, FNF, SNMP, YANG, etc.)
- ✓ 3 Core IGP links & peers (4 IGP speakers)
- ✓ 1 Access STP links & peers (5 STP speakers)
- ✓ Failover based on MEC (150-200 milliseconds)

Resilient Campus Deployment Options

FHRP Based

HSRP | GLBP | VRRP

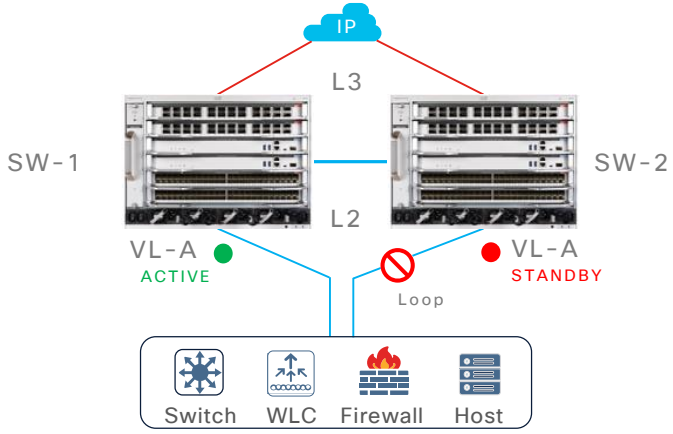

ESI mLAG

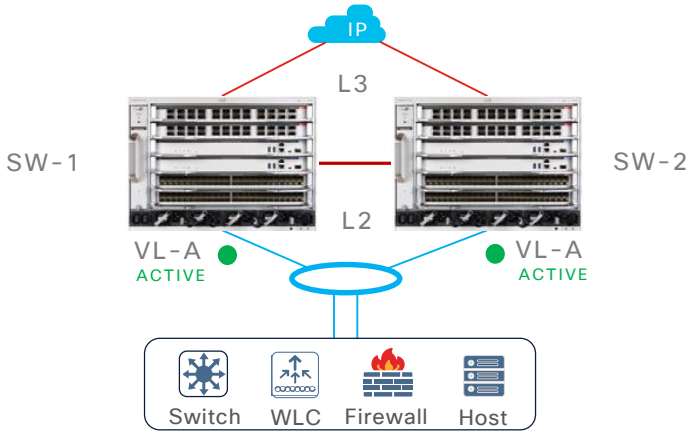

SWV MEC

Best

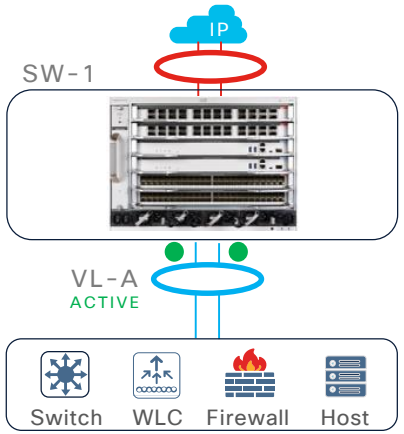

Best-In-Class Resiliency

Broad L3 IP gateway redundancy design alternatives

- Traditional FHRP-Based IP gateway redundancy - HSRP, GLBP and VRRP
- Industry-standard Layer 2 Multipath Network with Multi-Chassis LAG (mLAG)
- Cisco StackWise Virtual unified system for resilient, scalable and simplified networks

Design HA with Catalyst 9K

BRKENS-2095

Designing Highly Available Networks Using Cisco Catalyst 9000 Switches

Minhaj Uddin - Leader Technical Marketing, Cisco

This session will explore both new and existing high-availability features in IOS XE on Catalyst 9000 Series switching platforms.

We will begin by highlighting the significance of high availability across various layers of the hierarchical network. Following this, we will delve into different levels of resiliency, including standalone platform/hardware, design, and software. The session will conclude with a summary of these capabilities, illustrated through various real-world customer use cases and requirements.

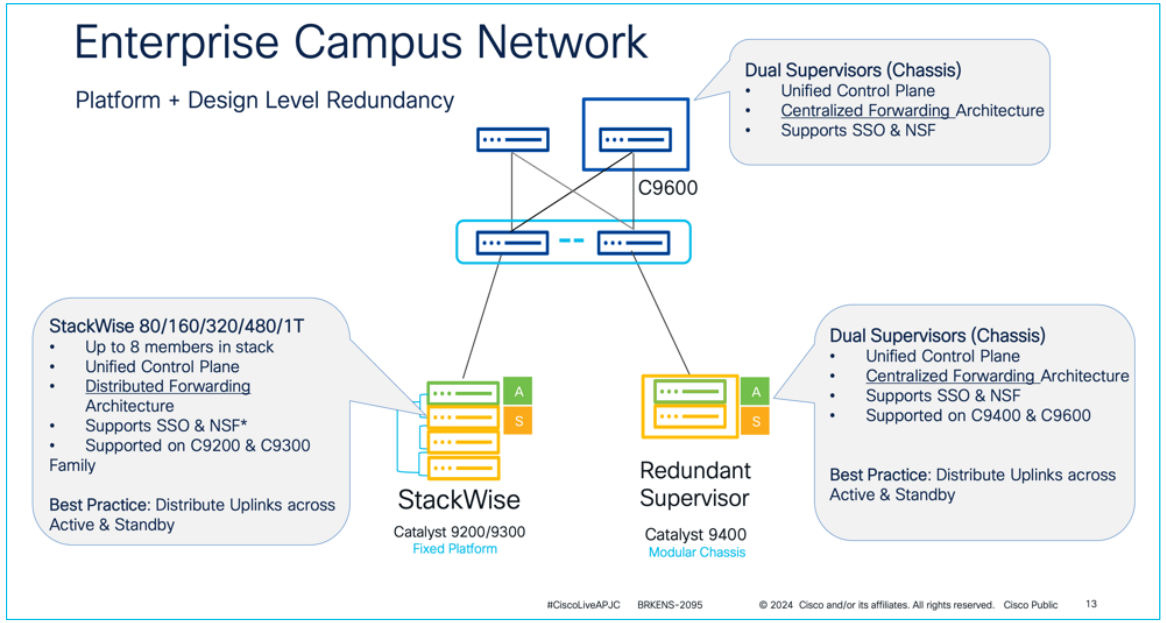

Platform Design

- LAN High Availability
- LAN Security
- Cisco Zero Trust (Overview)
  - User Authentication
  - First Hop Security
  - Protocol Authentication
  - Data Encryption
  - Segmentation
- Virtual Networking

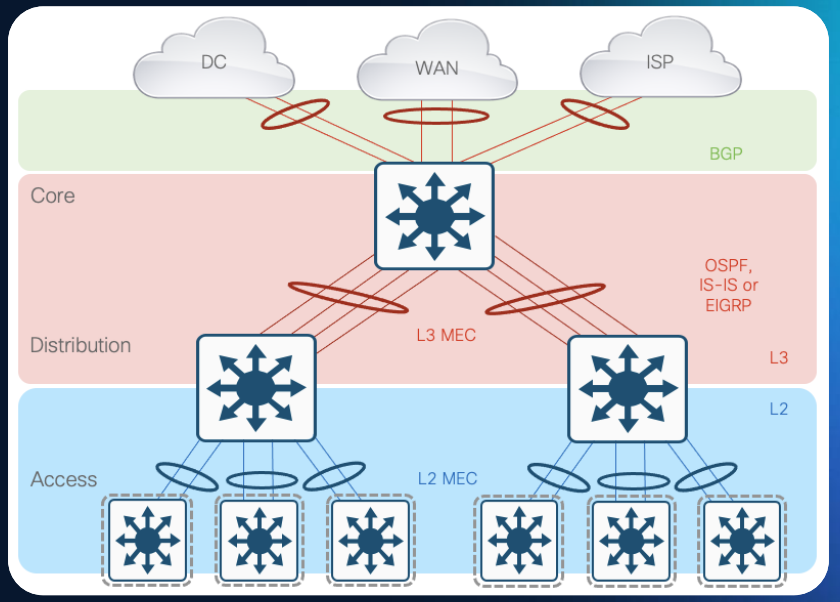

Campus Security

End to end Security occurs at different network layers

- Cisco Zero Trust
- User Authentication
  - 802.x, MAB, CWA
- First Hop Security
- Protocol Authentication
- Data Encryption
- Segmentation

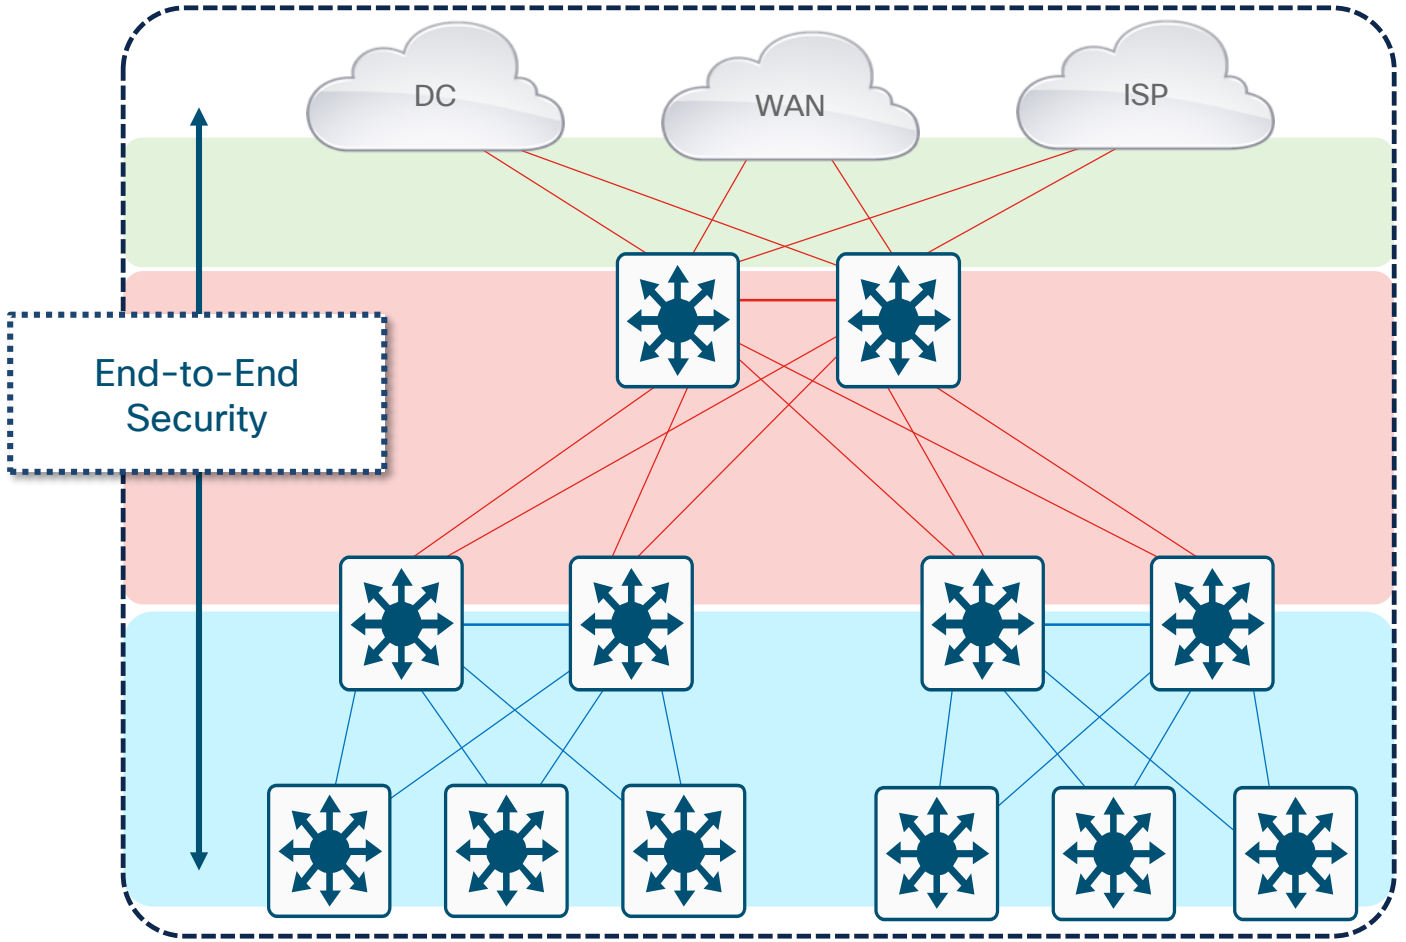

Three pillars of Zero Trust Workplace

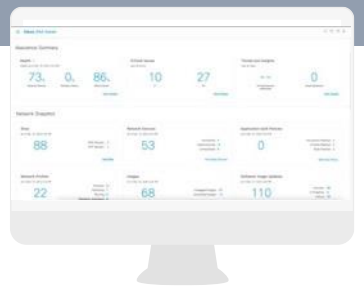

Cisco Catalyst Center

Endpoint Visibility

Segmentation

Trust Assessment

Cisco Identity Services

Enabled on Cisco **Catalyst 9K** Infrastructure

Zero Trust Workplace

Design Considerations

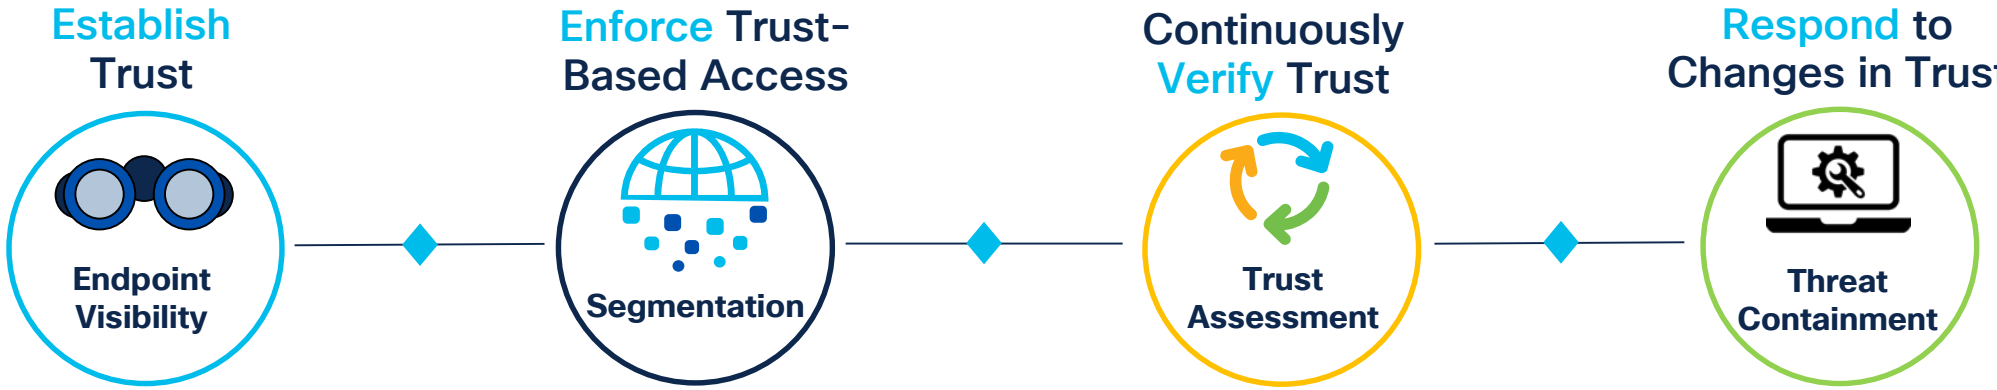

- User/Device Authentication
- MFA thru Integrations
- Endpoint Profiling
- Posture + Context
- Guest Access
- BYOD Onboarding

- Network-based Authorization Policies
- Micro-segmentation
- Compliance-based Change of Auth (CoA)
- Device Administration with TACACS+

- Integrations for Threat Detection
- Behavior Analysis, Vulnerabilities

- Adaptive Network Access Control
- Orchestrated Remediation (CoA)

Expand Threat Remediation - Secure Analytics

Secure The Access – Software Defined Network

Start with the Right Foundation – Access Policy

Security Best Practices

BRKENS-1500

Also protects limited Hardware & Software sources

Complexity

- **Cisco Umbrella** ➢ uses DNS as a security tool to identify and block threats
- **802.1x User Authentication** ➢ forces users to authenticate before allowing them on network
- **IP Source Guard / v6 RA Guard** ➢ prevents IP/MAC Spoofing and IPv6 Man-in-the-Middle attacks
- **Dynamic ARP Inspection** ➢ prevents current ARP/NDP attacks
- **DHCP Snooping** ➢ prevents Rogue DHCP Server attacks
- **Port Security** ➢ prevents CAM attacks or DHCP Starvation attacks

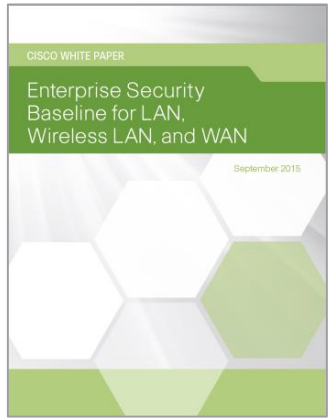

Security Best Practices

Port Security - Cutting Off MAC-Based Attacks

- Cisco Umbrella
- 802.1x
- IP Source Guard / IPv6 RA Guard
- Dynamic ARP Inspection
- DHCP Snooping
- Port Security

Protect your switch from CAM table overflow attacks
(Content Addressable Memory)

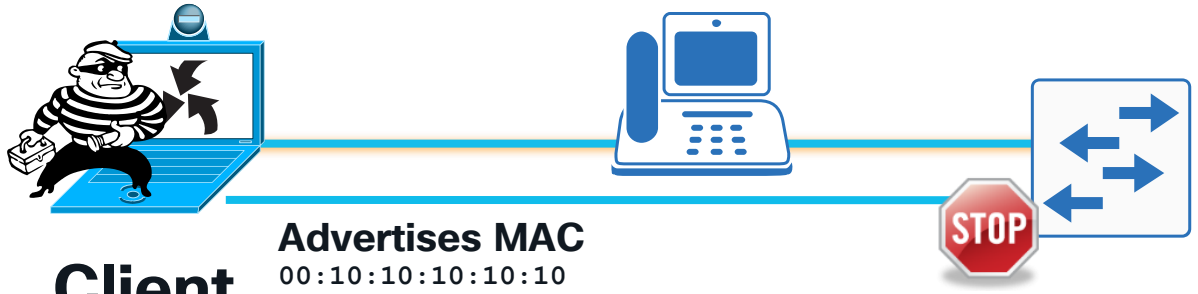

Client

Advertises MAC

```
00:10:10:10:10:10
00:10:10:10:10:11
00:10:10:10:10:12
00:10:10:10:10:13
00:10:10:10:10:14
00:10:10:10:10:15
00:10:10:10:10:16
00:10:10:10:10:17
00:10:10:10:10:18
00:10:10:10:10:19
00:10:10:10:10:1A
00:10:10:10:10:1B
```

Configure on the client interface:

```
switchport port-security
switchport port-security maximum 11
switchport port-security aging time 2
switchport port-security aging type inactivity
switchport port-security violation restrict
```

Exceeds Maximum

Port security and **dot1x authentication** are mutually exclusive

Security Best Practices

DHCP Snooping - Protection Against Rogue/Malicious DHCP Server

- Cisco Umbrella
- 802.1x
- IP Source Guard / IPv6 RA Guard
- Dynamic ARP Inspection
- DHCP Snooping
- Port Security

- DHCP requests (discover) and responses (offer) are tracked
- Rate-limit requests on trusted interfaces - limits DoS attacks on DHCP server

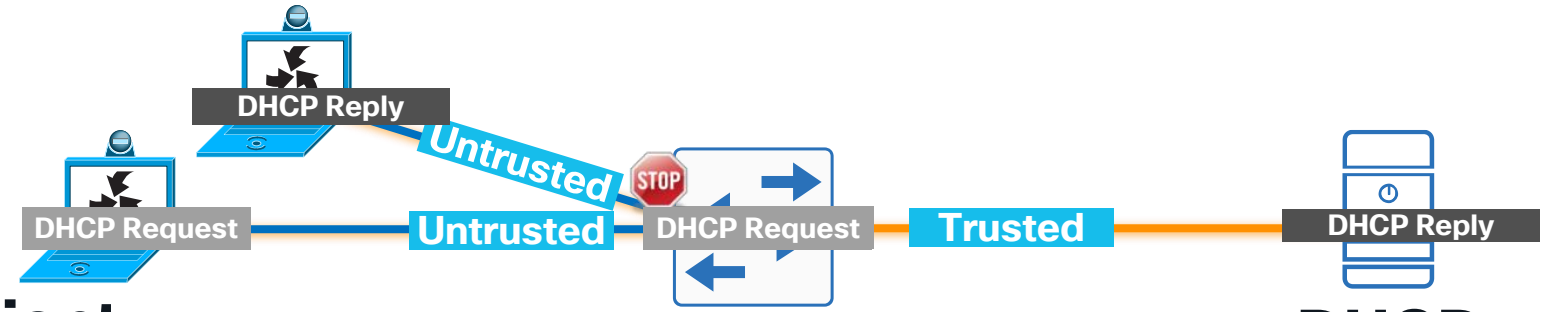

DHCP Snooping Binding Table

MAC Address	IP Address	VLAN	Interface
00:50:56:BA:13:DB	10.4.80.10	10	GigabitEthernet2/0/1

Configure in the global configuration:

```
ip dhcp snooping vlan [data vlan], [voice vlan]
no ip dhcp snooping information option
ip dhcp snooping
```

Configure on the client interface:

```
ip dhcp snooping limit rate 100
```

Security Best Practices

Dynamic ARP Inspection

- Cisco Umbrella
- 802.1x
- IP Source Guard / IPv6 RA Guard
- Dynamic ARP Inspection
- DHCP Snooping
- Port Security

- Dynamic ARP Inspection prevents ARP poisoning (ettercap, dsnif, arpspoof)
- Uses the DHCP Snooping binding table
- Tracks MAC to IP from DHCP transactions

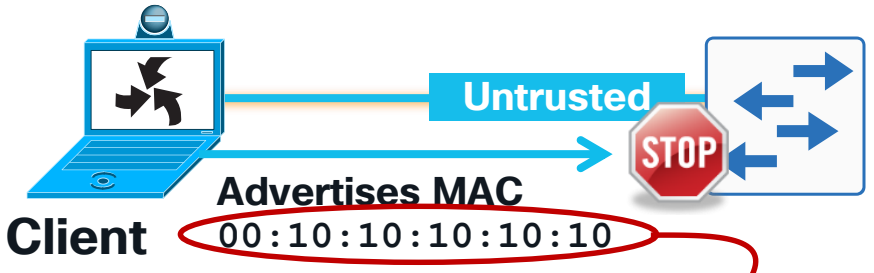

Does Not Match

DHCP Snooping Binding Table

MAC Address	IP Address	VLAN	Interface
00:50:56:BA:13:DB	10.4.80.10	10	GigabitEthernet2/0/1

Configure in the global configuration:

```
ip arp inspection vlan [data vlan], [voice vlan]
```

Configure on the client interface:

```
ip arp inspection limit rate 100
```

Security Best Practices

IP Source Guard

- Cisco Umbrella
- 802.1x
- IP Source Guard / IPv6 RA Guard
- Dynamic ARP Inspection
- DHCP Snooping
- Port Security

- IP source guard protects against spoofed IP addresses
- Uses the DHCP snooping binding table
- Tracks IP address to port associations

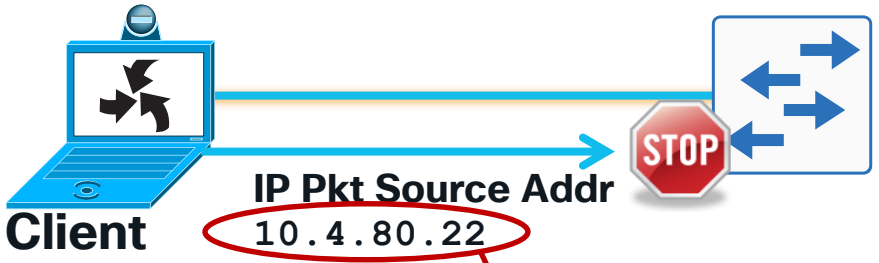

Does Not Match

DHCP Snooping Binding Table

MAC Address	IP Address	VLAN	Interface
00:50:56:BA:13:DB	10.4.80.10	10	GigabitEthernet2/0/1

Configure on the client interface:

```
ip verify source
```

Security Best Practices

IPv6 Router Advertisement Guard

- Cisco Umbrella
- 802.1x
- IP Source Guard / IPv6 RA Guard
- Dynamic ARP Inspection
- DHCP Snooping
- Port Security

Define policy in the global configuration:

```
ipv6 nd raguard policy HOST_POLICY
  device-role host
```

Attach policy configuration to the client interface:

```
ipv6 nd raguard attach-policy HOST_POLICY
```

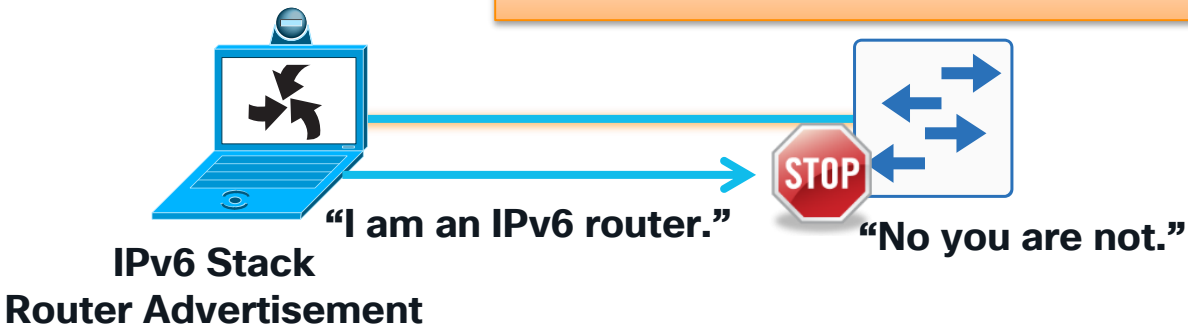

- If a port device role is configured as host, IPv6 First Hop Security (FHS) RA Guard drops all IPv6 Router Advertisement messages
- Useful even for IPv4-only networks
- Other port device role options are: monitor, router, and switch

Security Best Practices

IEEE 802.1x – RADIUS: Access Request

- Cisco Umbrella
- 802.1x
- IP Source Guard / IPv6 RA Guard
- Dynamic ARP Inspection
- DHCP Snooping
- Port Security

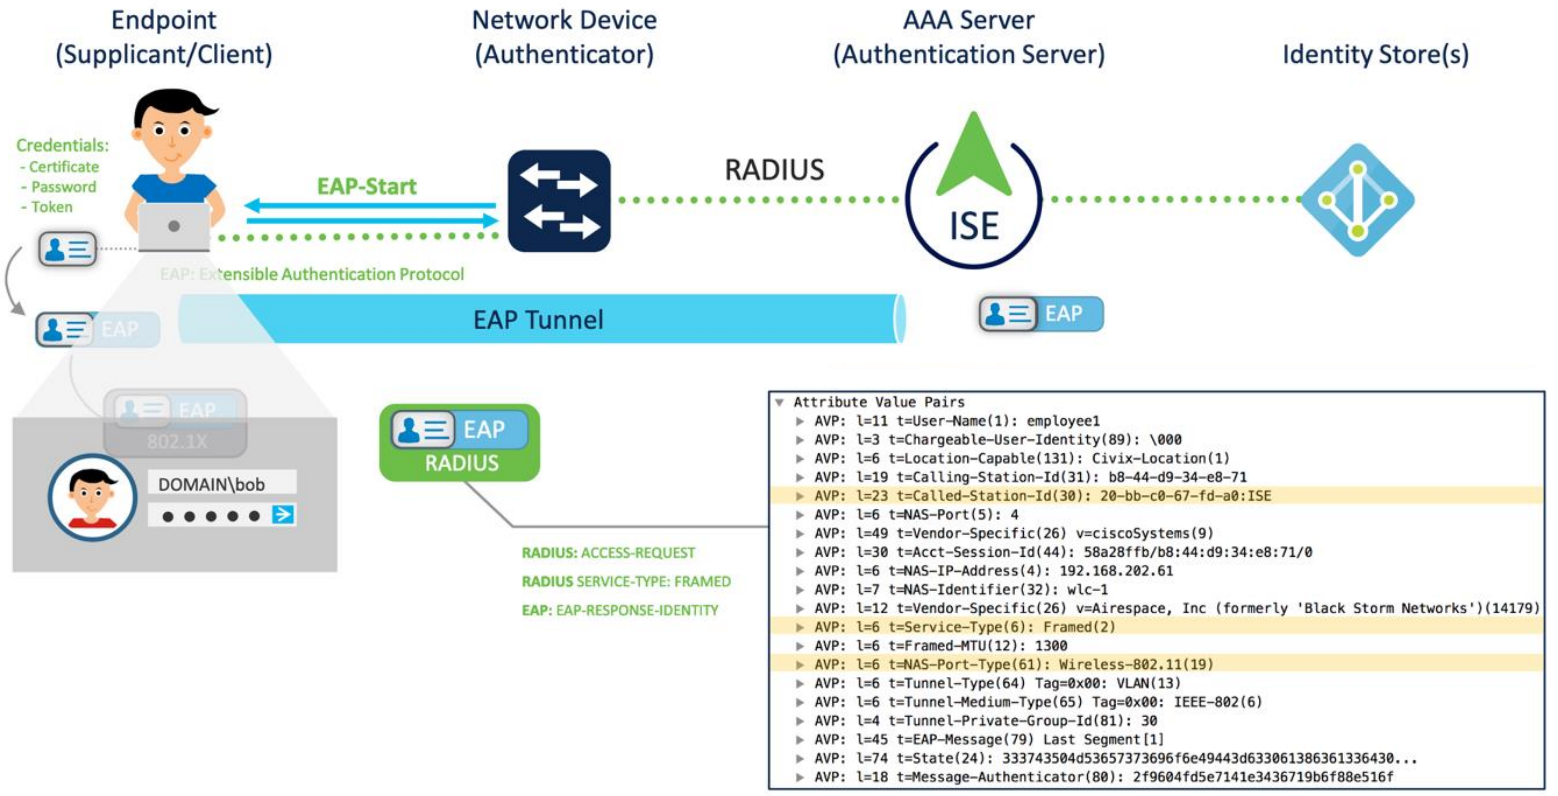

Security Best Practices

IEEE 802.1x – RADIUS: Access Accept

- Cisco Umbrella
- 802.1x
- IP Source Guard / IPv6 RA Guard
- Dynamic ARP Inspection
- DHCP Snooping
- Port Security

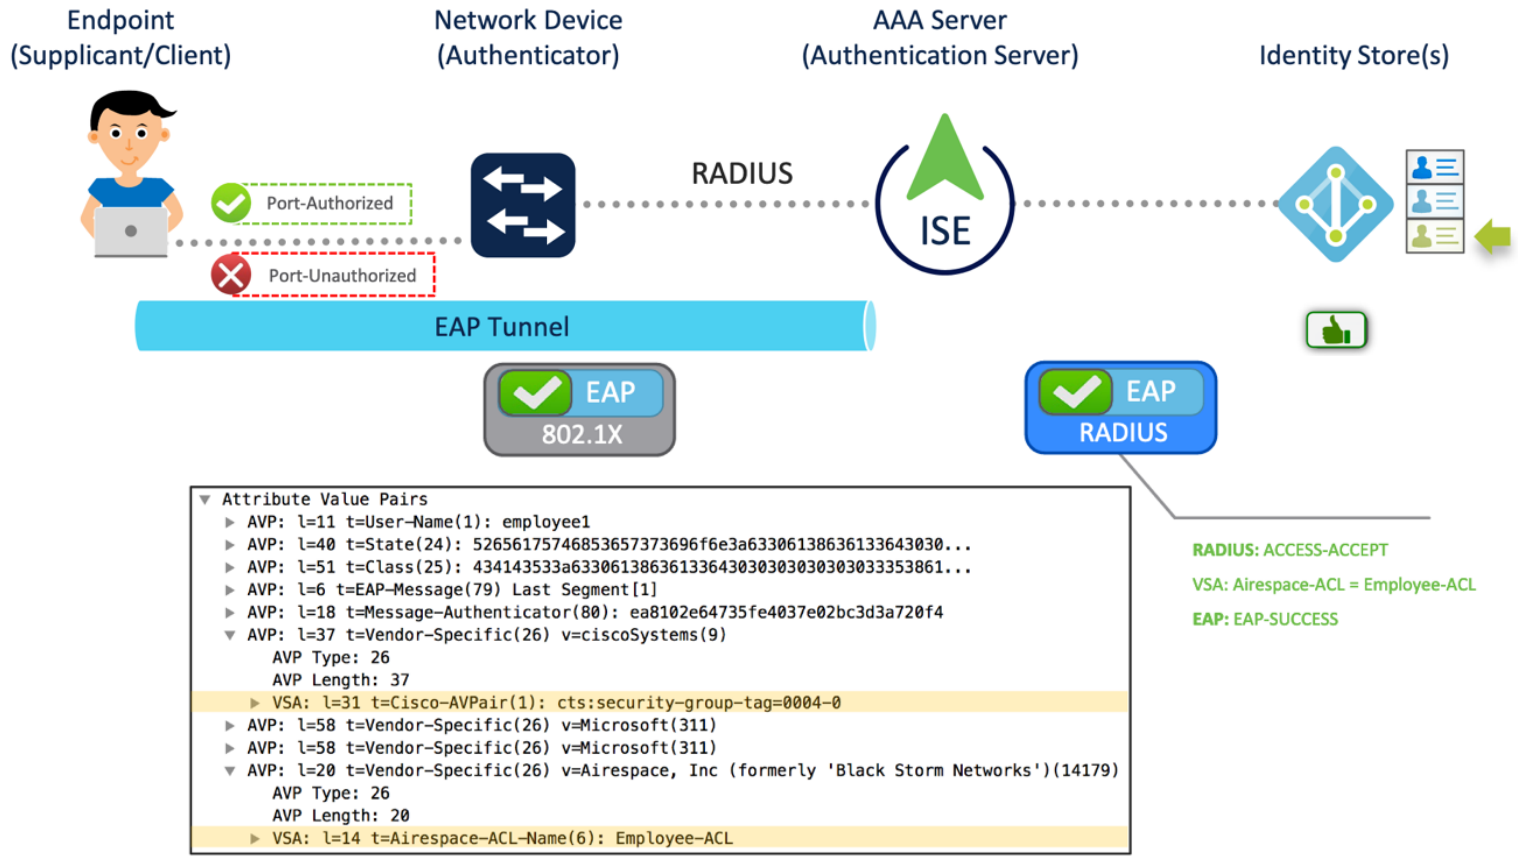

Security Best Practices

Cisco Umbrella – DNS Protection

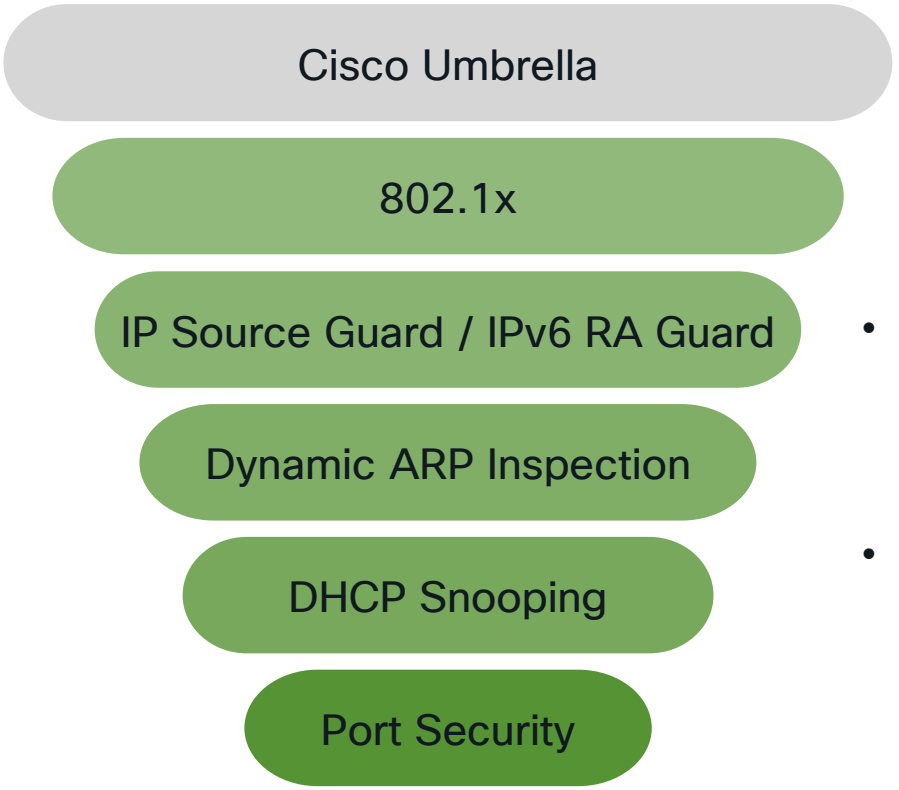

- Cisco Umbrella provides a first line of security for whenever users access the internet (by using DNS as a security tool).
- Since DNS is a core part of the internet – DNS is used to block requests to malicious domains and IP addresses before a connection is established.

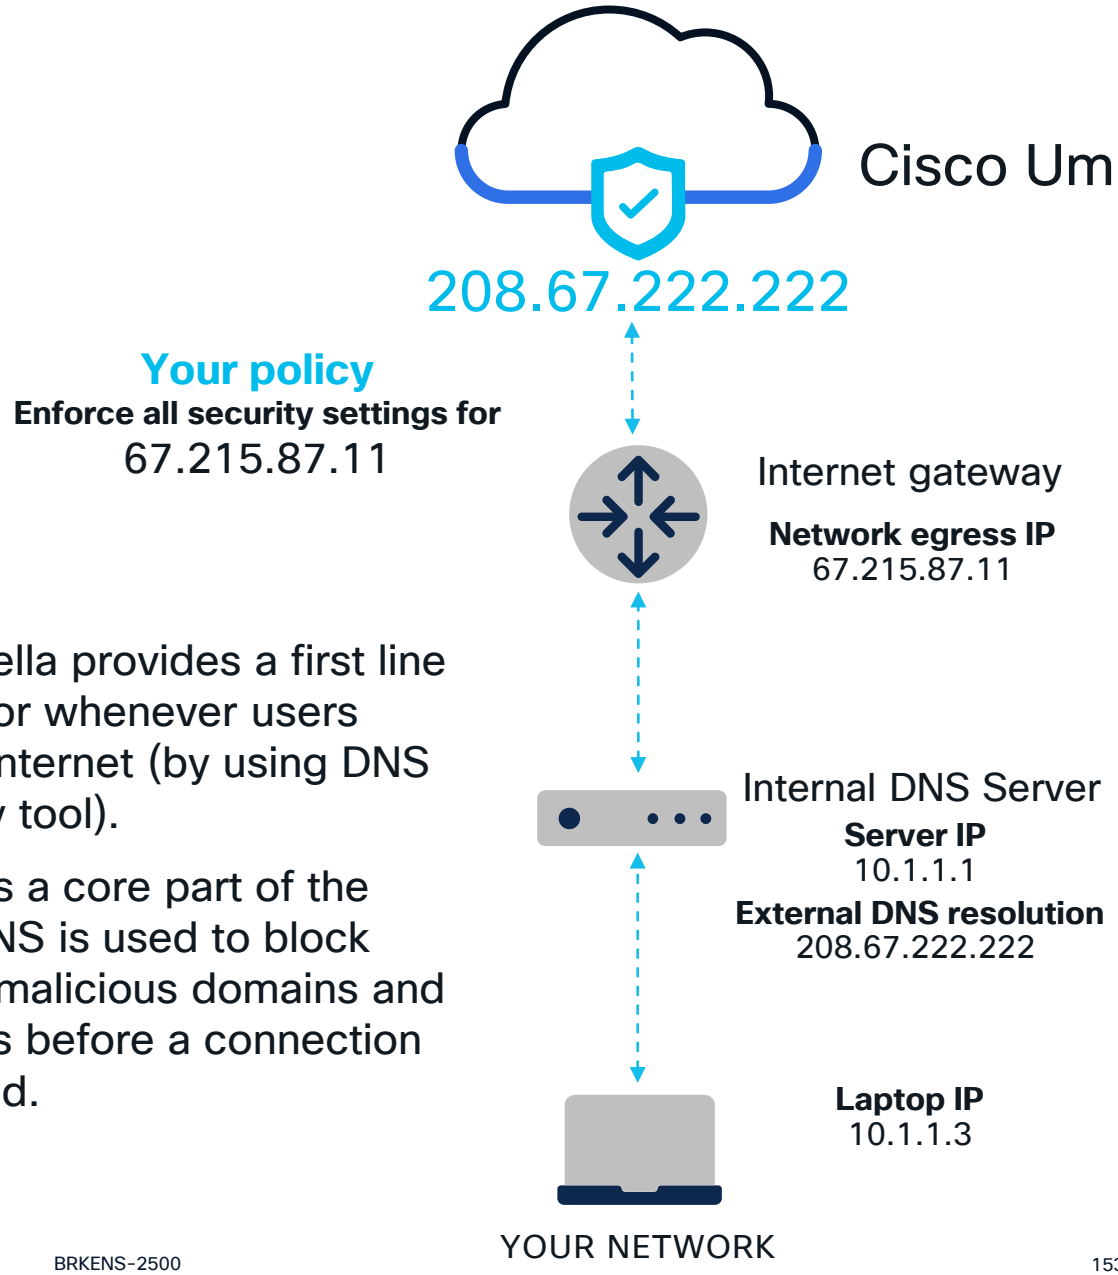

Switch Integrated Security Features

IPv4 & IPv6 Device Tracking

Switch Integrated Security Features (SISF) is a framework developed to optimize security in Layer 2 domains.

SISF merges IP Device Tracking (IPDT) and IPv6 first-hop security (FHS) functionality[1], to simplify the migration from IPv4 to IPv6 stack or a dual-stack.

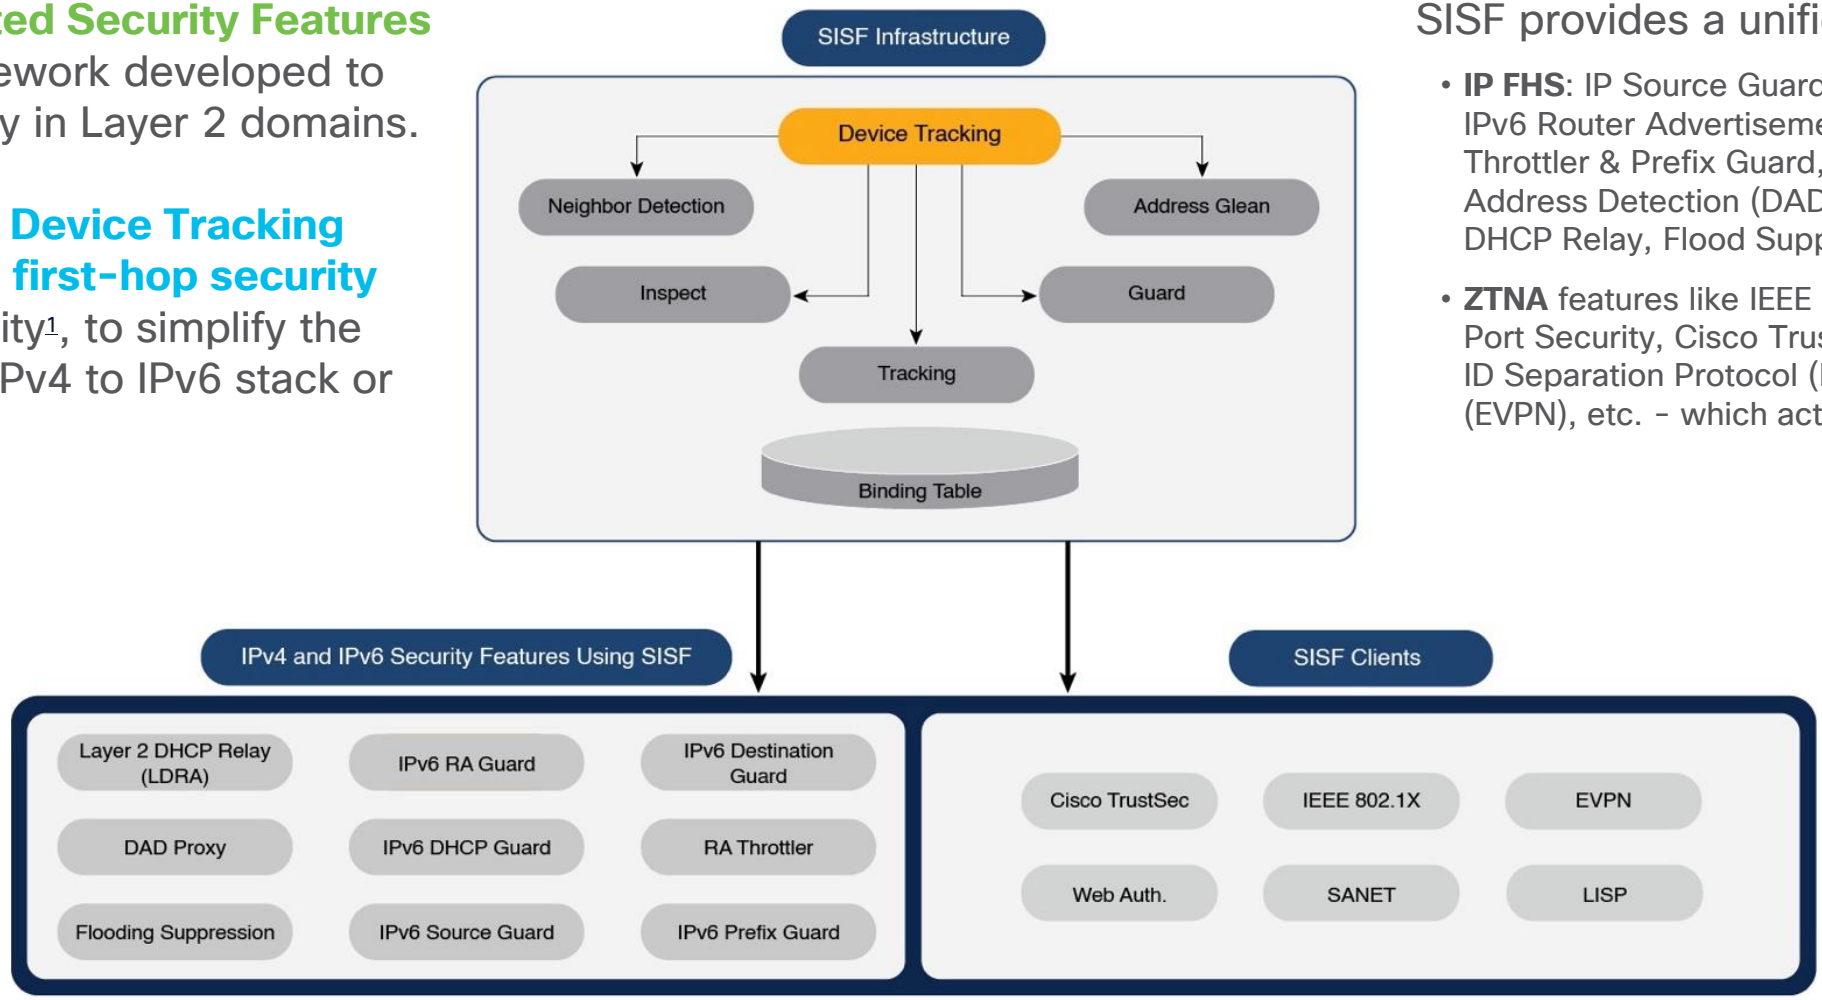

SISF provides a unified IP database:

- **IP FHS**: IP Source Guard, Destination Guard, IPv6 Router Advertisement (RA) Guard, RA Throttler & Prefix Guard, IPv6 Duplicate Address Detection (DAD), DHCP Guard, DHCP Relay, Flood Suppression, etc.
- **ZTNA** features like IEEE 802.1x, WebAuth, Port Security, Cisco TrustSec (CTS), Locator ID Separation Protocol (LISP), Ethernet VPN (EVPN), etc. - which act as clients of SISF.

SISF IP Binding Table

IP First-Hop Security (FHS)

Switch	IPv4	MAC	IF	VLAN	STATE
Acc-A	H_1	MAC_{H1}	P_1	100	REACH
Acc-A	H_2	MAC_{H2}	P_2	100	REACH
Acc-B	H_3	MAC_{H3}	P_1	200	REACH
Acc-B	H_4	MAC_{H4}	P_2	200	REACH
Distro	$H_{1,2}$	$MAC_{H1,2}$	P_A	100	REACH
Distro	$H_{3,4}$	$MAC_{H3,4}$	P_B	200	REACH

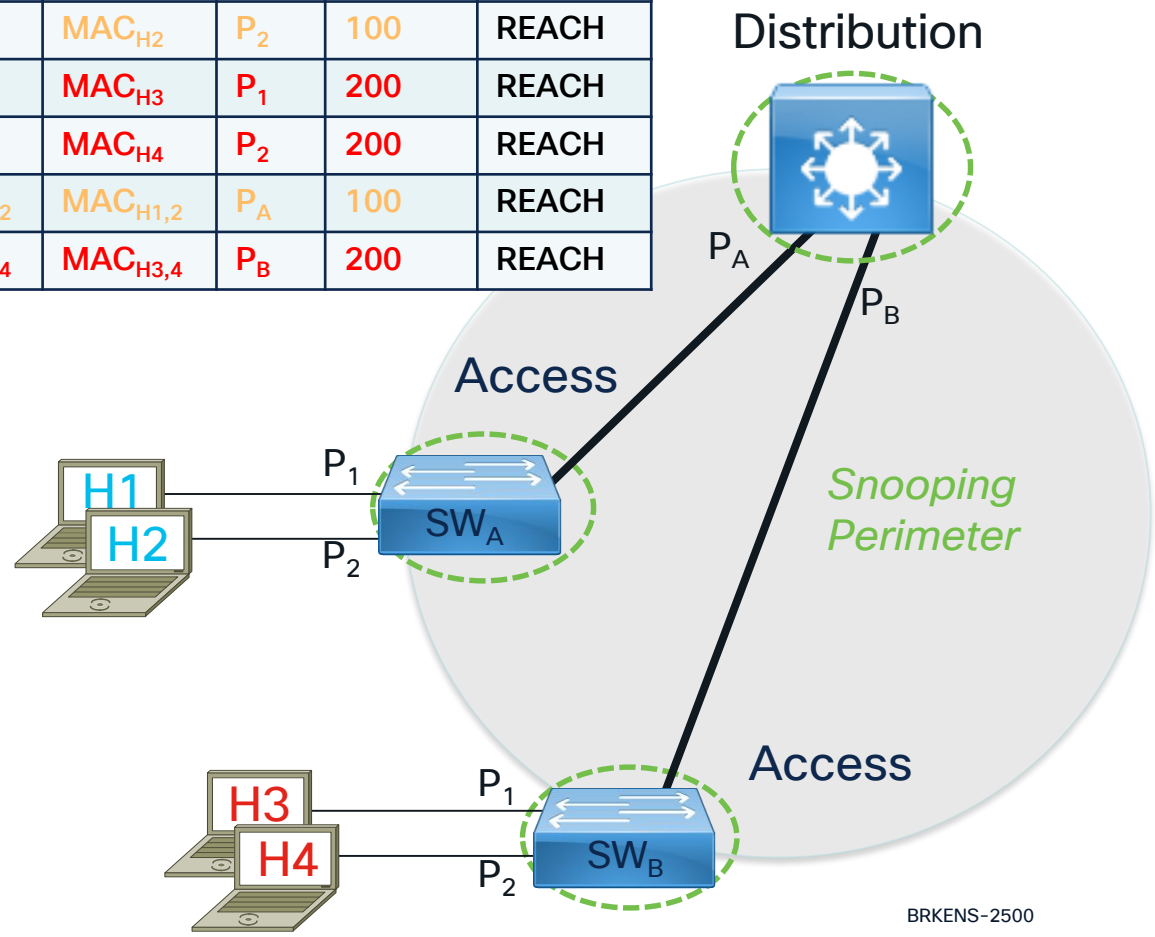

- A database of IP (v4/v6) neighbors connected to the L3 switch is created from various information sources, such as DHCP, ARP & NDP, etc.
- This database, or binding table is used by various L3 IP features (such as IP Source Guard) to validate the L2 link-layer address (LLA) and the IP address
- Also provide L2/L3 probes (e.g. Gratuitous ARP) to ensure reachability of the 'silent' devices
- Learns the L2/L3 bindings for stateless auto configuration (SLAAC) addresses in L2 tables.
- Former IP Device-Tracking (IPDT) CLI migrated to Switch Integrated Security (SISF) CLI
  - `device-tracking upgrade-cli`

Cisco TrustSec (CTS)

Simplifies Segmentation and Access Control

Traditional Access Control

- Static ACL
- Routing
- DHCP Scope
- Address
- VLAN

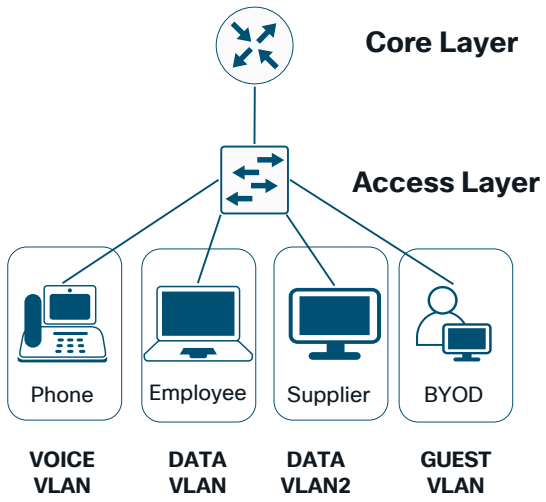

Security Policy based on addresses

High cost and complex maintenance

Cisco TrustSec

- Micro/Macro Segmentation
- Central Policy Provisioning
- No Topology Change
- No Vlan Change

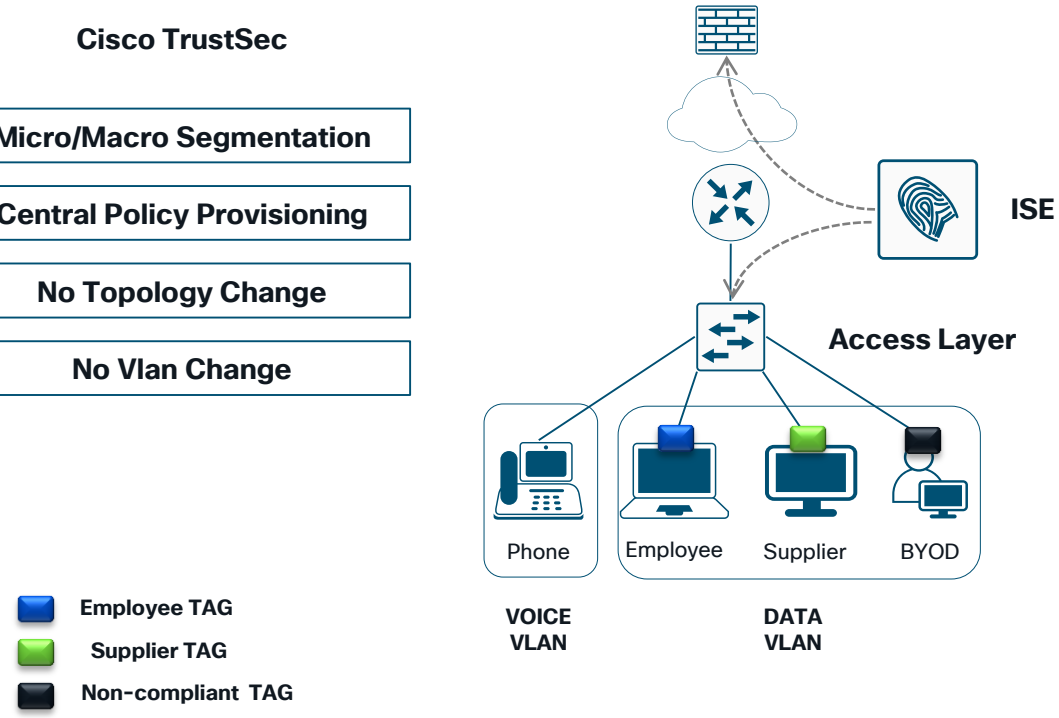

Use existing topology and

Automate security policy to reduce OpEx

The value of TrustSec

Policy Objective: “Allow employees web access to production servers”

Which policy is easier?

```
Switch-1# show ip access-list CorPolicy
Extended IP access list CorpPolicy
    10 permit tcp 10.1.100.0 0.0.0.255 172.16.100.0 0.0.0.255 eq 80
    20 permit tcp 10.1.100.0 0.0.0.255 172.16.100.0 0.0.0.255 eq 443
    30 permit tcp 10.1.101.0 0.0.0.255 172.16.100.0 0.0.0.255 eq 80
    40 permit tcp 10.1.101.0 0.0.0.255 172.16.100.0 0.0.0.255 eq 443
    50 permit tcp 10.1.100.0 0.0.0.255 172.16.101.0 0.0.0.255 eq 80
    …
```

Traditional ACL Policy

```
Switch-1# show cts role-based permissions
IPv4 Role-based permissions default:
        Deny IP-00
IPv4 Role-based permissions from group 10:Employee_SGT to group 100:ProdServer_SGT:
        Web_Only-10
```

TrustSec SGACL Policy

MACSEC

Layer2 P2P Encryption

- Higher Speed compared to IPSEC. MACSEC can reach line-rate interface speeds
- Encryption done at the physical layer (L2, MAC) of the ethernet port
- MACSEC encrypts Layer 2 Frame
- It is a hop-to-hop protocol
- MACSEC = 802.1AE Standard

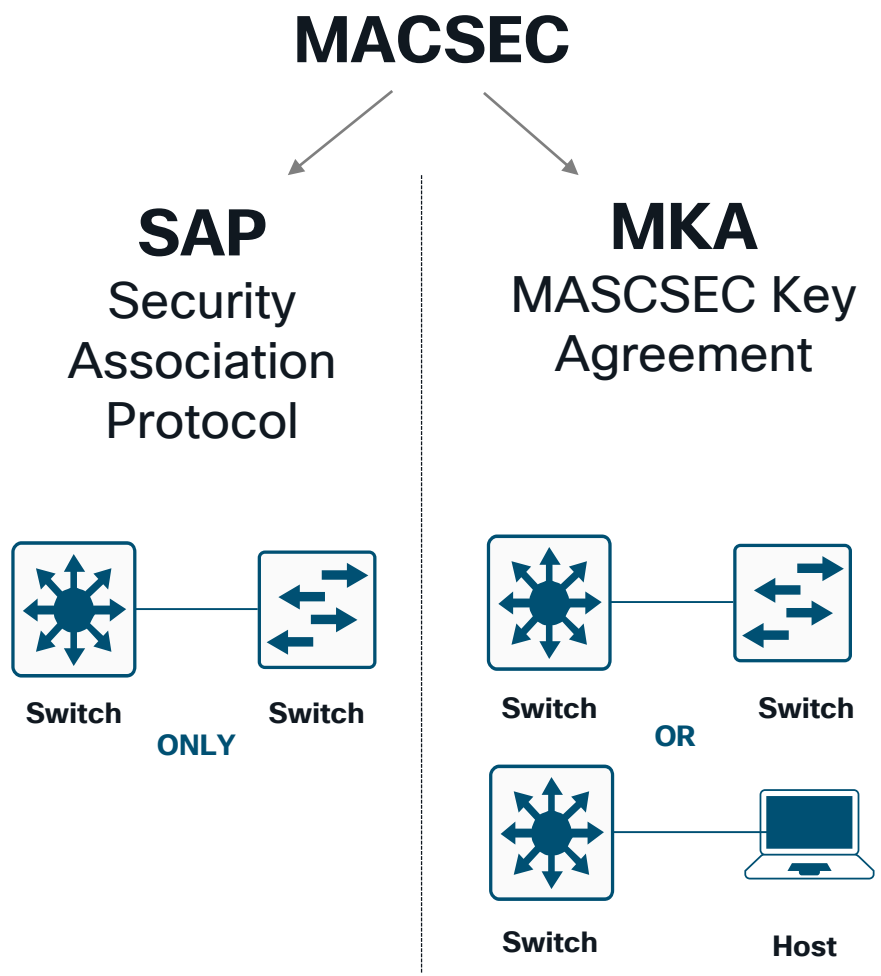

Control Plane Policing

- CoPP increases security on the switch by protecting the CPU from unnecessary Control-plane or DDOS traffic
- ON by default on Catalyst 9000 switches
- The best when implemented in hardware
- Catalyst 9000 platforms CoPP values adjusted to CPU of the model

```
!
control-plane
 service-policy input system-cpp-policy
!
```

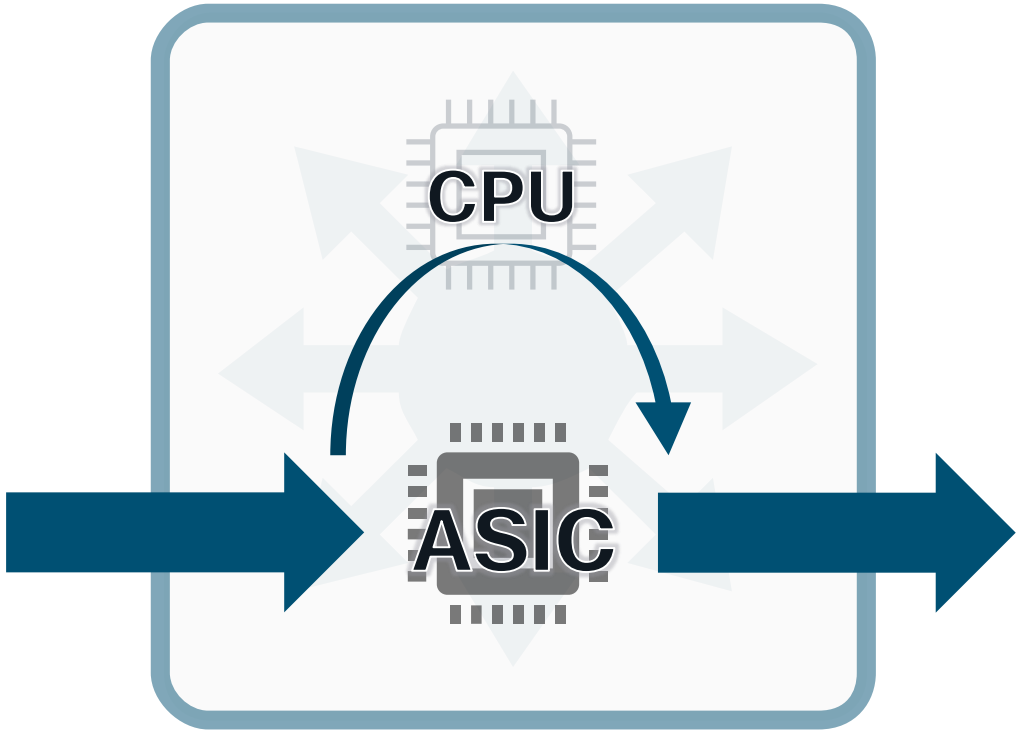

Catalyst 9000 Switching Security

BRKENS-2602

End-to-End Security Strategy for Enterprise Campuses with Catalyst 9000 Series Switches

Ama Owusu-Hammond - Technical Marketing, Cisco

This session is focused primarily on the Cisco Catalyst 9000 Series Switches, which can provide end-to-end security from campus and branch to cloud.

The session also covers secure infrastructure with Cisco Trustworthy Solutions, secure transport with MACsec and IPsec (site-to-site, site-to-cloud), secure endpoints with native connectors, Cisco Secure Network Analytics (Stealthwatch), Cisco Umbrella, auto-profile and secure endpoints, and using endpoint analytics and trust analytics. After this session, you will be able to take away how Catalyst 9000 Series Switches are built with security in mind for fulfilling various use cases.

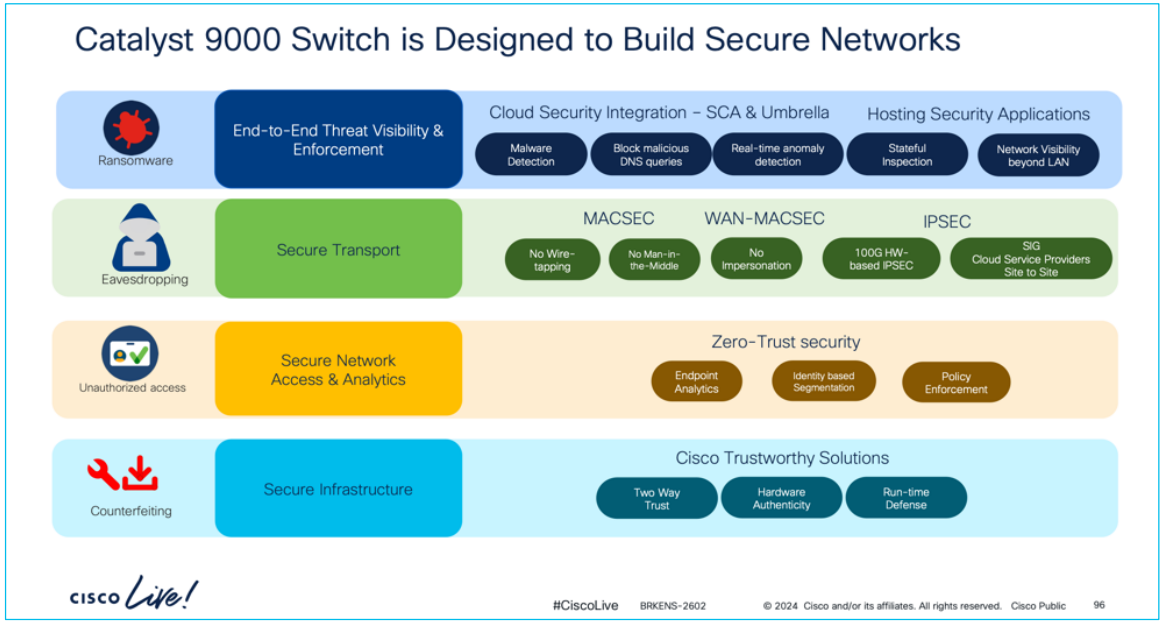

Platform Design

- LAN High Availability
- LAN Security
- Virtual Networking
  - Value of Overlays
  - MPLS VPN
  - SDA LISP
  - SDA EVPN

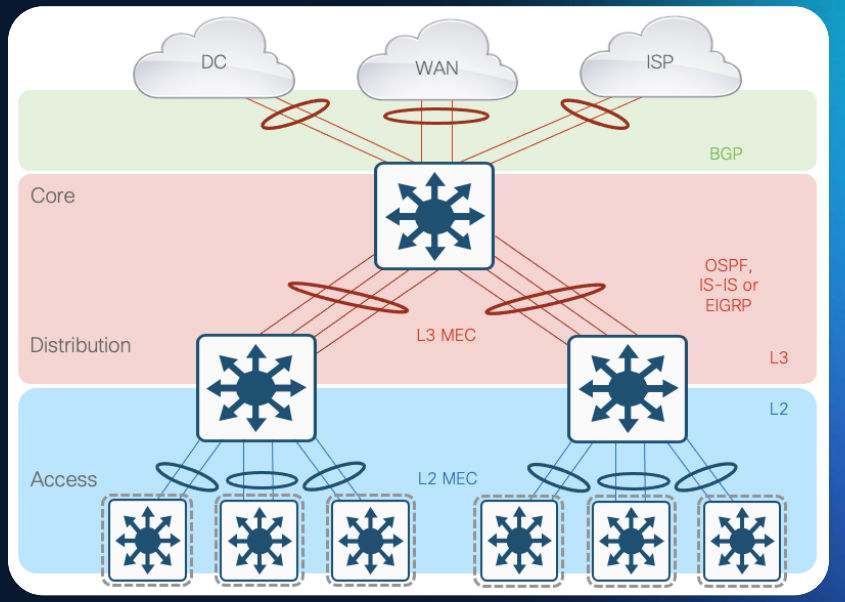

Design Options

How about something else? Fabric Overlay Networks!

Live! BRKENS-2501
BRKENS-2502

- Fabric Overlay enables:
  - ☑ Host Mobility
  - ☑ Layer 2 & Layer 3
  - ☑ Network Segmentation
  - ☑ Role-based Access Control

- It is a **virtual overlay** on top of the **physical underlay**
  - Control plane based on LISP or EVPN
  - Data plane based on VxLAN-GPO
  - Policy plane based on Cisco TrustSec

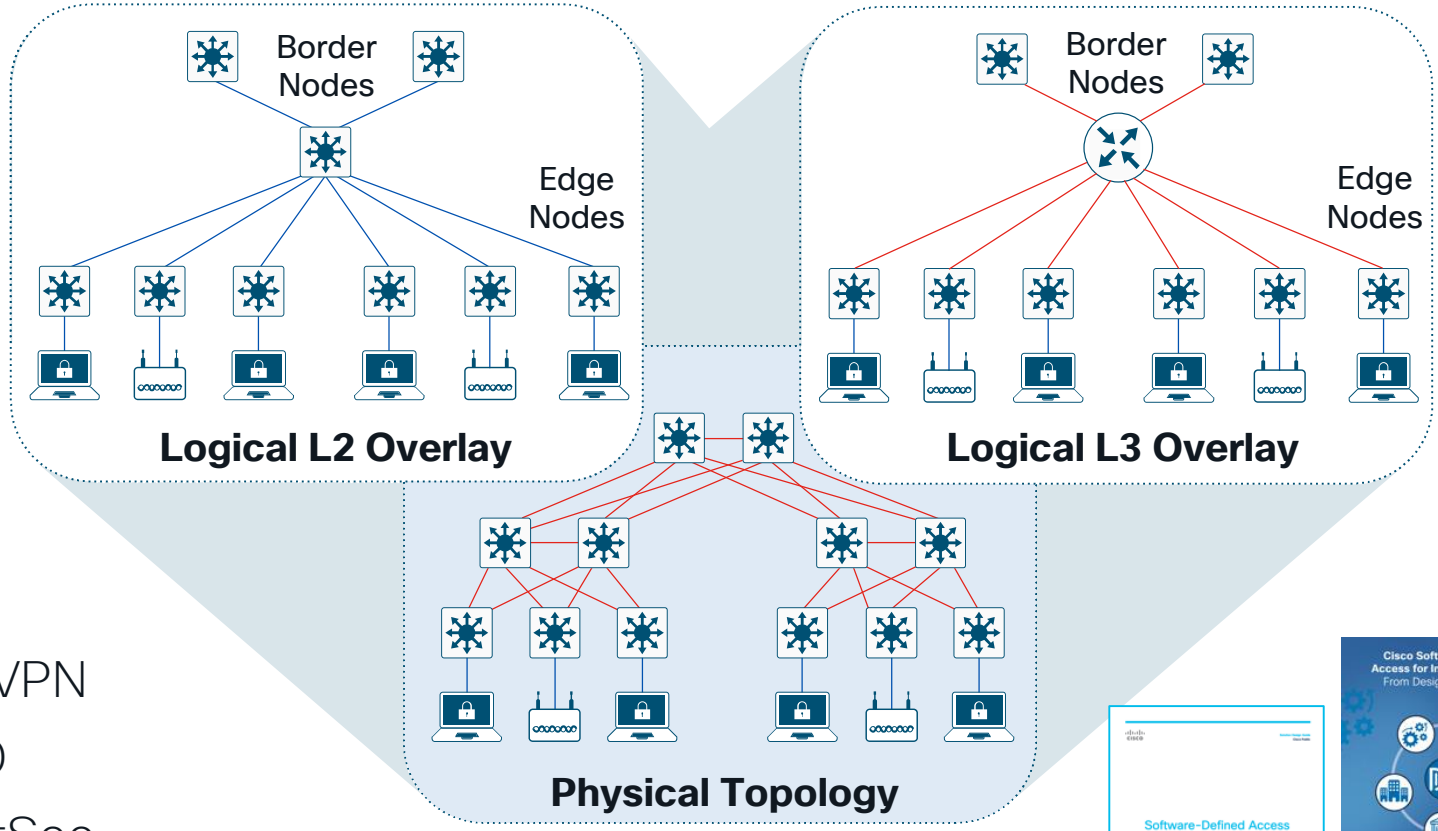

Design Options

Cisco Software-Defined Access

BRKENS-2501
BRKENS-2502

Logical topology—
L2/L3: flexible overlays

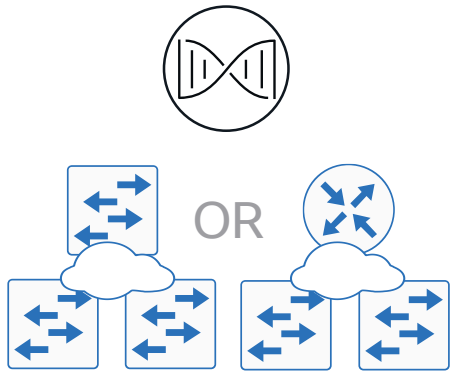

Physical topology:
2 core
2 dist./acc.

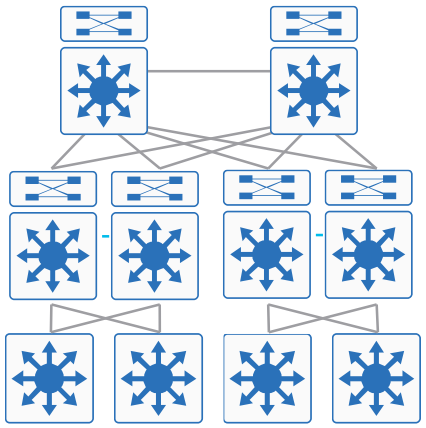

- Uses advantages of a routed access physical design, with Layer 2 capable logical overlay design
- Provisioning and policy automation
- Integrates wireless into the same policy
- Requires automation to simplify configuration

Survives device and link failures	✔
Easy mitigation of Layer 2 looping concerns	✔
Rapid detection/recovery from failures	✔
Layer 2 across all access blocks within distribution	✔
Device-level CLI configuration simplicity	
Automated network and policy provisioning included	✔

A Fabric is an Overlay

A logical mesh topology used to virtually connect devices

Live! BRKENS-2501
BRKENS-2502

- Built on top of physical underlay topology – using encapsulation
- Provides additional L2/L3 services not provided by the underlay

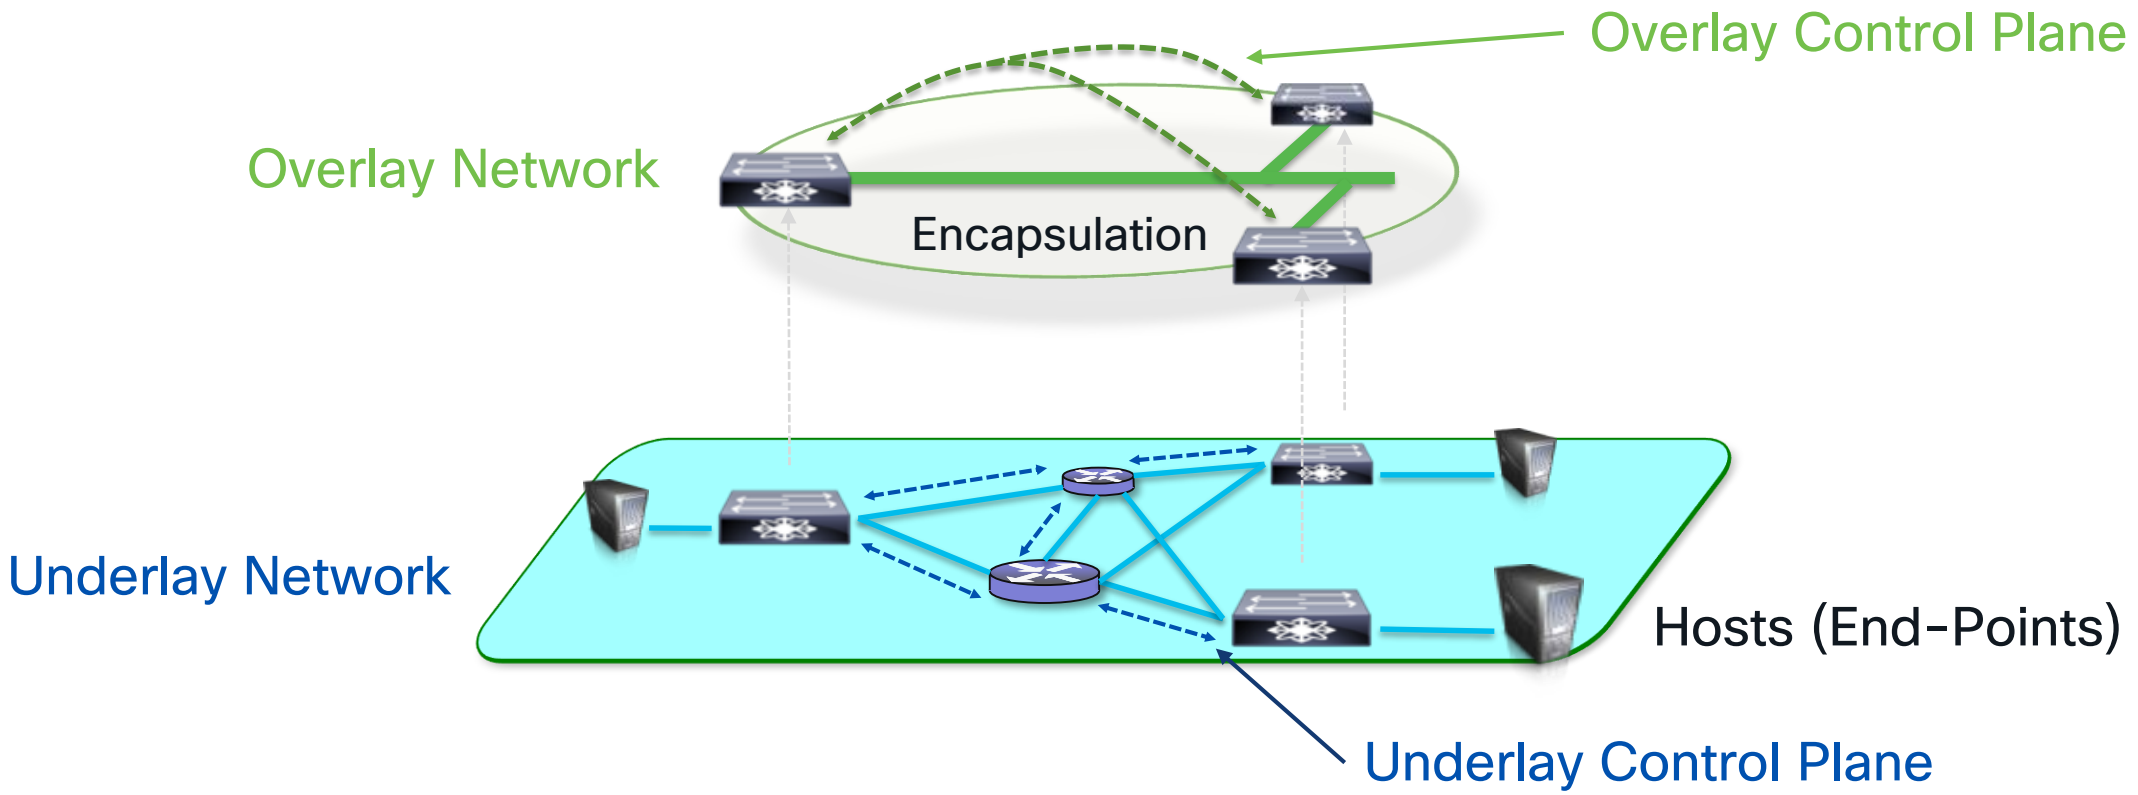

Fabric Solves Network Problems

Solve all your L2/L3 and Policy challenges with a single design

Live! BRKENS-2501
BRKENS-2502

- ☑ Fabric to the Access removes L2 protocols (STP, DTP, VTP)
- ☑ Reliable Layer 2 extension - over simple Layer 3 network
- ☑ Simple, scalable, reliable –and- automatable
- ☑ Consistent Access layer configurations
- ☑ Convergence of Wired and Wireless

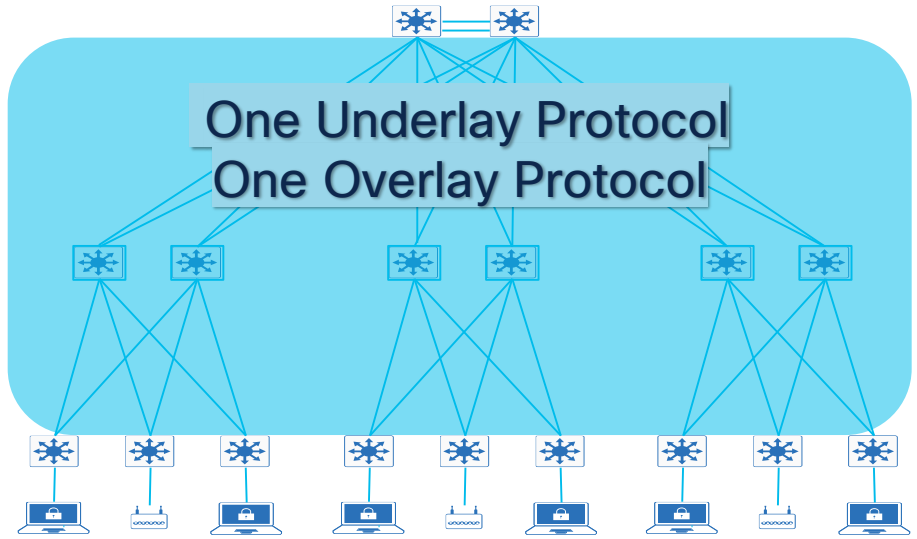

Campus Solutions & Designs

Providing design options (beyond basic PINs)

1 MPLS (L2/L3VPN)

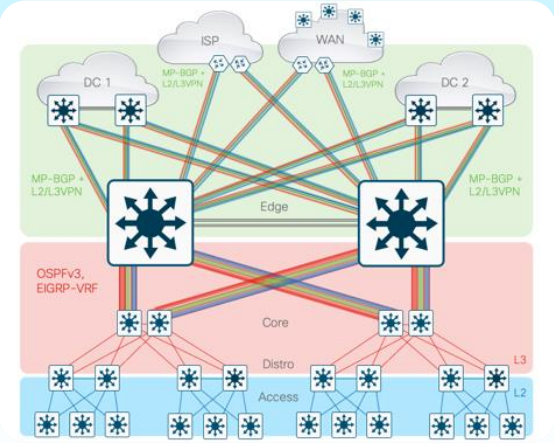

- L3 Underlay + L2/L3 VPN Overlay
- Virtual Private Networks
- L3 VRF-based Segmentation
- WAN/Edge + **VPN Services**

MPLS/VPLS, **LDP**, SR, MP-BGP, PIC
MVPN, LSM, Extranet, MSR
SSO, NSF/NSR, ECMP, GIR
VPN-FNF, Uniform/Pipe QoS, PBR, IPACL

2 EVPN (L2/L3VNI)

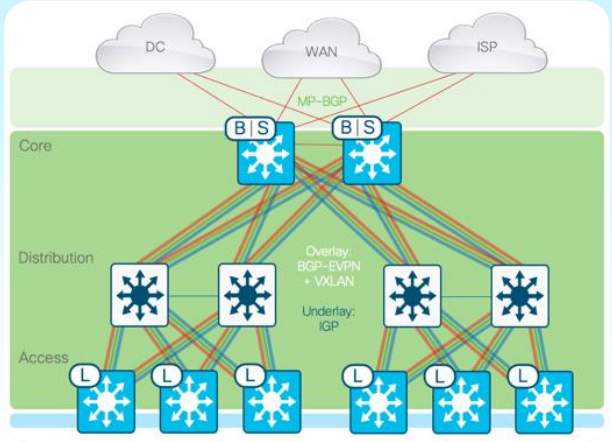

- **L3 Underlay + L2/L3 VNI Overlay**
- Virtual Network Instances
- L2/L3 VNI-based Segments
- Common WAN/LAN Services

MP-BGP/EVPN, **VXLAN**, VRF-Lite
L2 TRM, L3 TRM, L2 BUM
SSO, NSF/NSR, ECMP, GIR
Fabric-FNF, Uniform QoS, IPACL/OGACL

3 LISP (L2/L3VNI + SGT)

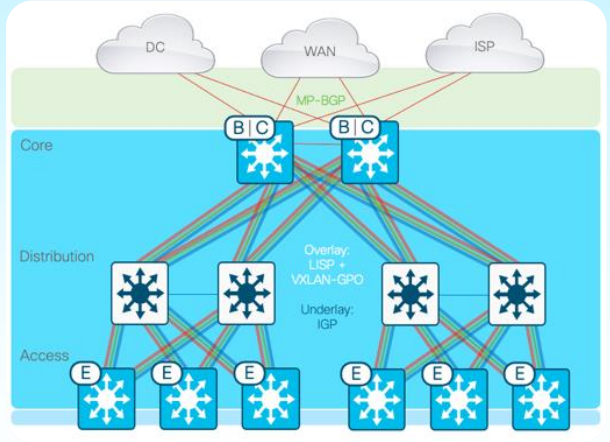

- **L3 Underlay + L2/L3 VNI Overlay**
- VNIs + Scalable Group Tagging
- L2/L3 VNI + SGT Segments
- LAN Services + Group-Based Policy

LISP, **VXLAN-GPO**, MP-BGP, VRF-Lite
LISP HER, Native, L2 BUM
SSO, NSF/NSR, ECMP, GIR
Fabric-FNF, App QoS, SGACL

MPLS-VPN Provider Edge

The Provider-Edge PIN (Tier 3-4) focuses on connecting multiple Campus areas to remote domains (SP/WAN) using MPLS-VPN.

Main goal is to connect EVPN fabric to other networks

Uses a **L3 Underlay + L3 Hand-off**

- North (outside): L3 MP-BGP + Inter-AS, PIM + MSDP
- South (inside): L3 IGP, PIM + MSDP

Uses a **Virtualized L2/L3 Overlay**

- Control-Plane: **MPLS, EoMPLS/VPLS, MVPN**
- Data-Plane: **LDP, mLDP**
- Policy-Plane: **VPN ID**

Tends to use **Overlay-aware Features**

- IP or OG ACLs (e.g. destined Outside)
- Uniform/Pipe QoS (e.g. separate Inner vs. Outer)
- Inter-VRF Routing (e.g. VRF-Lite, Leaking)
- MPLS-aware NetFlow (e.g. VPN ID in FNF)

May require **multiple encapsulation(s)**

Tends to require **high L2/L3 & feature scale**

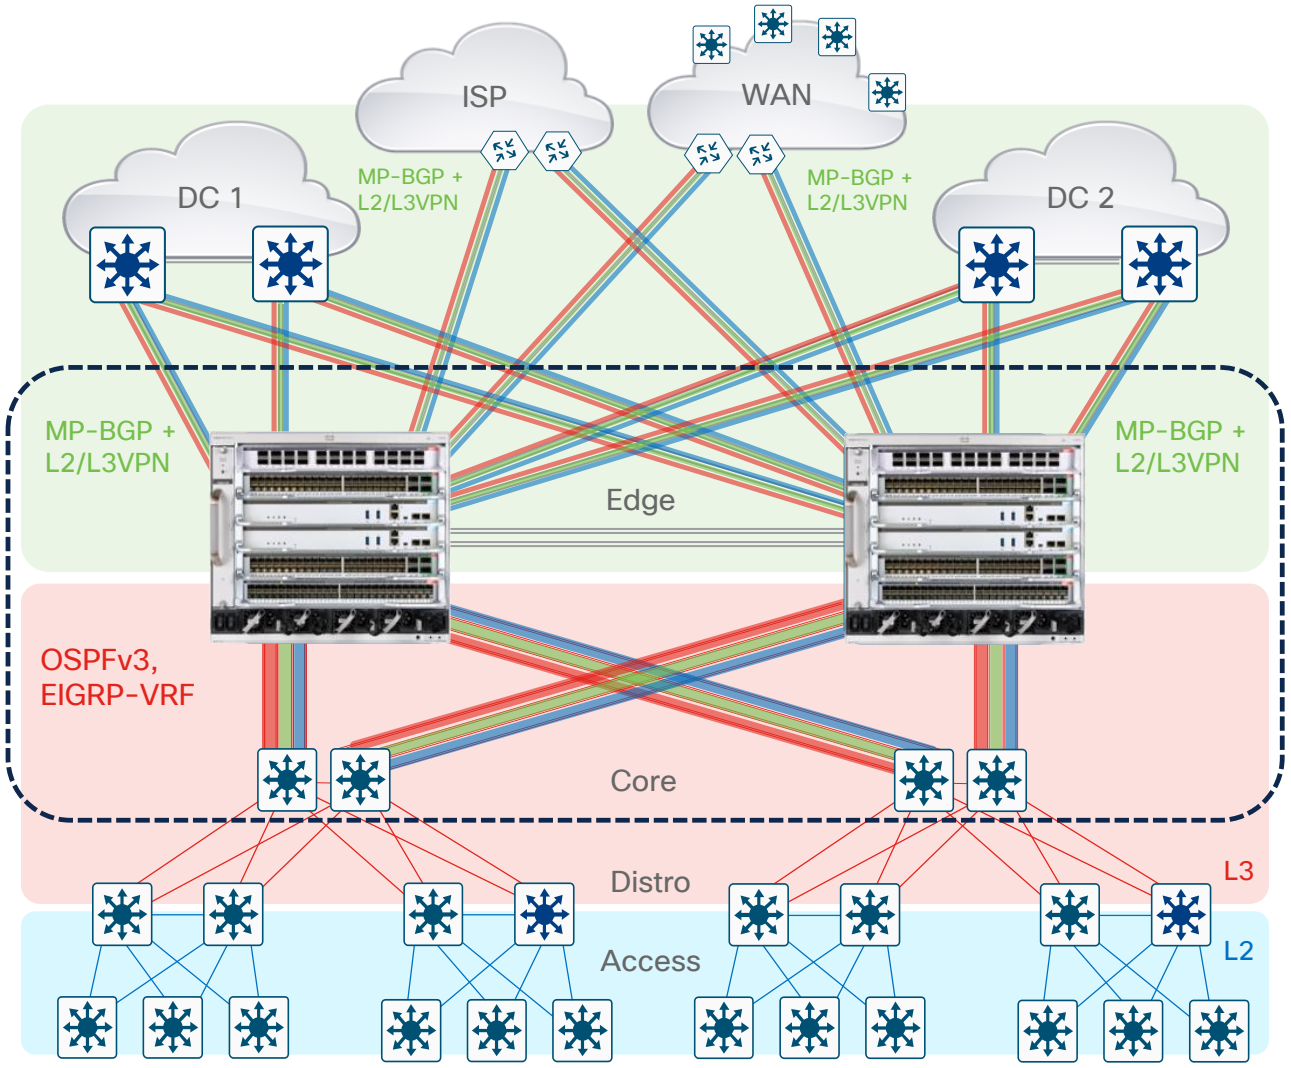

EVPN Border & Spine

The **EVPN Border & Spine PIN** focuses on connecting an **EVPN Fabric** and/or **other network domains**.

- Typically, the same layer as Core or Edge (Tier 3-4)

Main goal is to connect EVPN fabric to other networks

Uses a **L3 Underlay + L3 Hand-off**

- North (outside): L3 MP-BGP + Inter-AS, PIM + MSDP
- South (inside): L3 IGP, PIM + MSDP

Uses a **Virtualized L2/L3 Overlay**

- Control-Plane: **BGP-EVPN (RR), TRM**
- Data-Plane: **VXLAN**
- Policy-Plane: **L2/L3 VNID**

Tends to use **Overlay-aware Features**

- IP/OG ACLs (e.g. destined Outside)
- Uniform QoS (e.g. copy Inner, queue Outer)
- Inter-VRF Routing (e.g. VRF-Lite, Leaking)
- Fabric NetFlow (e.g. VRF/VNID in FNF)

May require **multiple encapsulation(s)**

Tends to require **high L2/L3 & feature scale**

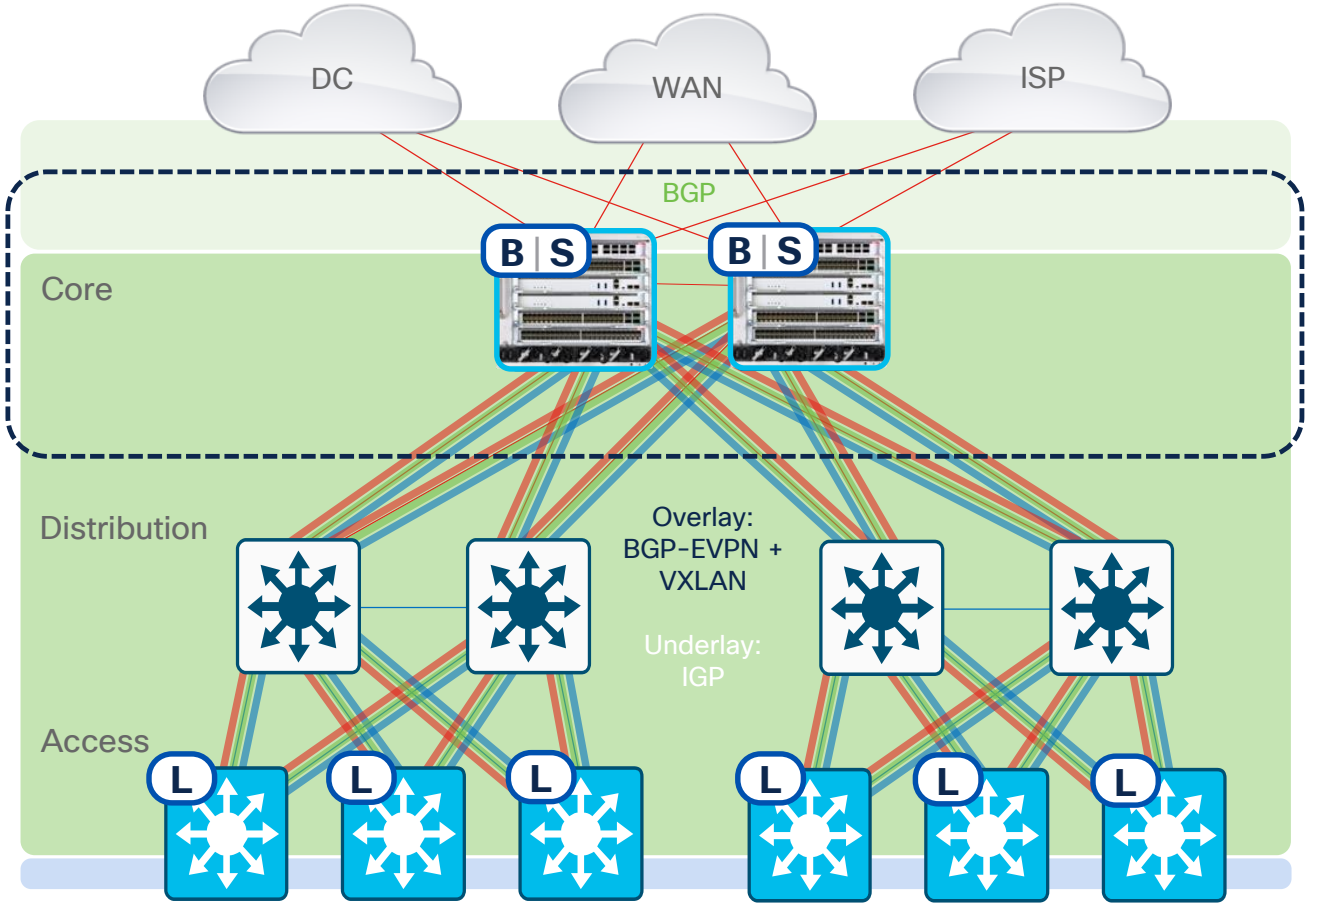

EVPN Leaf

The **EVPN Leaf PIN** focuses on connecting Wired endpoints to an **EVPN Fabric domain**.

- Typically, the same layer as Access or Extended (Tier 1)

Main goal is to connect Endpoints to EVPN network

Uses a **L3 Underlay + L2 Hand-off**

- North (inside): L3 IGP, PIM + MSDP
- South (outside): L2 VLAN (L3 SVI), STP, IGMP

Uses a **Virtualized L2/L3 Overlay**

- Control-Plane: **BGP-EVPN**, **TRM**
- Data-Plane: **VXLAN**
- Policy-Plane: **L2/L3 VNI**

Tends to use **Overlay-aware features**

- IP/OG ACLs (e.g. destined outside)
- Uniform QoS (e.g. copy inner, queue outer)
- Inter-VRF Routing (e.g. VRF Leaking)
- Fabric NetFlow (e.g. FNF + VNID)

Tends to require **med-high L2/L3 & feature scale**

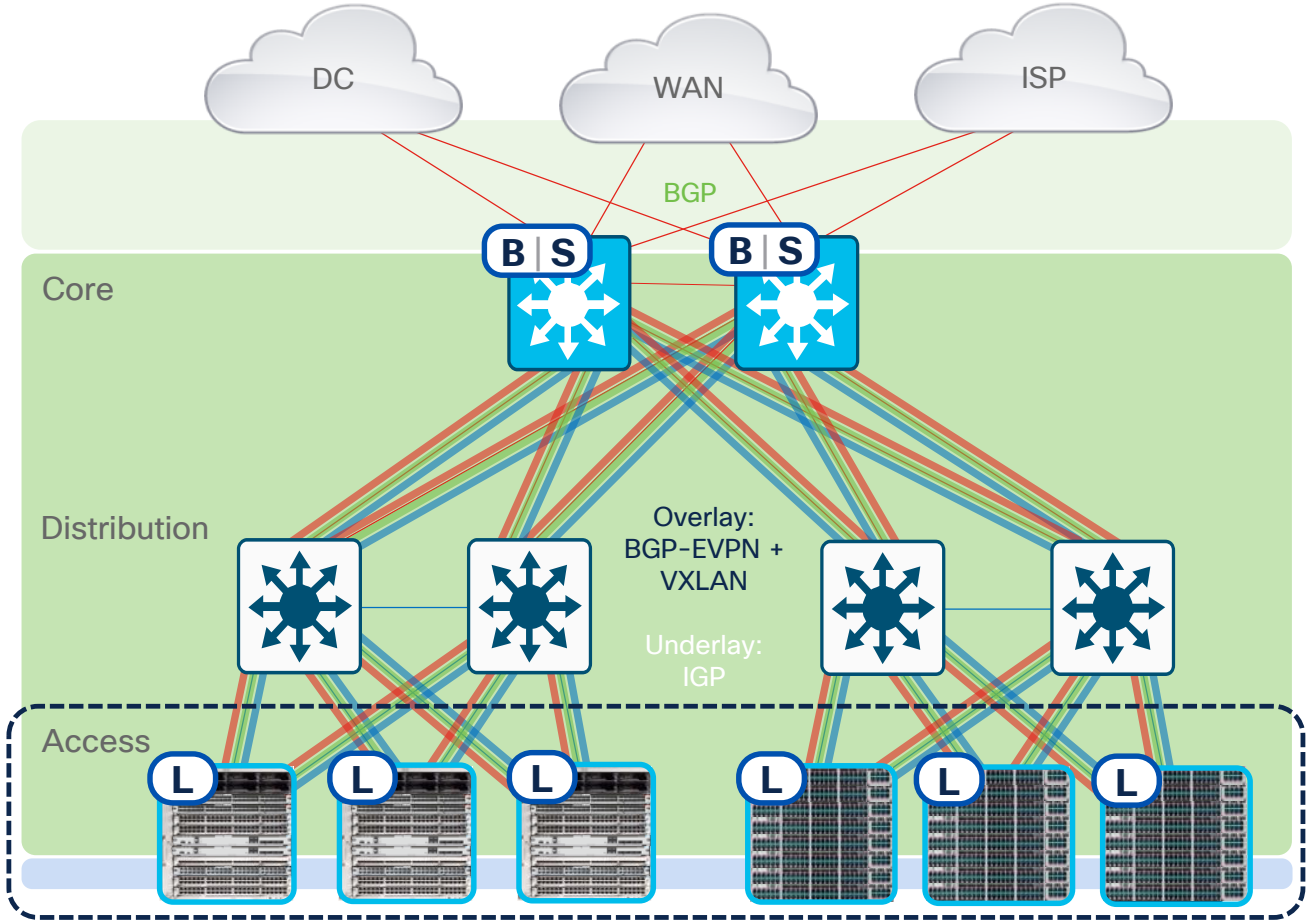

SD-Access Border & CP

The **SDA Border & CP PIN** focuses on connecting an **SDA Fabric** and/or **other network domains**.

- Typically, the same layer as Core or Core/Edge (Tier 3-4)

Main goal is to connect SDA fabric to other networks

Uses a **L3 Underlay + L3 Hand-off**

- North (outside): L3 MP-BGP + Inter-AS, PIM + MSDP
- South (inside): L3 IGP, PIM + MSDP

Uses a **Virtualized L2/L3 Overlay**

- Control-Plane: **LISP (XTR, MS/MR), PIM**
- Data-Plane: **VXLAN-GPO**
- Policy-Plane: **L2/L3 VNI + SGT**

Tends to use **Overlay-aware features**

- Security Group ACLs (e.g. destined outside)
- Uniform Pipe QoS (e.g. copy inner, queue outer)
- Inter-VRF Routing (e.g. VN Extranet, or VRF-Lite)
- Fabric NetFlow (e.g. VRF/VNID + SGT FNF, NaaS/ETA)

May require **multiple encapsulation(s)**

Tends to require **higher L3 & feature scale**

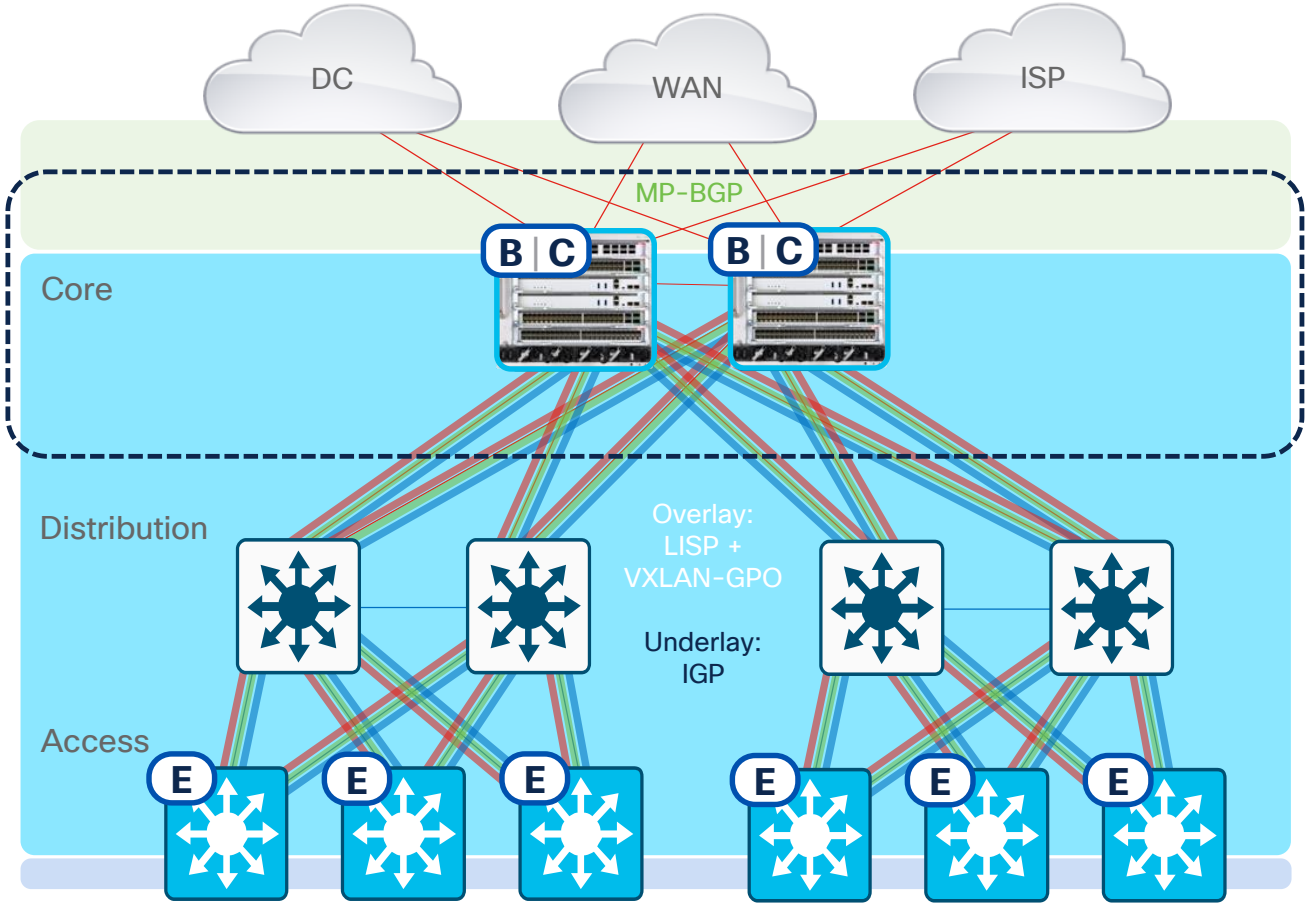

SD-Access Edge

The **SDA Edge PIN** focuses on connecting Wired/Wireless endpoints to an **SDA Fabric domain**.

- Typically, the same layer as Access or Extended (Tier 1)

Main goal is to connect Endpoints to SDA network

Uses a **L3 Underlay + L2 Hand-off**

- North (inside): L3 IGP, PIM + MSDP
- South (outside): L2 VLAN (L3 SVI), STP, IGMP

Uses a **Virtualized L2/L3 Overlay**

- Control-Plane: **LISP (XTR), PIM**
- Data-Plane: **VXLAN-GPO**
- Policy-Plane: **VN + SGT**

Tends to use **Overlay-aware features**

- Security Group ACLs (e.g. destined outside)
- Uniform Pipe QoS (e.g. copy inner, queue outer)
- Inter-VRF Routing (e.g. VN Extranet)
- Fabric NetFlow (e.g. FNF, NaaS)

Tends to require **higher L3 & feature scale**

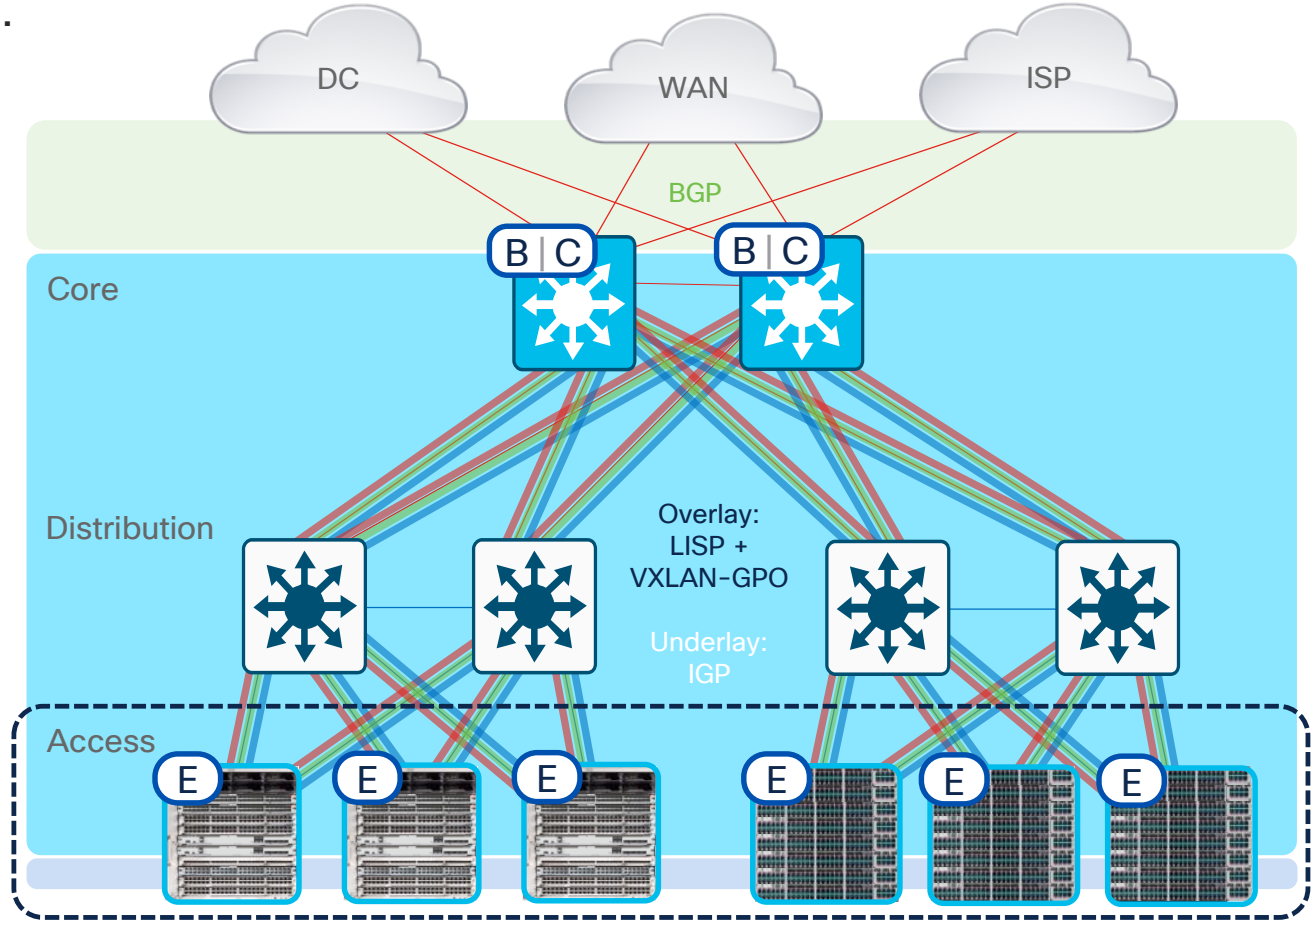

MTU Consideration

MTU is the Maximum Transmit Unit a device can forward.

- In general, this “Transmission Unit" is the IP packet Length – including the IP Header.
- L2 headers e.g., 802.1Q tag, MACsec, SVL, etc, are not accounted in this calculation
- System MTU –vs– Port MTU –vs– IP MTU
  - System MTU – System MTU is a global configuration, which sets the MTU of the whole device
  - Per-port MTU – Per-port MTU allows setting an MTU value on a per-interface basis – and takes precedence over the system MTU. If the per-port setting is removed, interface will fall back to the system MTU.
  - IP MTU – is only applicable to IP packets. Other non-ip packet sizes will not be accounted for using this command.

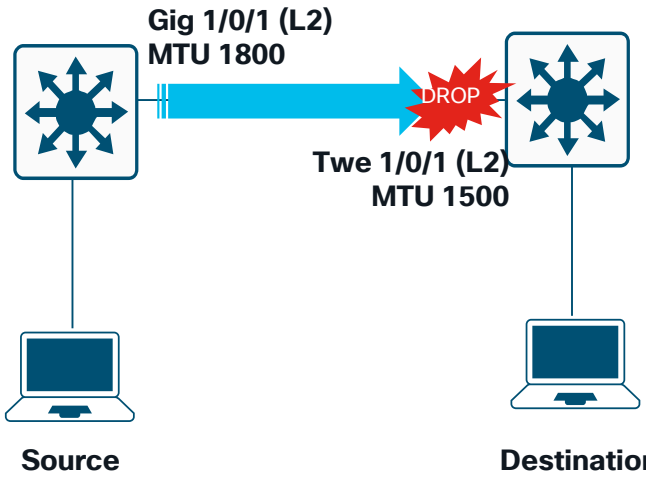

Catalyst 9000 switches handle packet sizes from **64** bytes to **9238** bytes

Catalyst 9000 Overlay Fabrics

BRKENS-2501

Overlay Design Options for Campus Networks

Raj Kumar Goli - Technical Marketing, Cisco

This presentation will delve into the various overlay design options available with the state-of-the-art Catalyst 9000 Switching Platforms. We'll start by defining the concept of network overlay and its critical role in modern network architecture, especially in the era of cloud computing and virtualization.

The focus will then shift to the Catalyst 9000 series, exploring how these switches leverage advanced technologies to support multiple overlay design options. This includes discussion on technologies like Virtual Extensible LAN (VXLAN) and Software-Defined Access (SD-Access), which are integral for creating network overlays.

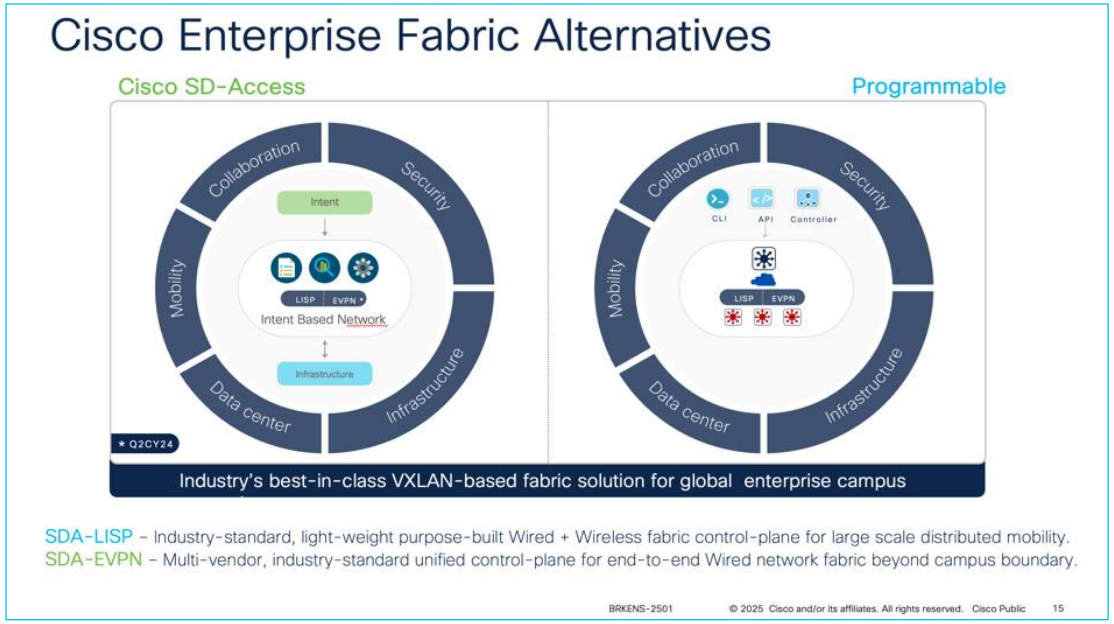

Catalyst 9000 SDA Design

BRKENS-2502

Cisco SD-Access LISP VXLAN Fabric Best Practices: Design and Deployment

Mahesh Nagireddy - Technical Marketing, Cisco

This session includes a brief introduction of Cisco SD-Access components, and dives into design and scale considerations and deployment options, for single-site designs covering greenfield and brownfield converged wired and wireless infrastructures.

Participants will gain insights into how Cisco SD-Access can provide a journey to digitalization and immediate benefits at every step of embracing the zero-trust architecture. This session will focus on multi-site design and deployment options, with the intent to provide end-to-end segmentation with consistent policy across the enterprise.

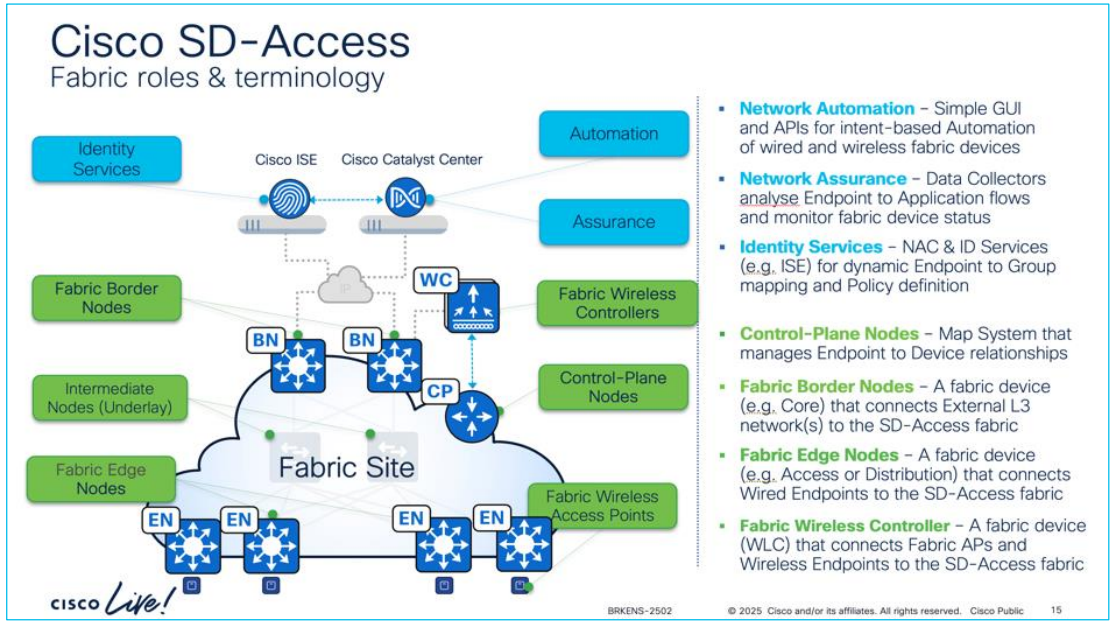

Session Agenda

Design Fundamentals | **Design Considerations**

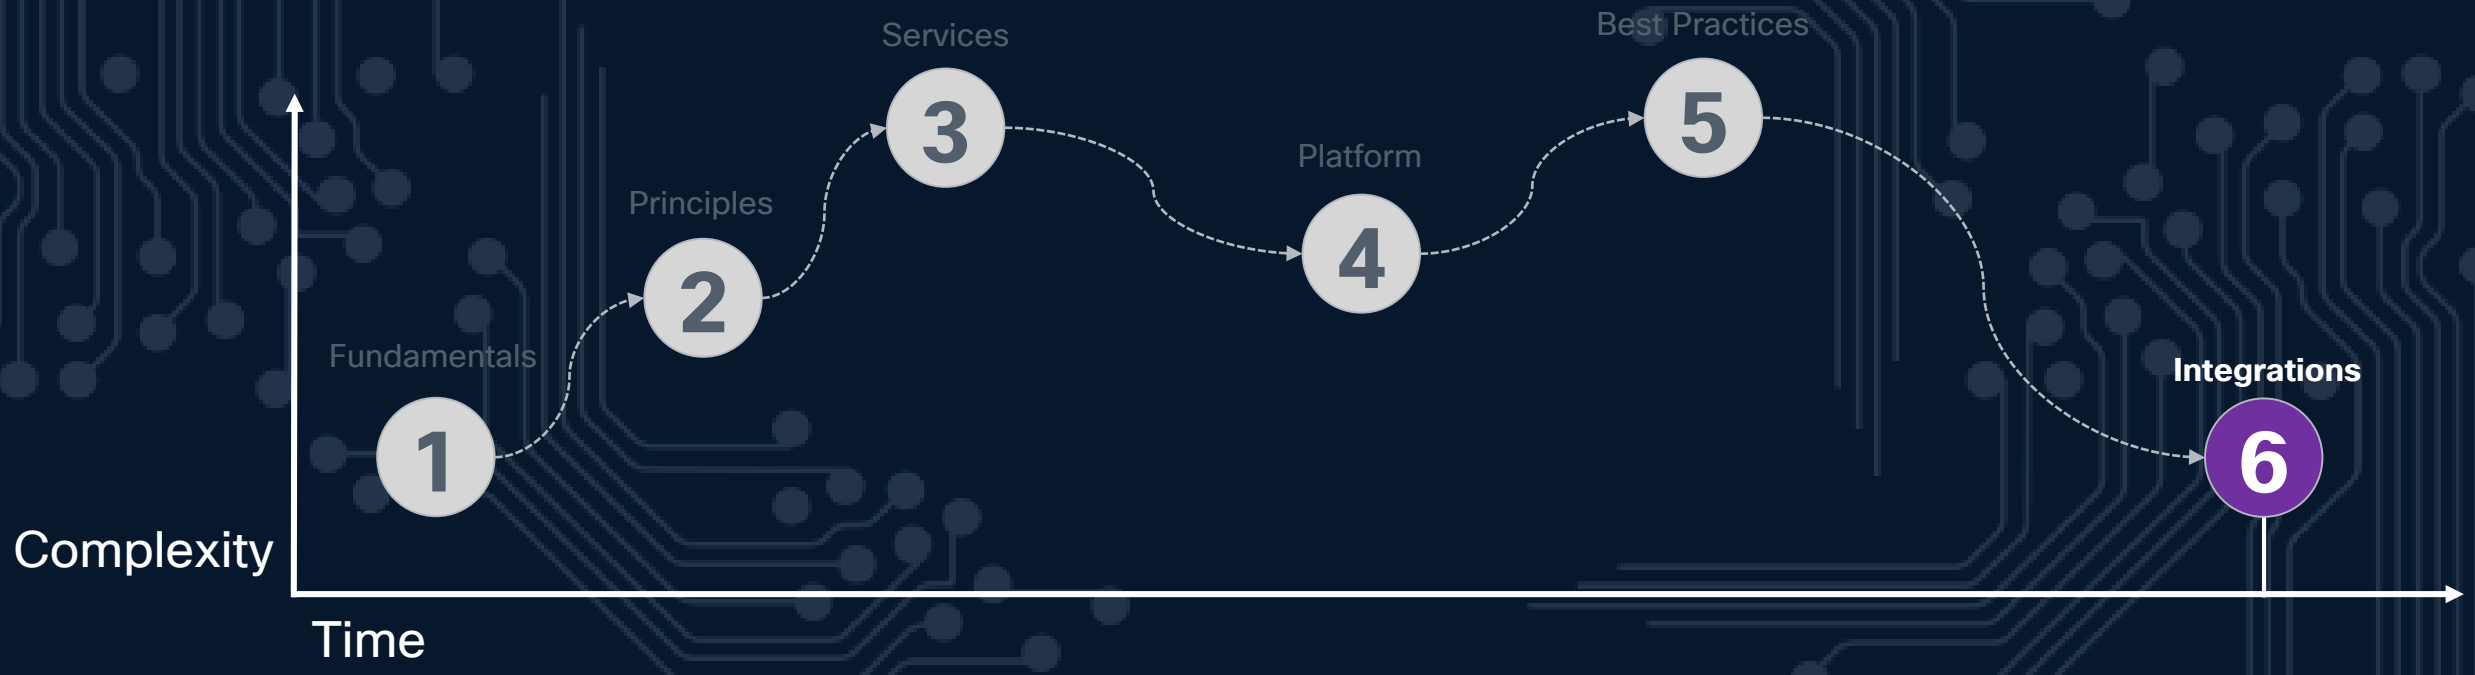

Wireless & Firewall

- **Central Wireless LAN**
- **Firewalls, VRFs & ACLs**

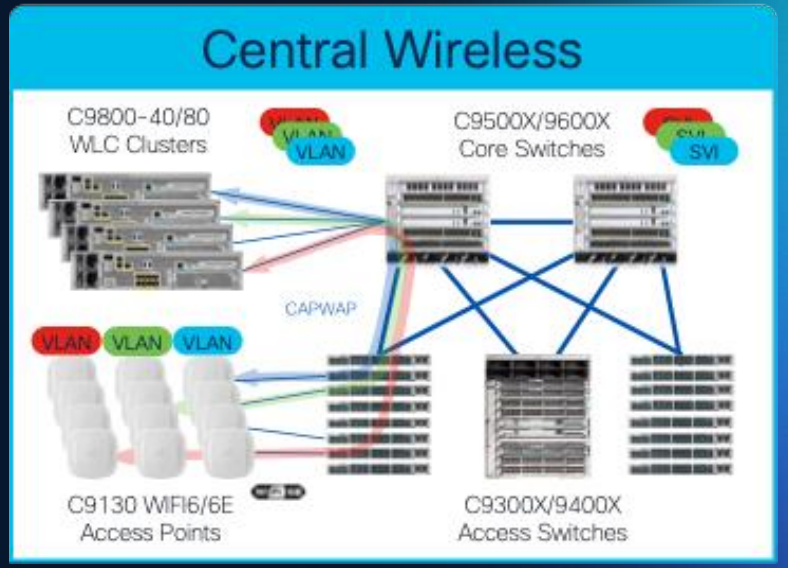

Wireless LAN

The **Central Wireless PIN** focuses on connecting Wireless APs centrally to one or multiple WLCs.

- WLC is typically connected to Core, Edge or DC (Tier 3+)
- APs are typically connected to Access (Tier 1)

Main goal is to connect Wireless Endpoints (via APs) to a Wireless LAN (WLAN) - centrally in the network

Uses a **L2/L3 Underlay + L2 Hand-off**

- North (to WLC): L2 VLAN + 802.1Q, L3 SVI, IGP
- South (to AP): L2 VLAN + 802.1Q, STP, IGMP

Uses a **Tunneled L2 Overlay**

- Control-Plane: **CAPWAP, DTLS, LWAPP**
- Data-Plane: **CAPWAP, DTLS**

Tends to require **L2 (WLAN) features**

- L2 ACLs (e.g. VACL, MAC ACL)
- L2 QoS (e.g. VLAN QoS)
- L2 NetFlow (e.g. FNF, AVC, EPA & ETA)

Tends to require **higher L2/L3 + feature scale**

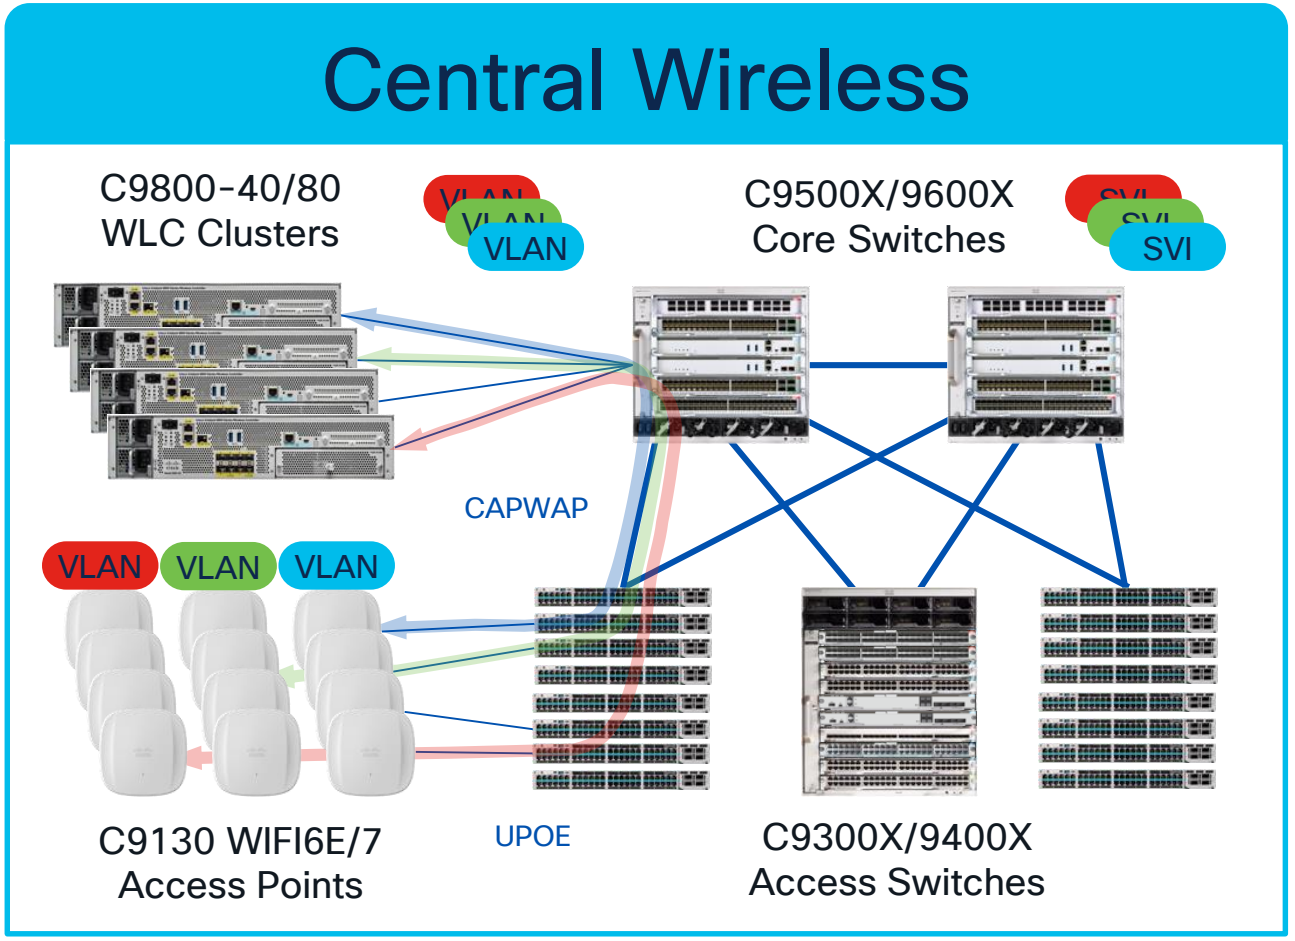

Wireless – Local Mode

Design Considerations

- Recommended primarily for large site deployments
- Simplified (central) configuration and troubleshooting, and roaming at scale
- The WLAN controller and APs are both located (local) at the same site
- If any of the following are true, you should consider deploying a controller locally at the site
  - The site has a **LAN distribution layer**
  - The site has **more than 100 APs**
  - The site has a **WAN latency ≥100ms** round-trip to the central WLC

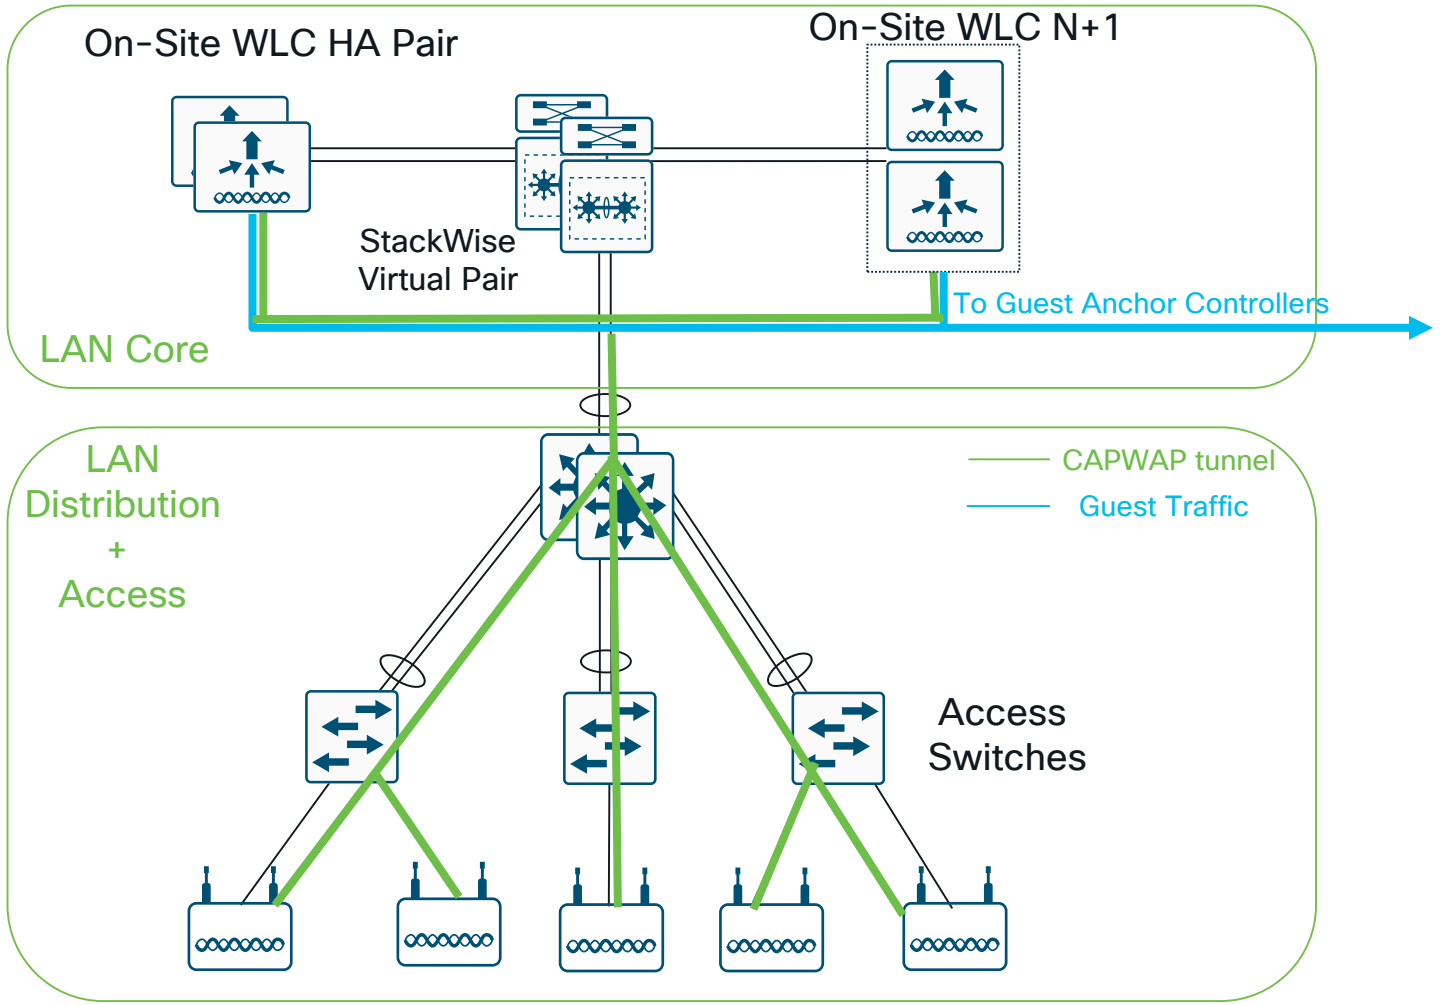

Wireless – FlexConnect Mode

Design Considerations

- Solution for deployments that consist of many small remote sites (branches) connected to a central site
- Enables control of remote-site APs from HQ, through the WAN, without deploying a controller in each site
- If all the following are true (per site), consider deploying FlexConnect:
  - Site has **fewer than 50 APs**.
  - Site LAN is a **single access-layer** (switch or stack)
    – connected to a central location
  - Site has a **WAN latency ≤ 100ms** round-trip to the central WLC

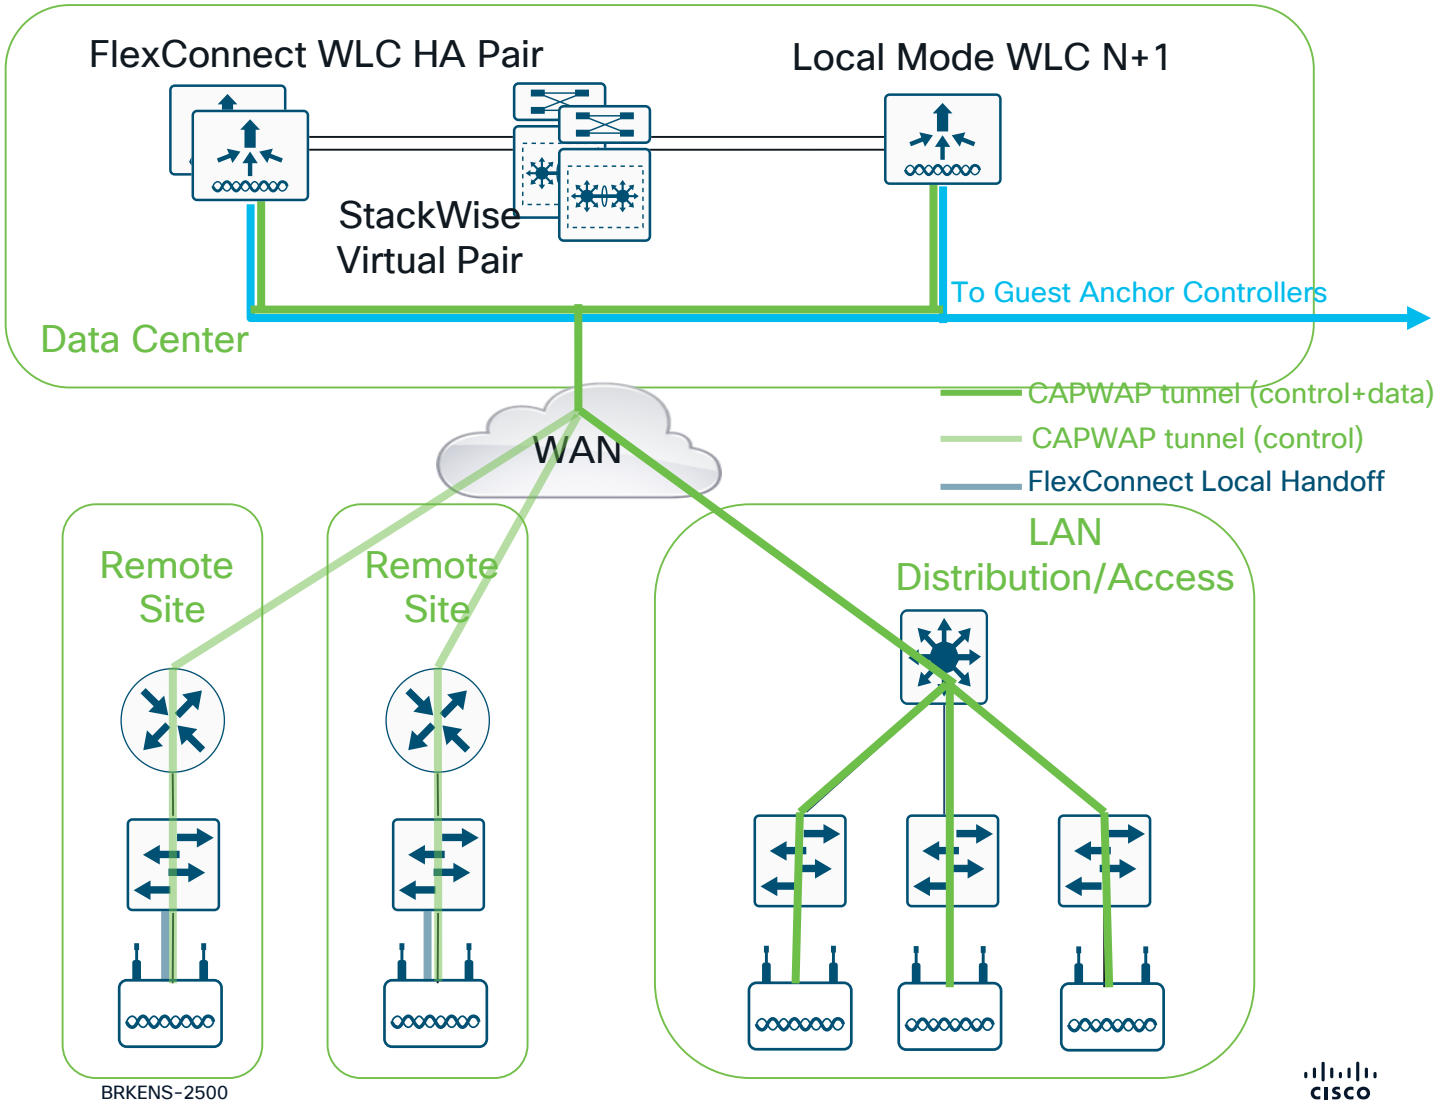

SD-Access Wireless Model

- The primary benefit of Cisco SD-Access Wireless for seamless L2 roaming and common policy across both wired and wireless
- SD-Access Wireless control plane traffic is managed by the WLC, while data plane traffic is passed directly by the fabric.
- Requirements for SD-Access Wireless:
  - Cisco **SD-Access wired fabric** deployment
  - **APs with fabric mode** support, directly connected to Cisco SD-Access **fabric edge nodes**
  - **WLC with fabric mode** control plane support
  - Fabric WLC with **≤ 20ms latency** to the fabric APs

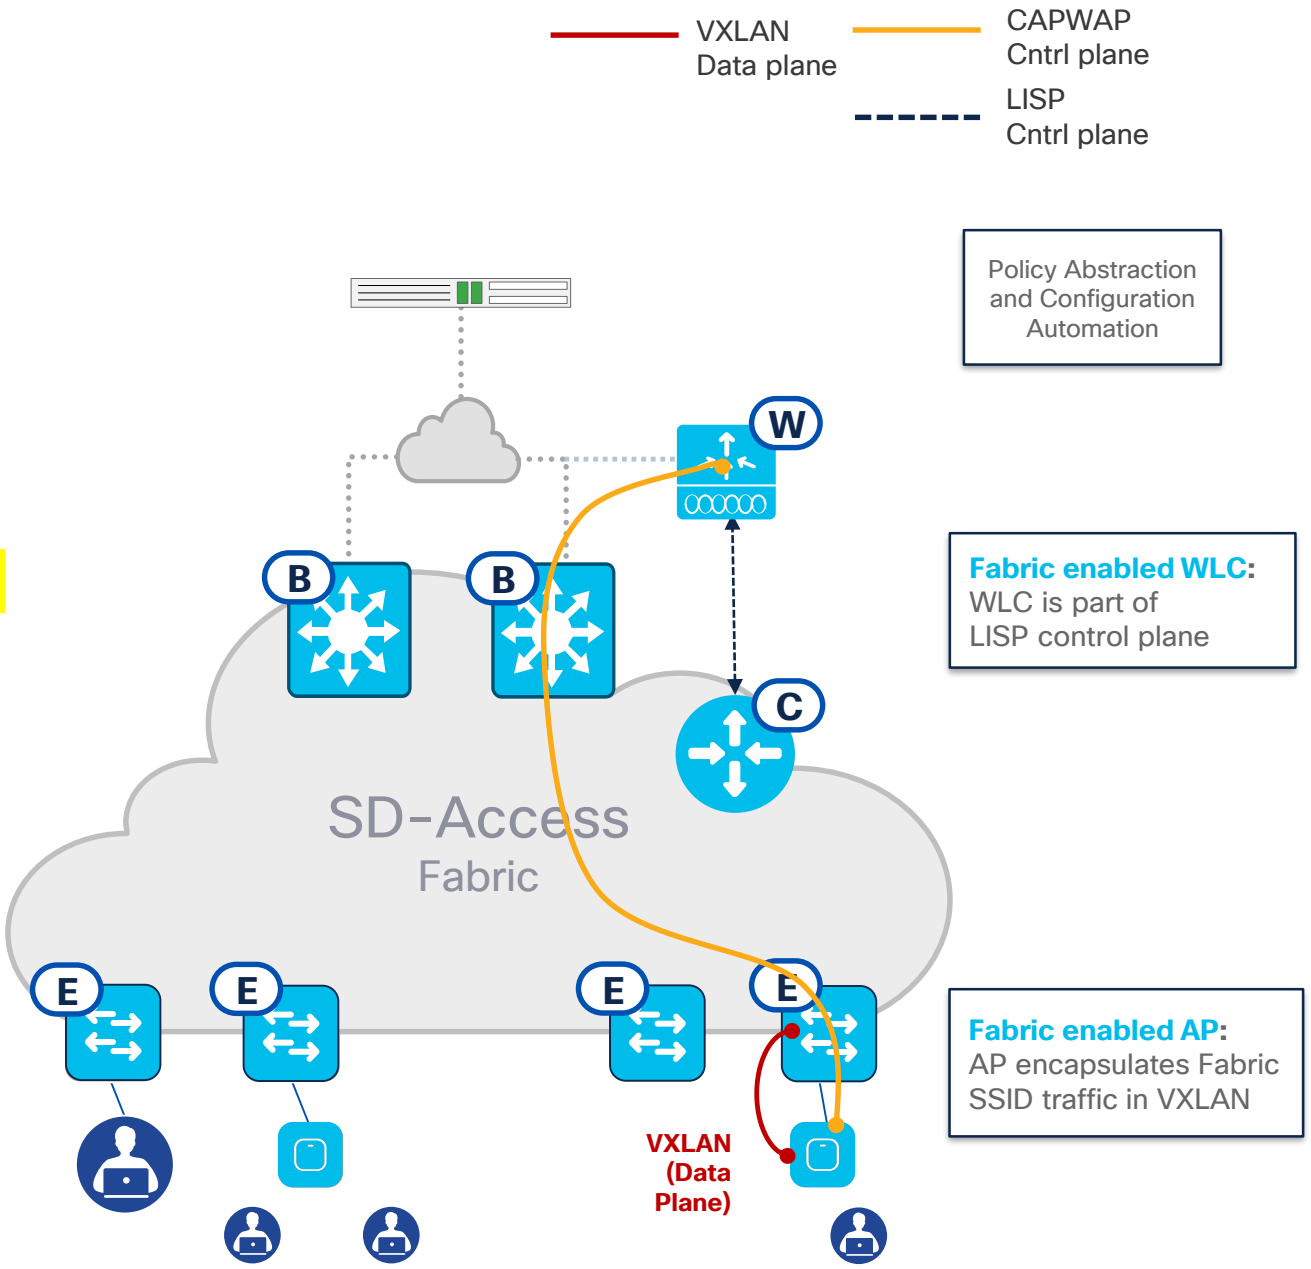

Wireless QoS

QoS is like a chain – It's only as strong as its weakest link

The WLAN is one of the weakest links in enterprise QoS designs

For three basic reasons:

1. Normal **downshift in input vs. output speed**
   - For example: From 1Gbps to ≤ 500Mbps
2. Shift from **full-duplex vs. half-duplex media**
   - Cable is TX & RX vs RF is TX, then RX
3. Shift from **dedicated media to shared media**
   - Switching + P2P Cable vs. Radio Broadcast

WLAN QoS policies control both wireless jitter and packet loss

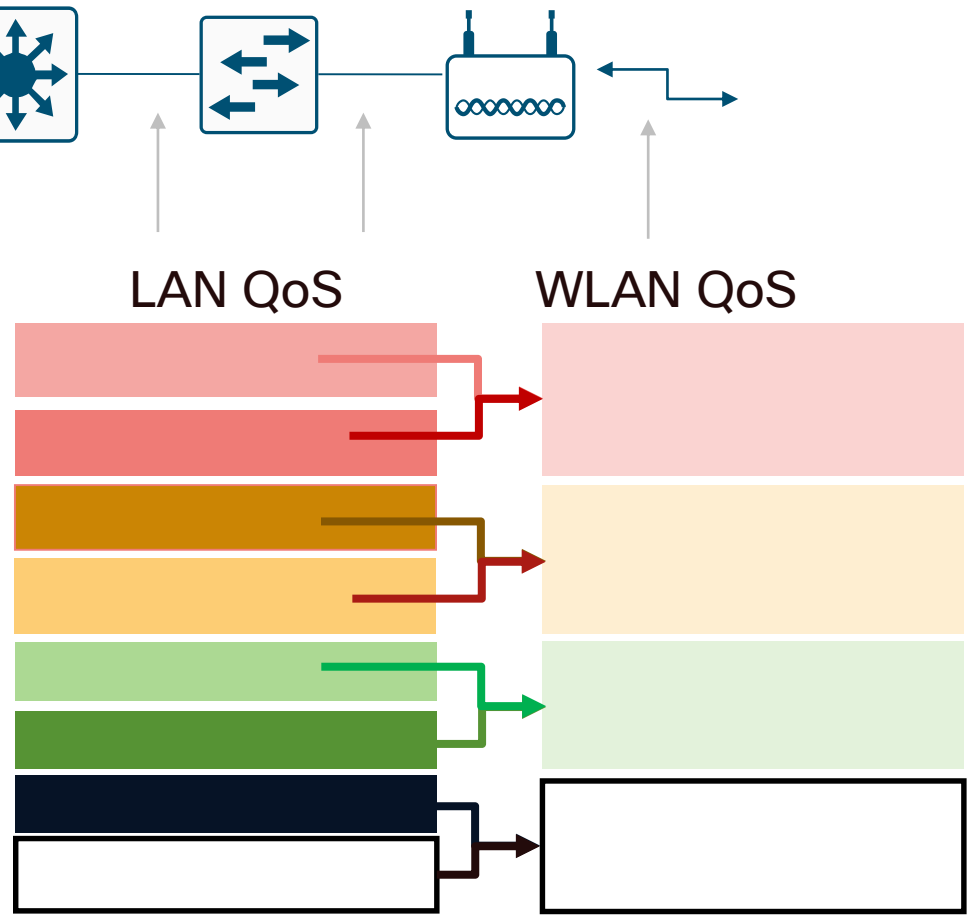

Catalyst High-Density Wireless Design

BRKEWN-2087

High-Density Wi-Fi Design, Deployment and Optimization

Matt Swartz - Distinguished Engineer, Cisco
Josh Suhr - Principal Architect, Cisco

This session will cover an array of detailed tips, tricks, and tools for configuring and optimizing high-density Wi-Fi networks in today's most challenging environments, including valuable real-world insights on how the latest Cisco Catalyst access points, antennas, and Catalyst wireless controllers can be leveraged to deliver an optimized and reliable user experience.

Attend this session to hear real-world hints directly from engineers who have deployed networks for some of the largest sports venues and events worldwide.

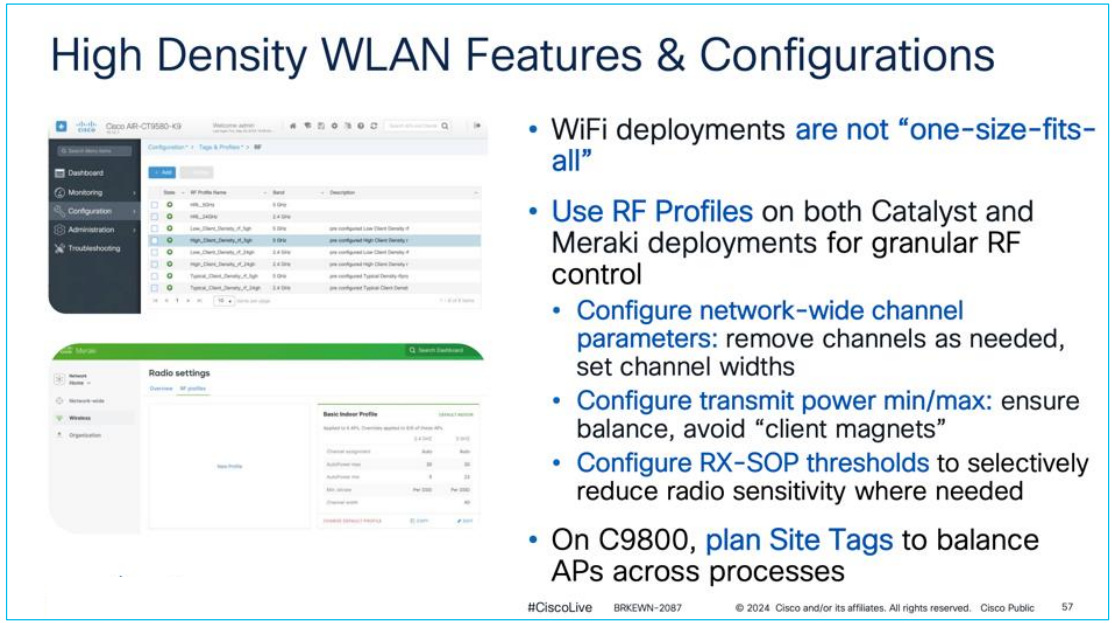

Wireless & Firewall

- Central Wireless LAN
- Firewalls, VRFs & ACLs

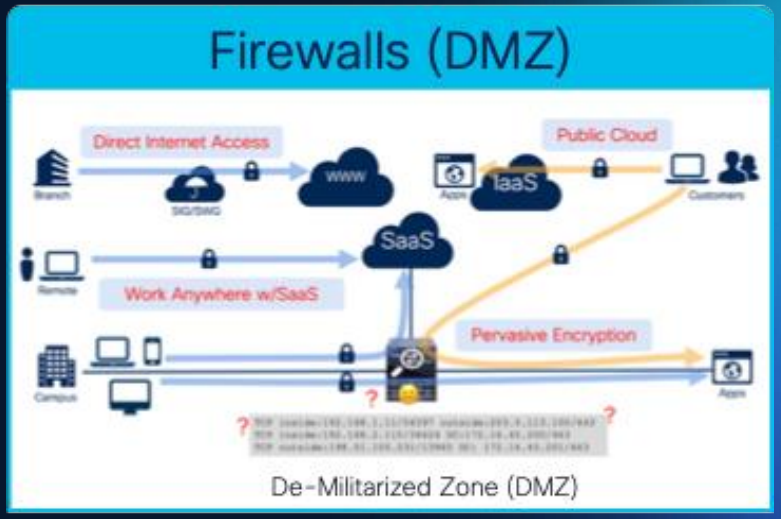

Firewalls, VRFs & ACLs

The **Firewall (DMZ) PIN** focuses on controlling access into or out of different network areas.

- Typically connected to Core, Edge or DC (Tier 3+)
- Complex designs may use Distro or Access (Tier 1-2)

Main goal is to prevent unauthorized access to different network domains (segments).

- Evolved from "Edge" Access-Control Lists (ACLs)
- Can be either L2, L3 or VRF-aware
- Tends to focus on L4-L7 flows (with or w/o DPI)

Uses a **L2 or L3/VRF + ACLs**

- North (outside): L2 802.1Q, L3 (SVI, Sub-Ints), IGP, BGP
- South (inside): L2 802.1Q, L3 (SVI, Sub-Ints), IGP, BGP

Tends to use **L2 & L3/VRF + DPI & ACL features**

- L4/App ACLs (e.g. VACL, MAC ACL)
- L4/App QoS (e.g. VLAN QoS)
- L4/App NetFlow (e.g. FNF, AVC, EPA & ETA)

Tends to require **med-high L2/L3 & feature scale**

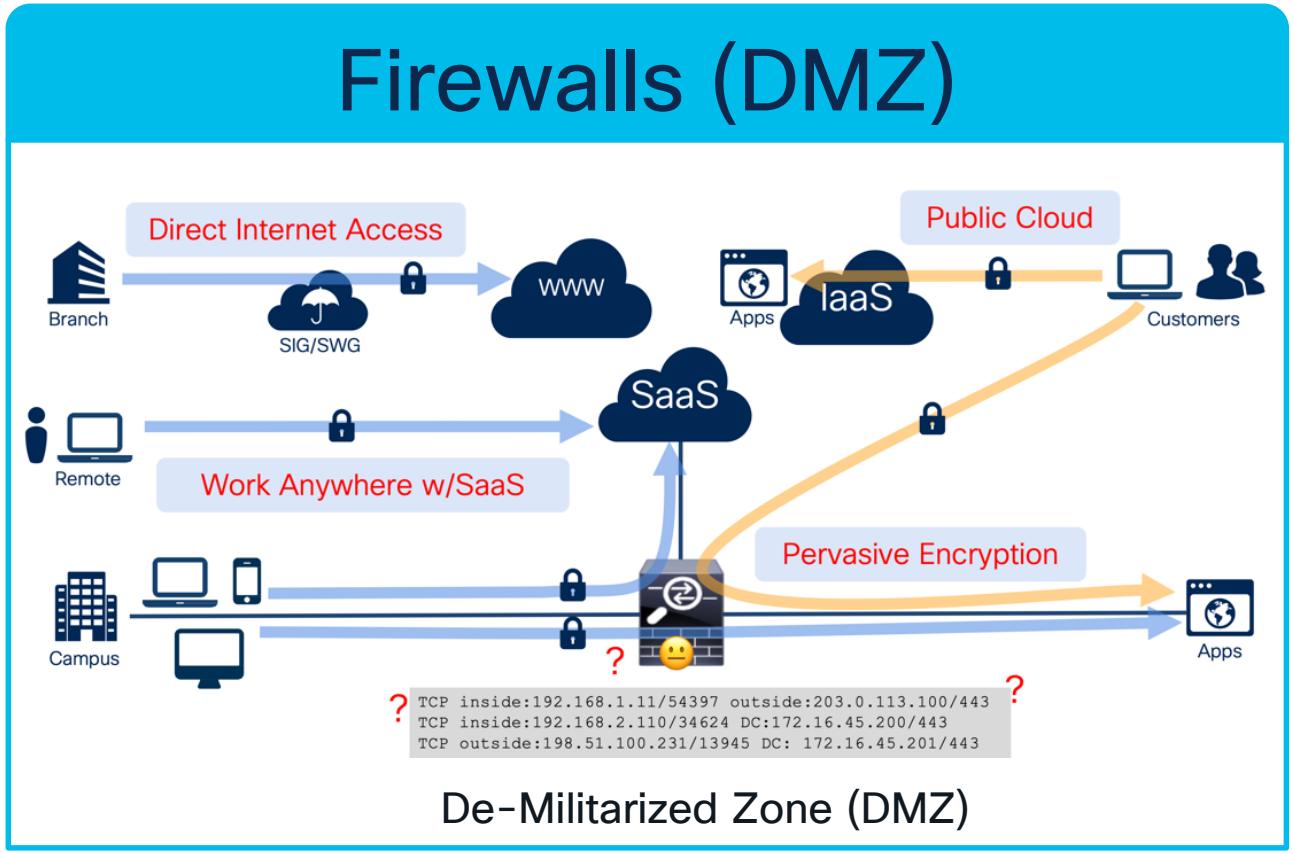

Direct FW Connection

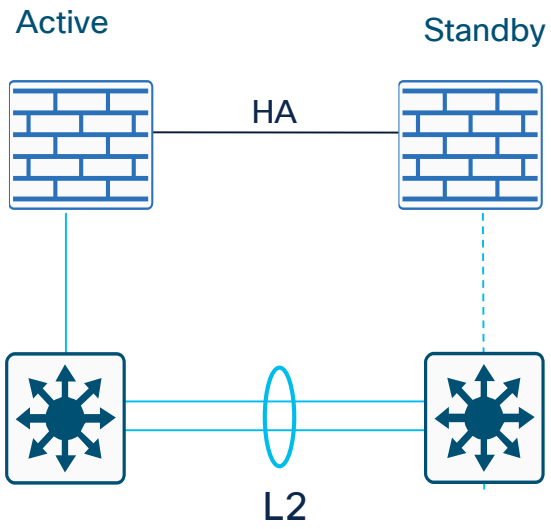

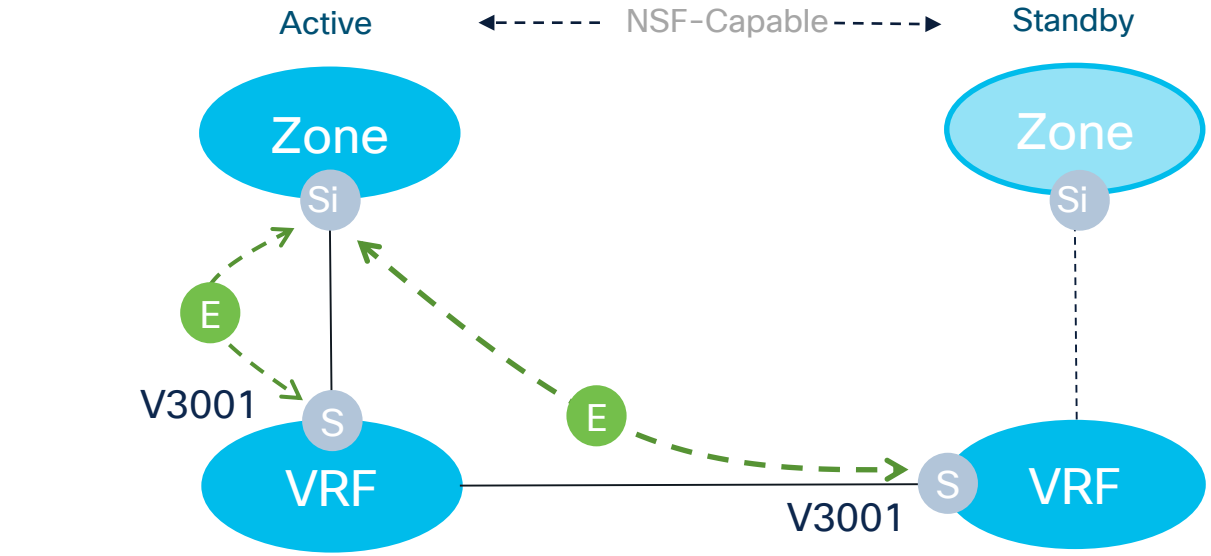


```
interface Vlan3001
 vrf forwarding IOT
 ip address 172.16.15.1/29
```

```
interface Vlan3001
 vrf forwarding IOT
 ip address 172.16.15.2/29
```

- L2 extension required between Core nodes – consider STP implications
- Active/Standby L2 adjacency with both Borders required
- FTD supports BGP with NSF

Si sub-int
S SVI
E eBGP

Direct FW Connection with ECMP

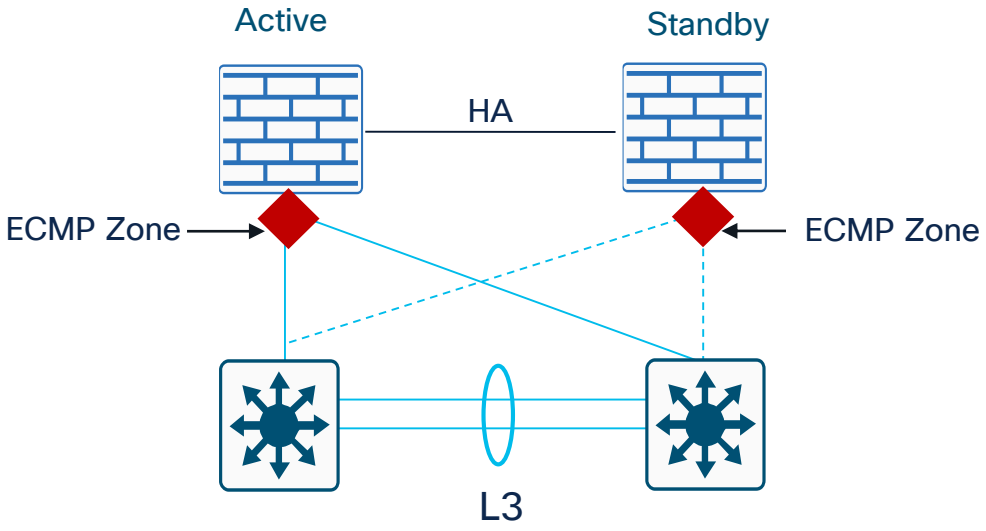

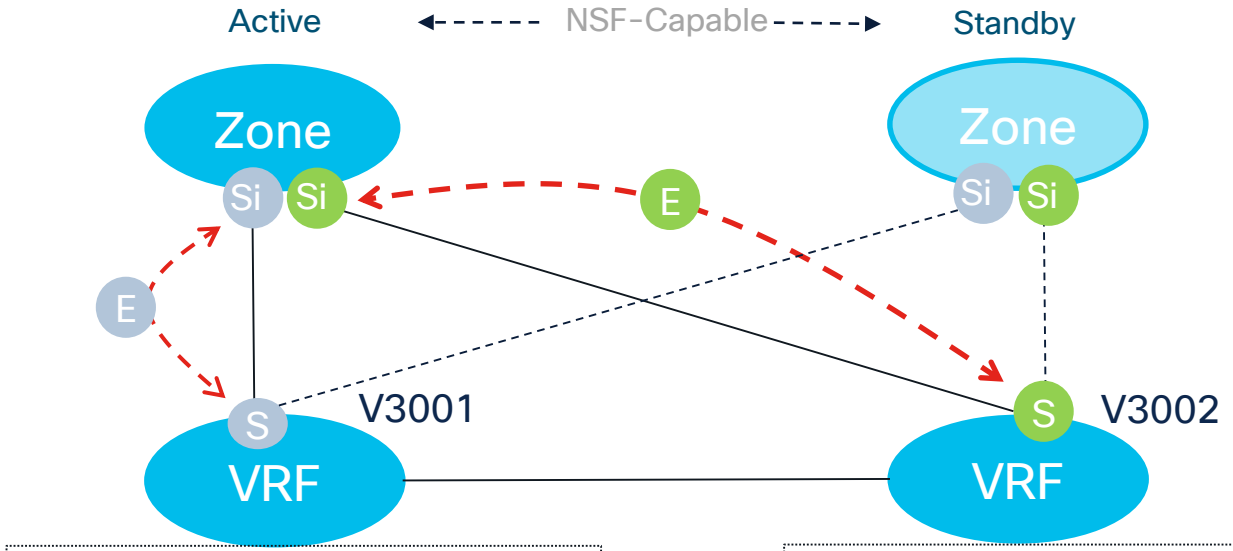


```
interface Vlan3001
 vrf forwarding IOT
 ip address 172.16.1.1/29
```

```
interface Vlan3002
 vrf forwarding IOT
 ip address 172.16.2.1/29
```

- ECMP Zone allows grouping of interfaces together – FTD maintains connections per zone table
- Both ECMP uplinks active at the same time load-sharing traffic to/from both Core nodes.
- ECMP is supported under global and user VRFs on FTD

Si sub-int
S SVI
E eBGP

FTD Cluster with MEC Switching Layer

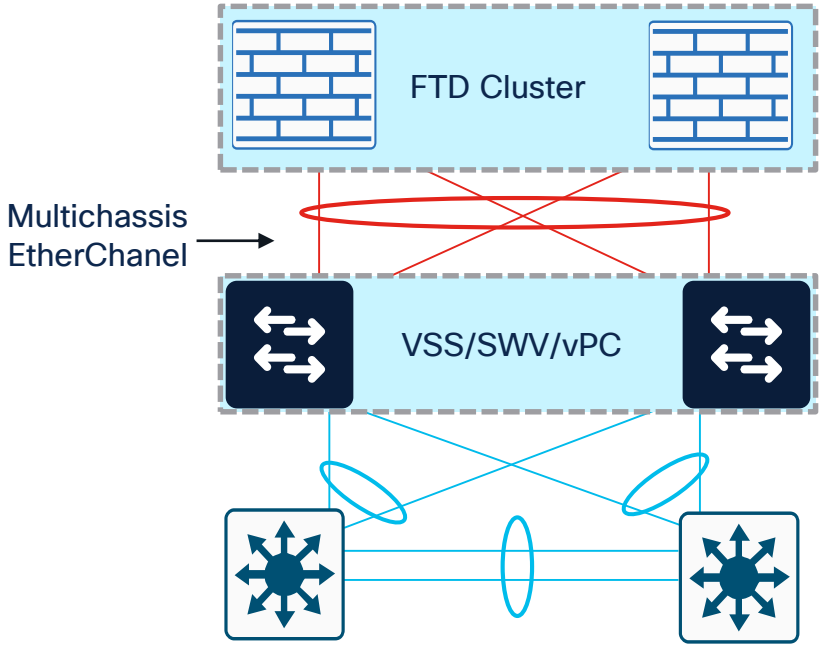

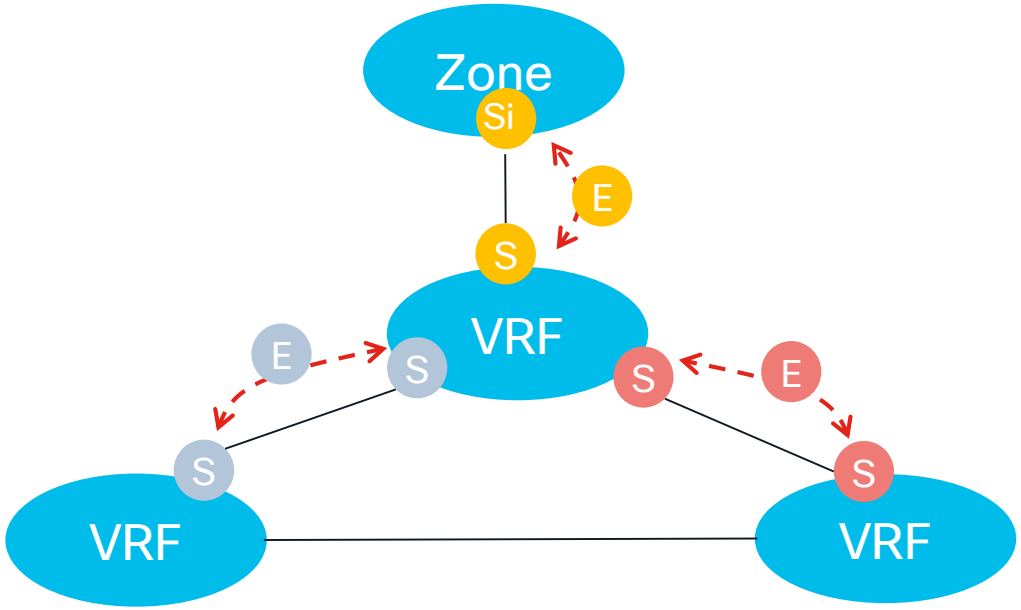

- An FTD Cluster behaves as a single firewall – packets are switched internally between members to ensure stateful and symmetric NGFW / NGIPS services.
- Intermediate switch layer load-balances packets to multiple FTD cluster members over Multichassis EtherChannel
- VRF and Multi-instance supported

Si sub-int
S SVI
E eBGP

FTD High Availability

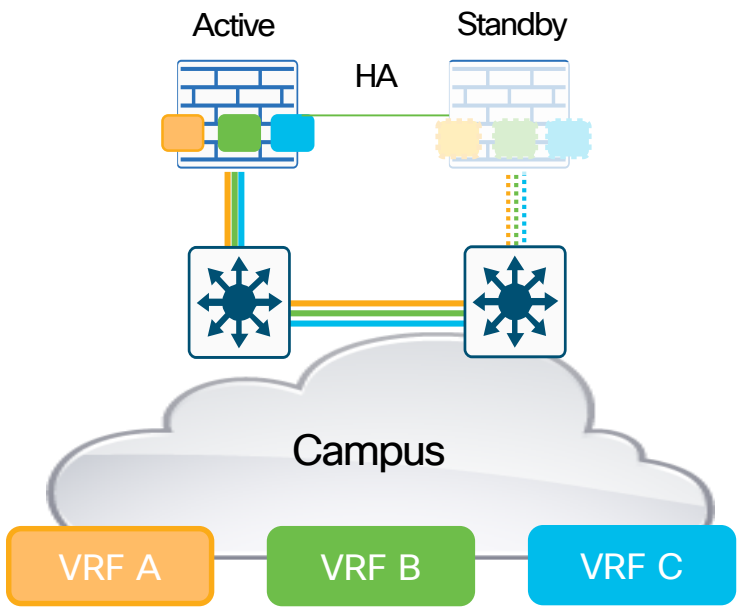

Active/Standby

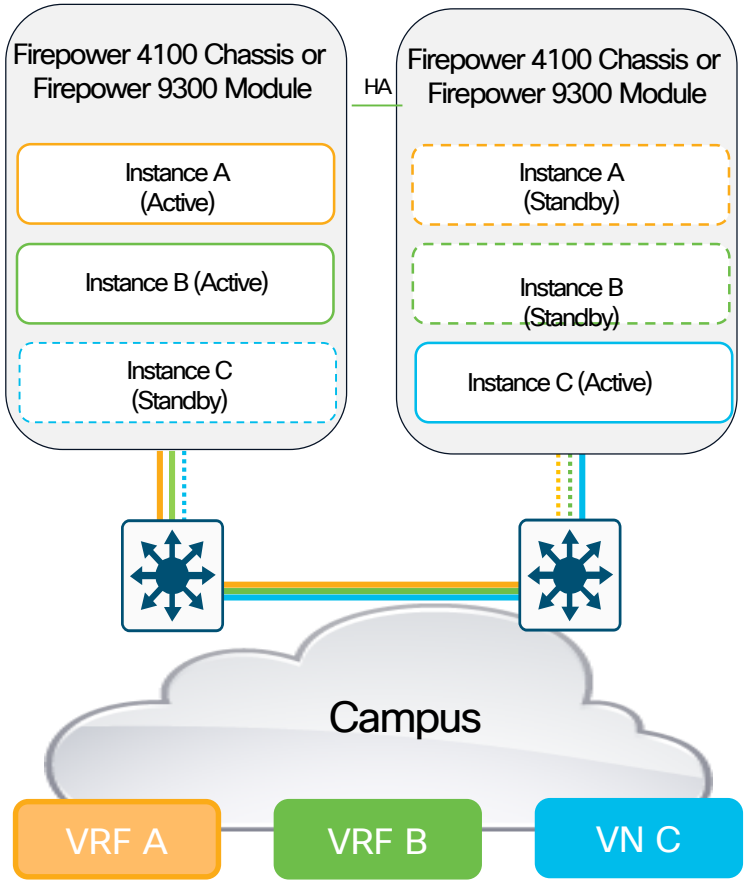

Semi-Active/Active
(w/ Multi-Instance)

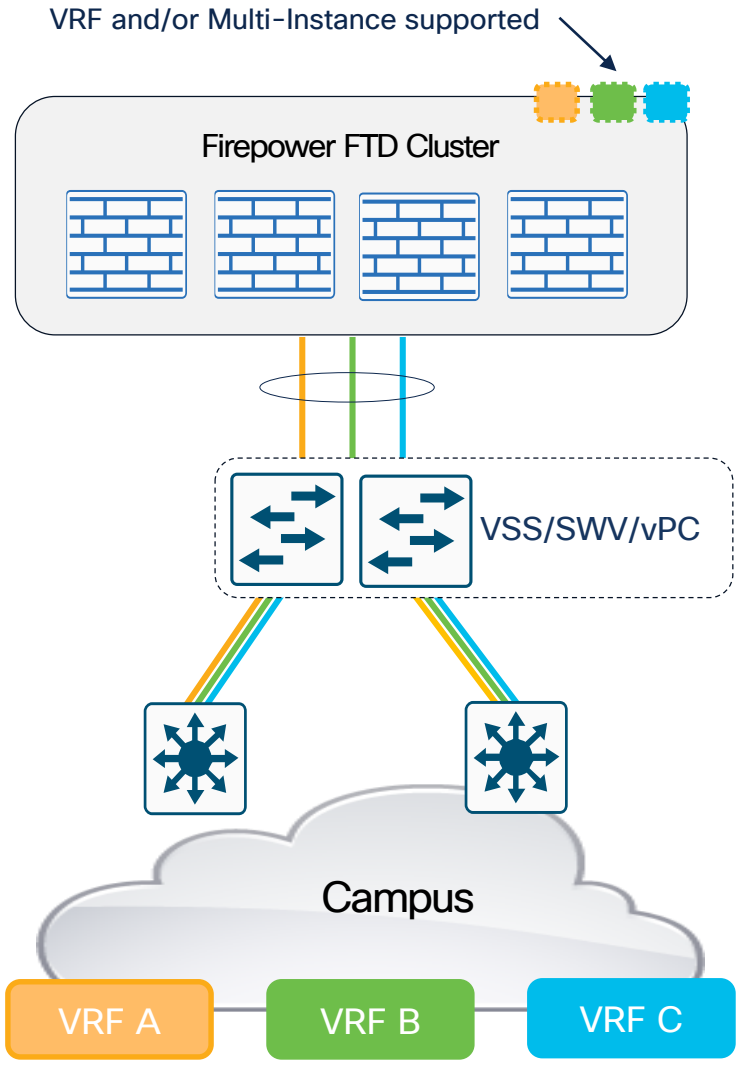

Cluster

Cisco SDA & Mult-Domain Segmentation

BRKENS-2819

Cisco SD-Access and Multi-Domain Segmentation

Jerome Dolphin - Technical Marketing, Cisco

Using Cisco SD-Access LISP Fabric as a foundation this session will explore and integrate important LAN, WAN and cloud segmentation solutions. Theory will be reinforced with real-world demonstrations of each concept.

Beginning with Catalyst Center, cabled networking devices, and a basic preliminary configuration, you will learn how and why to build end-to-end network segmentation while preserving and adapting security context.

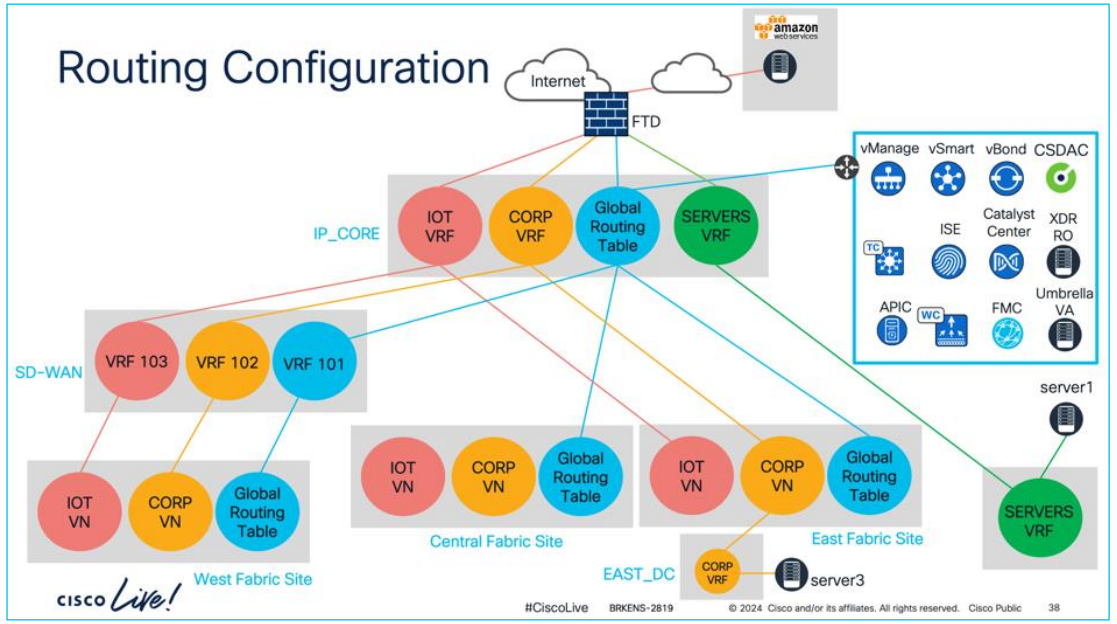

Wrap up

- Know the Campus PINs
- Other References
- Keep Learning!! ☺

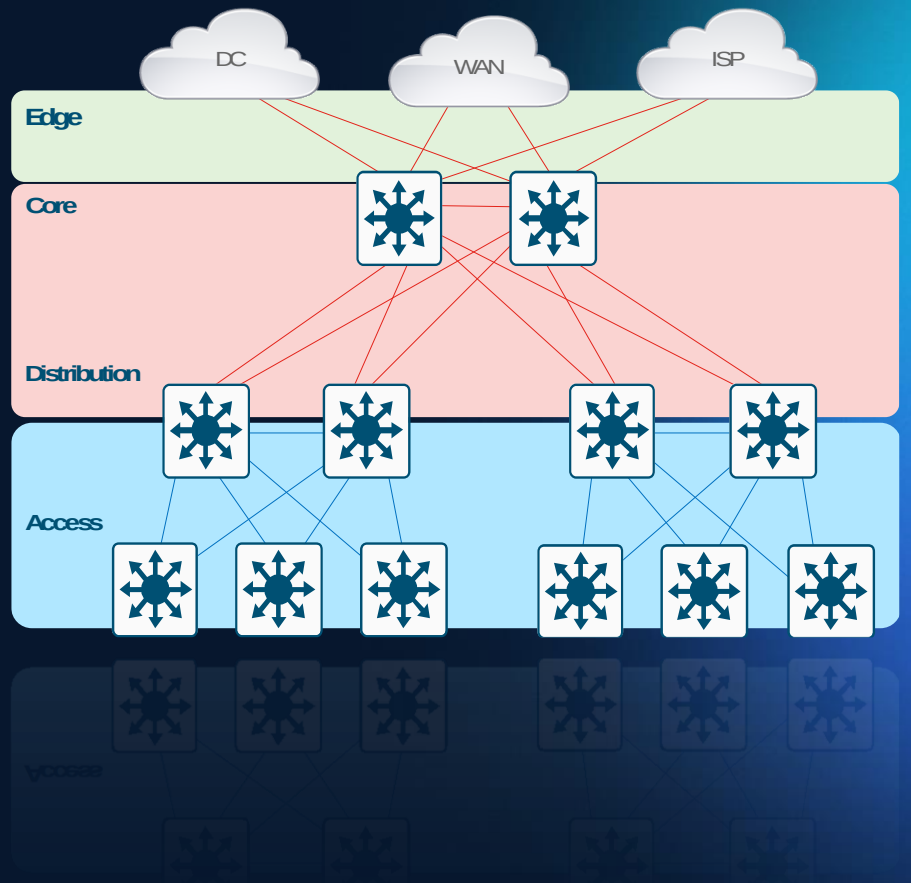

Remember: Campus PINs & Topology

RECAP

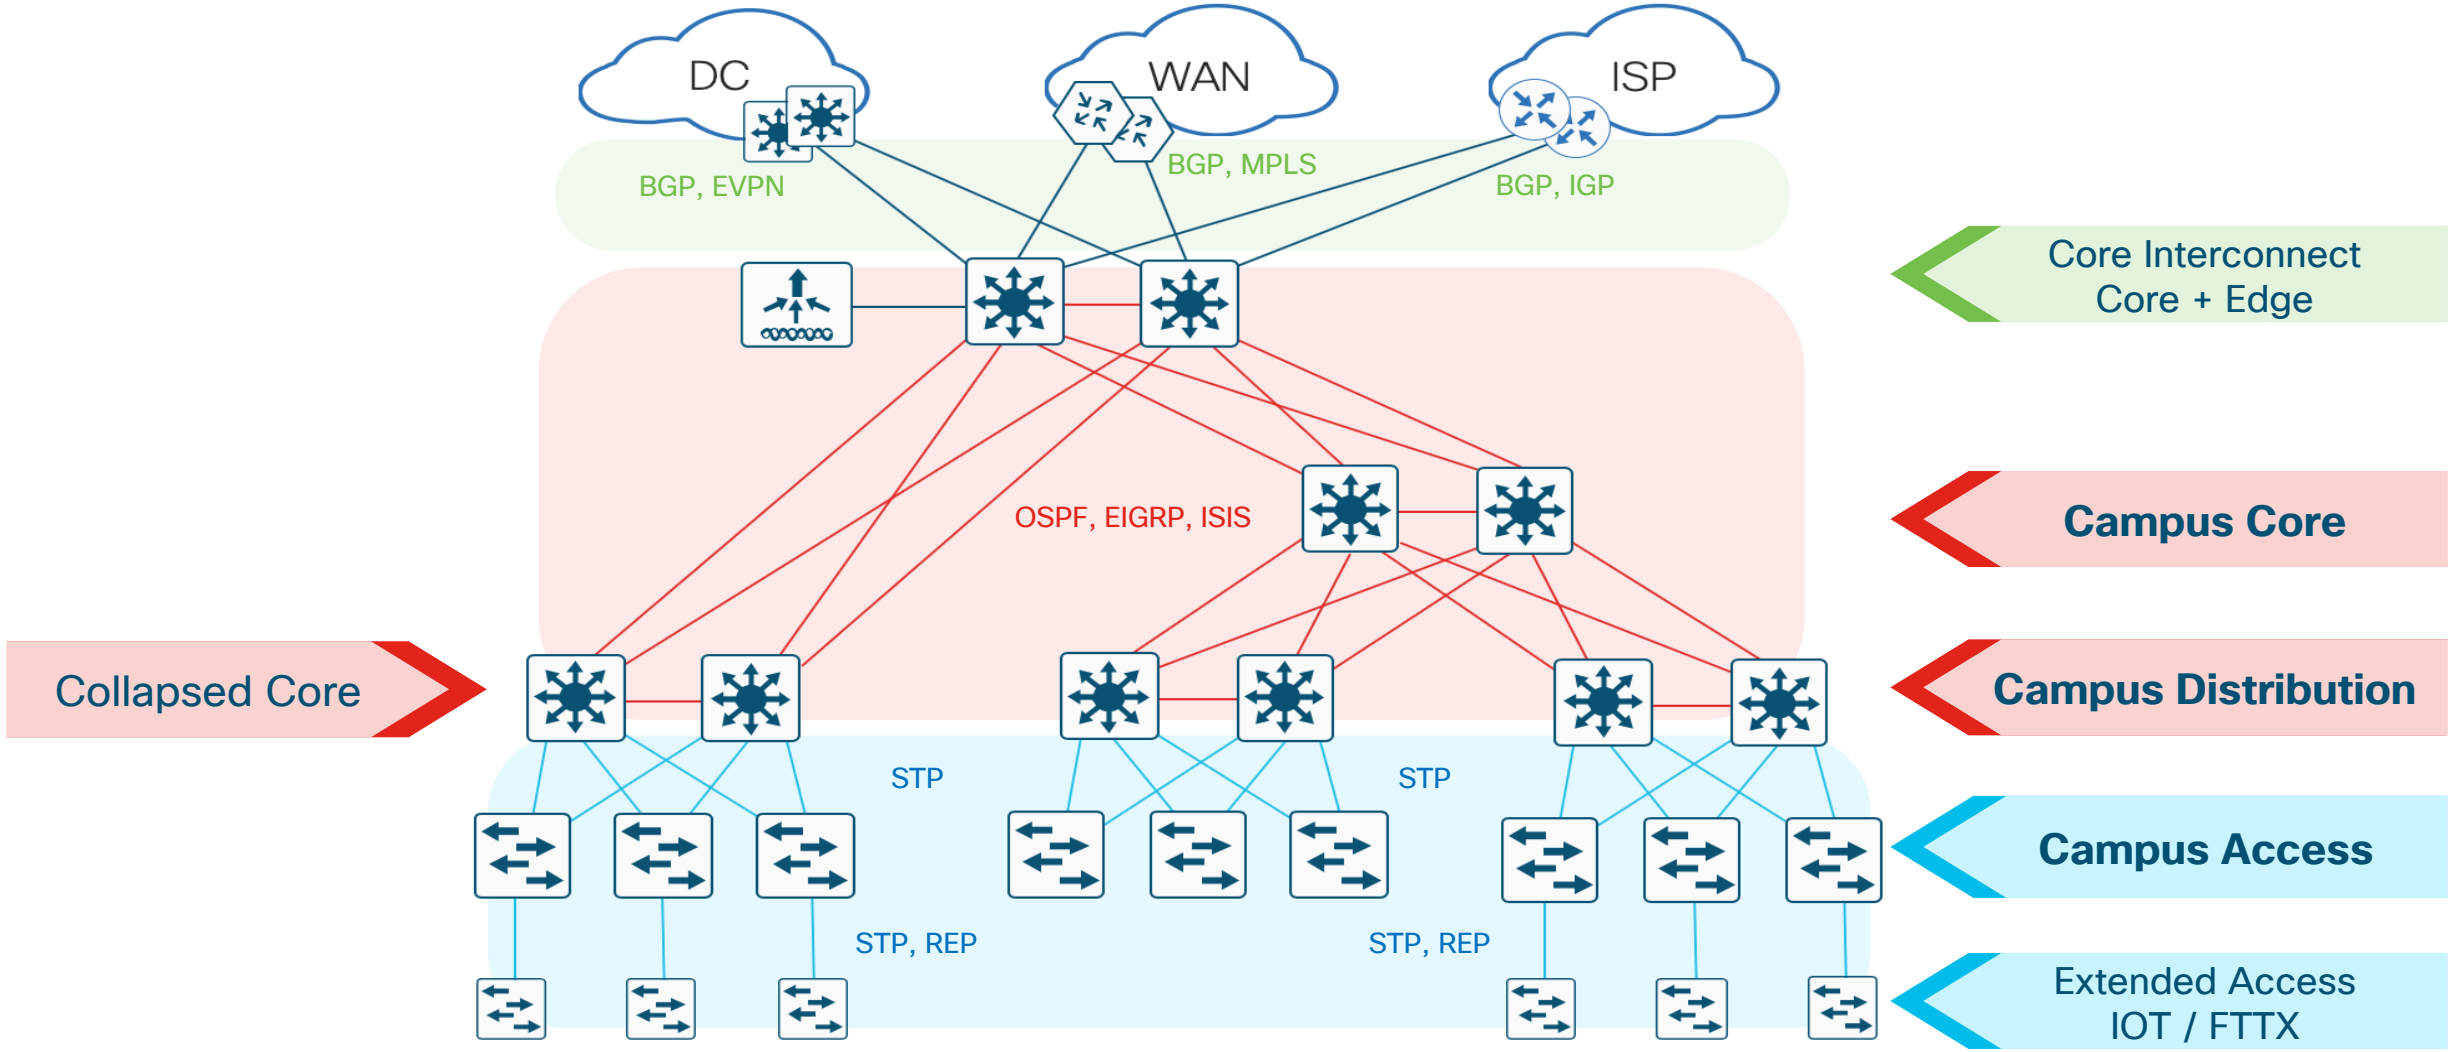

Remember: Hierarchical Campus

RECAP

Access

Distribution

Core

Distribution

Access

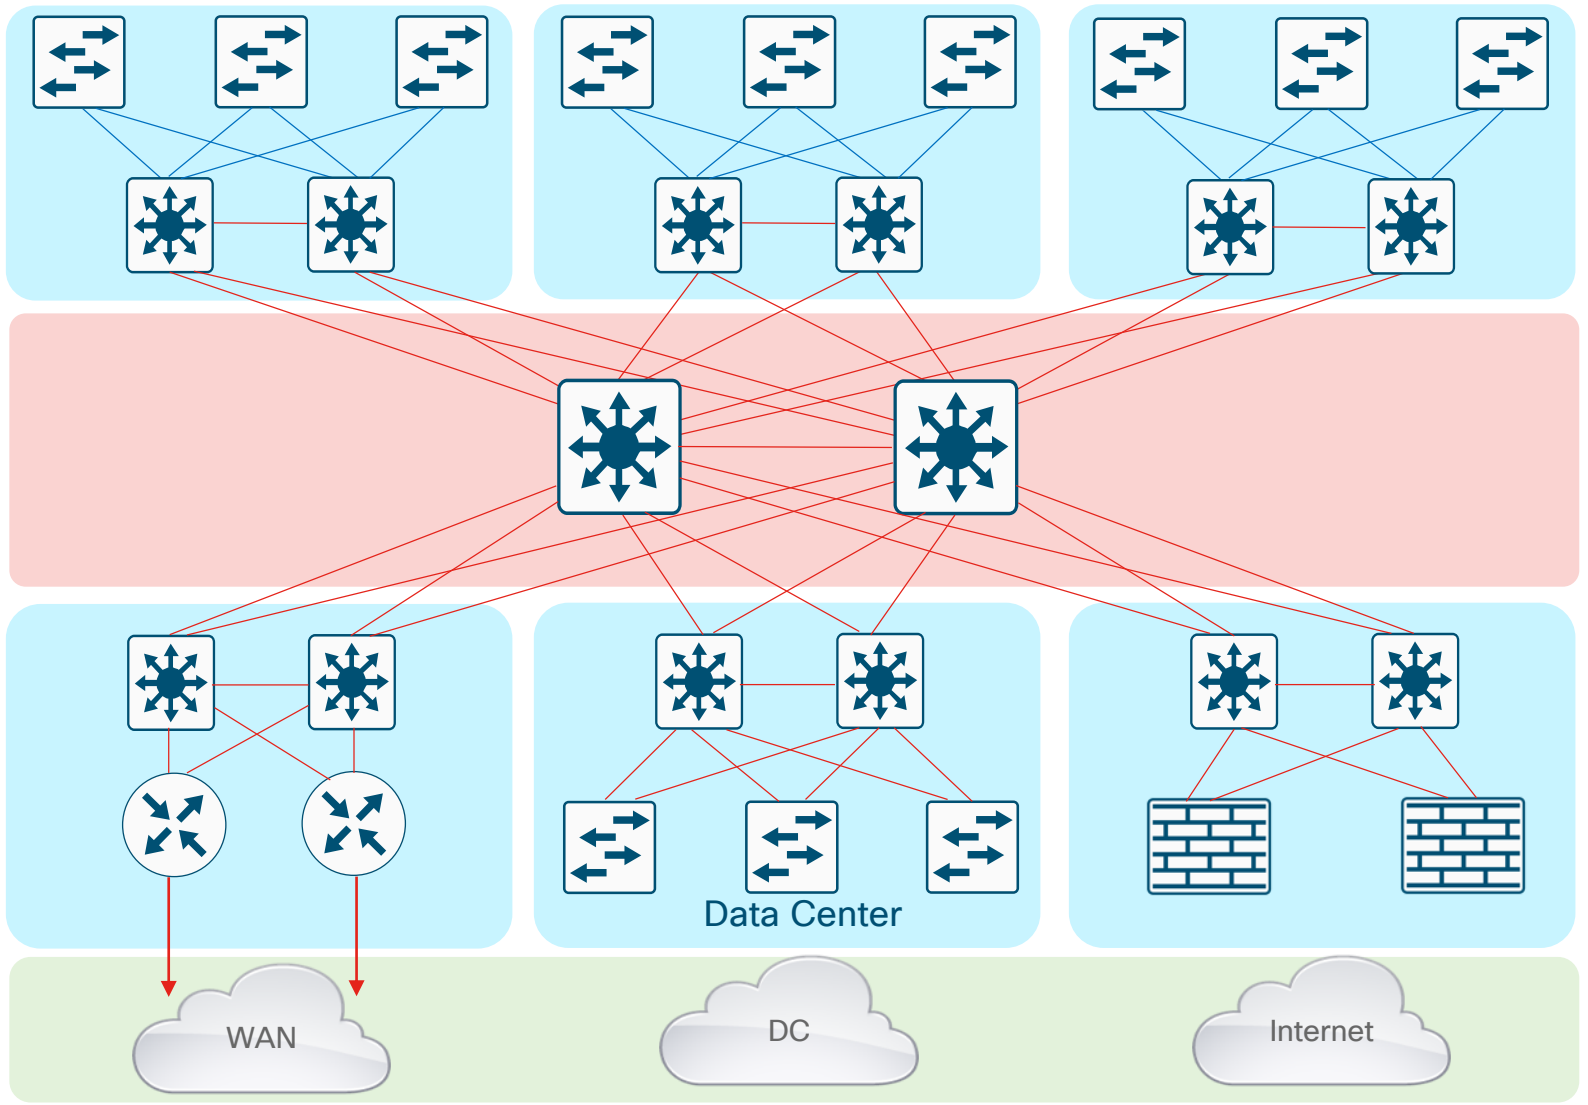

Remember: Campus Design Fundamentals

RECAP

Collapse or Expand Layers?

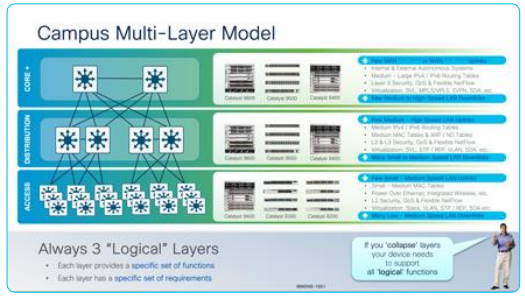

Modular or Fixed Platforms?

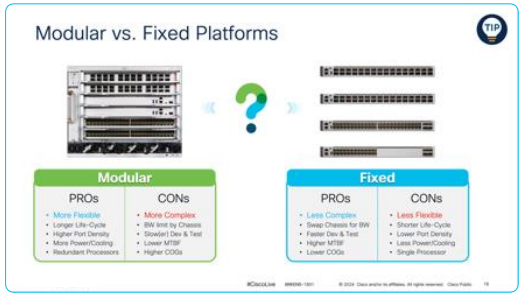

Fiber or Copper Links?

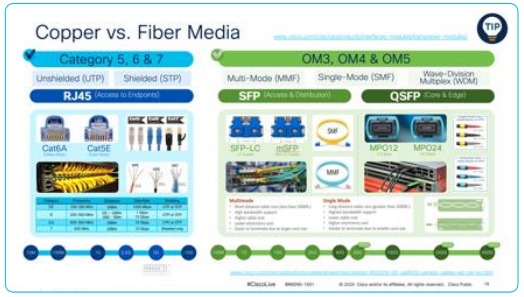

ECMP, EtherChannel or Stacking?

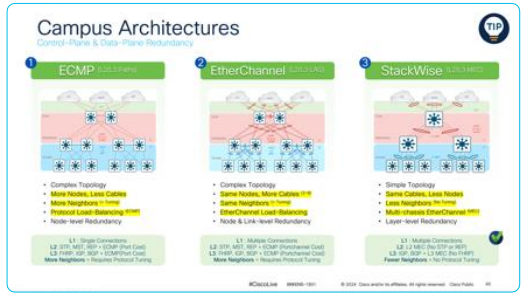

L2/L3 or MPLS or VXLAN?

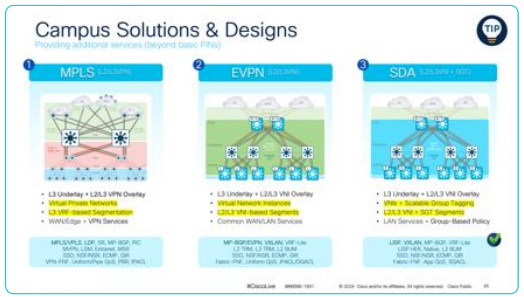

Wireless or Security Included?

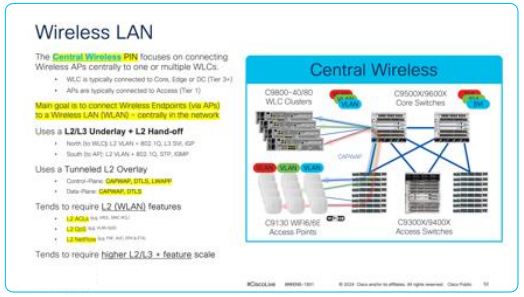

cisco

Catalyst Leadership in Enterprise Networks

New Feature
Enhanced

A Platform based Approach

Catalyst Center and Meraki Dashboard

28M Network Devices Managed

50% Y/Y

19M APs | 6M Switches | 2.5M Routers | 830M Clients

13M
Devices on Catalyst Center

15.3M
Devices on Meraki Dashboard

Catalyst 9000 Family

100,000+ Customers, Millions of Switches

"Catalyst 9K continues to be the fastest ramping product in the company's history"

- Chuck Robbins, CEO Cisco Systems

Secure Networking

- Common Policy
- Secure Equipment Access
- SD-Access (LISP & EVPN)
- High-speed Encryption

Digital Experience

- Campus Automation
- AI Endpoint Analytics
- ThousandEyes Digital Experience
- AI Ops & Assurance

Operational Simplicity

- Cloud Managed Catalyst
- Infrastructure as a Code
- S3 & CloudWatch Integration
- Visibility, Control & Rollback

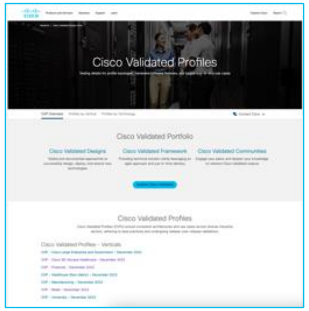

Cisco Validated Profiles (CVP)

Industry Validated Reports

Industry Certifications

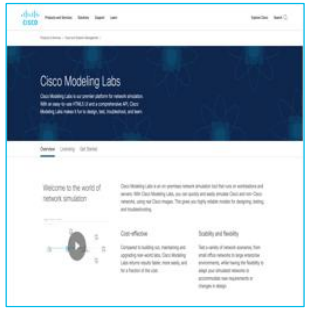

Cisco Modeling Labs

Keep Learning!

Cisco Validated Design (CVD)

cisco.com/go/cvd
cs.co/en-cvds

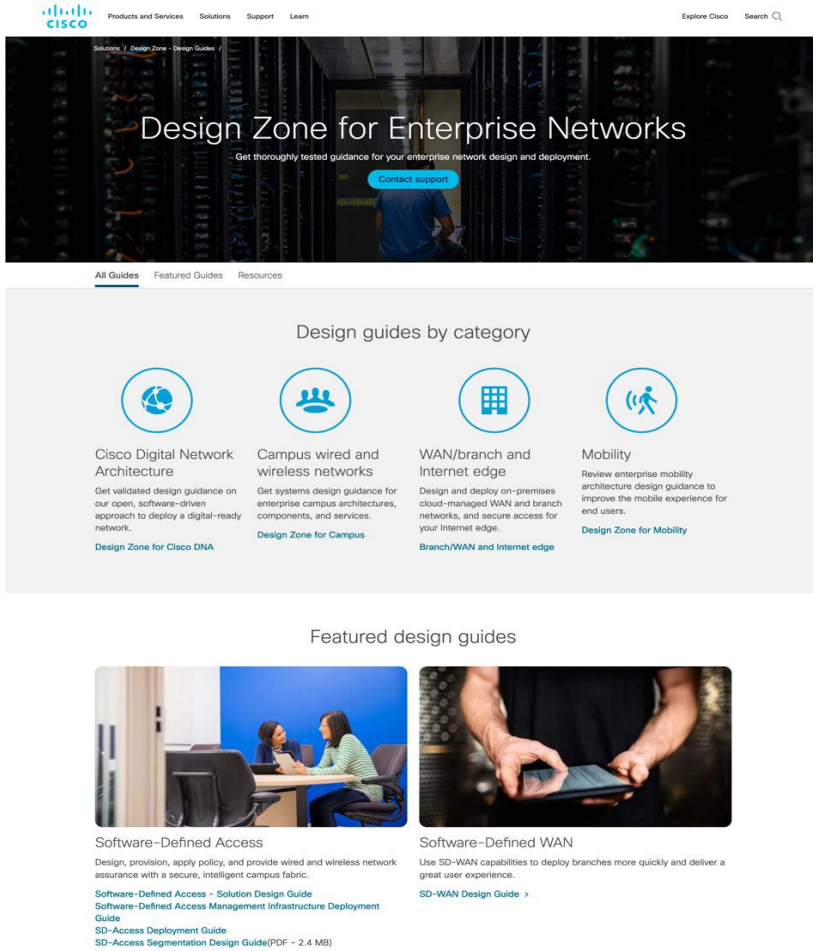

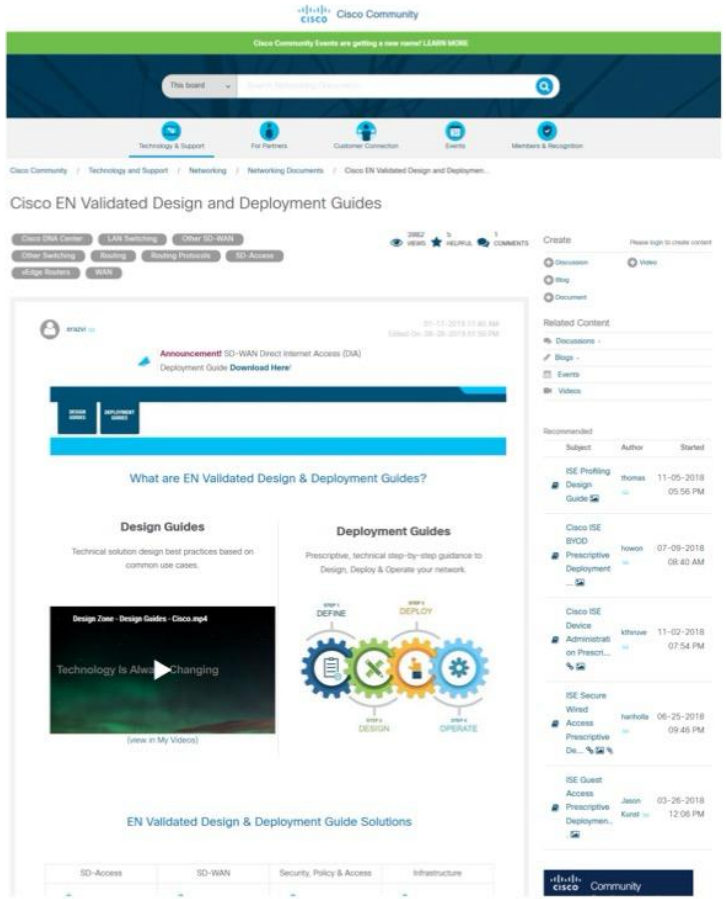

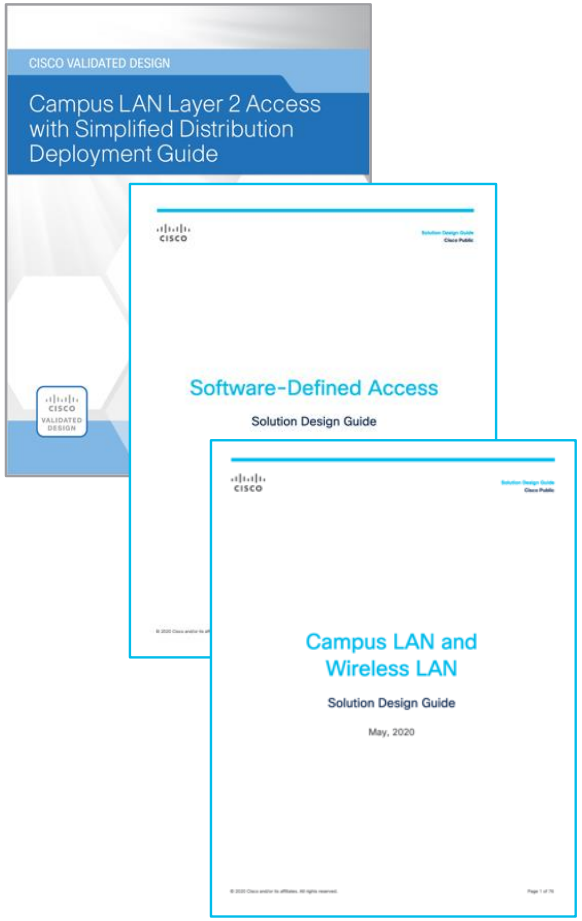

References – Multi-Layer Campus

Type	Sub-Type	References
General	Multi-Layer	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html www.ccexpert.us/network-design-2/designing-a-campus-network-design-topology.html networkdirection.net/articles/network-theory/hierarchicalnetworkmodel www.geeksforgeeks.org/types-of-area-networks-lan-man-and-wan/
Core	Edge	www.atlantic.net/managed-services/network-edge/ www.ccexpert.us/network-design/enterprise-edge-modules.html what-when-how.com/ipv6-for-enterprise-networks/enterprise-edge-network-design-ipv6/
	Interconnect	www.geeksforgeeks.org/difference-between-lan-and-man www.ti.com/solution/intra-dc-interconnect-metro en.wikipedia.org/wiki/Backbone_network
	Baseline	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#Corelayer www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html#wp1107724 www.ccexpert.us/network-design/campus-core-design-considerations.html en.wikipedia.org/wiki/Hierarchical_internetworking_model#Core_layer
Distribution	Collapsed Core	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#Twotierdesign www.econfigs.com/ccna-1-5-compare-and-contrast-collapsed-core-and-three-tier-architectures interestingtraffic.nl/2018/06/08/collapsed_core_design oreilly.com/library/view/ccna-data-center/9780133860429/ch01lev3sec4.html
	Baseline	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#Distributionlayer www.ccexpert.us/network-design/building-distribution-layer-design-considerations.html en.wikipedia.org/wiki/Hierarchical_internetworking_model#Distribution_layer
Access	Baseline	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#Accesslayer www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html#wp1107746 www.ccexpert.us/network-design/building-access-layer-design-considerations.html en.wikipedia.org/wiki/Hierarchical_internetworking_model#Access_layer
	Routed Access	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#Layer3routedaccesscampusdesign www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html#wp1108952
	Extended/IOT	www.cisco.com/c/en/us/td/docs/solutions/Verticals/CCI/CCI/DG/cci-dg/cci-dg.html#99480 www.geeksforgeeks.org/5-layer-architecture-of-internet-of-things/

References – ECMP & StackWise(Virtual)

Type	Sub-Type	References
General	Redundancy	www.cisco.com/c/en/us/solutions/hybrid-work/what-is-high-availability.html#~infrastructure-elements www.ccexpert.us/network-design/designing-link-redundancy.html www.geeksforgeeks.org/redundant-link-problems-in-computer-network/
Core	ECMP	www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5212-46.html www.ccexpert.us/routing-protocols/equalcost-load-balancing.html en.wikipedia.org/wiki/Equal-cost_multi-path_routing
	EtherChannel	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#EtherChannel en.wikipedia.org/wiki/Link_aggregation#Network_backbone en.wikipedia.org/wiki/Multi-chassis_link_aggregation_group
	SVL	www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/BRKCRS-2650.pdf www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/nb-06-cat-9k-stack-wp-cte-en.html www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#StackWiseVirtualTechnology
Distribution	ECMP	www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5212-46.html www.ccexpert.us/routing-protocols/equalcost-load-balancing.html en.wikipedia.org/wiki/Equal-cost_multi-path_routing
	EtherChannel	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#EtherChannel en.wikipedia.org/wiki/Link_aggregation en.wikipedia.org/wiki/Multi-chassis_link_aggregation_group
	SVL	www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/BRKCRS-2650.pdf www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/nb-06-cat-9k-stack-wp-cte-en.html www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#StackWiseVirtualTechnology
Access	ECMP	www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10555-15.html en.wikipedia.org/wiki/Spanning_Tree_Protocol#Path_to_the_root_bridge en.wikipedia.org/wiki/Flex_links
	EtherChannel	www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#EtherChannel en.wikipedia.org/wiki/EtherChannel
	Stacking	www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/BRKCRS-2650.pdf www.cisco.com/c/en/us/products/collateral/switches/catalyst-9300-series-switches/white-paper-c11-741468.html www.cisco.com/c/en/us/products/collateral/switches/catalyst-9200-series-switches/nb-06-stackwise-architecture-cte-en.html www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#SwitchStacksandCiscoStackWiseTechnology

References – SD-Access, EVPN & MPLS

Type	Sub-Type	References
General	SDN/IBN	www.cisco.com/c/en/us/solutions/intent-based-networking.html www.networkworld.com/article/3281447/a-new-era-of-campus-network-design.html www.geeksforgeeks.org/difference-between-software-defined-network-and-traditional-network/
Core	SDA	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKCRS-2810.pdf#page=27 www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html#BorderNode www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html#ControlPlaneNode
	EVPN	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2021/pdf/BRKENS-2003.pdf#page=12 www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-7/configuration_guide/vxlan/b_177_bgp_evpn_vxlan_9500_cg/bgp_evpn_vxlan_overview.html#id_126799 www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#AlternativevirtualizationdesignforcampusBGPEVPNVXLAN
	MPLS	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-1100.pdf#page=48 www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-2112.pdf#page=42 www.geeksforgeeks.org/multi-protocol-label-switching-mpls/
Distribution	SDA	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKCRS-2810.pdf#page=19 www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html#IntermediateNode
	EVPN	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2021/pdf/BRKENS-2003.pdf#page=12 www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-7/configuration_guide/vxlan/b_177_bgp_evpn_vxlan_9500_cg/bgp_evpn_vxlan_overview.html#id_126799
	MPLS	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-1100.pdf#page=48 www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-2112.pdf#page=42 www.geeksforgeeks.org/multi-protocol-label-switching-mpls/
Access	SDA	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKCRS-2810.pdf#page=24 www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html#EdgeNode www.cisco.com/c/dam/en/us/solutions/collateral/internet-of-things/nb-09-intent-based-iot-wp-cte-en.pdf www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-lan-wlan-design-guide.html#CiscoSoftwareDefinedAccesscampusdesign
	EVPN	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2021/pdf/BRKENS-2003.pdf#page=12 www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-7/configuration_guide/vxlan/b_177_bgp_evpn_vxlan_9500_cg/bgp_evpn_vxlan_overview.html#id_126799
	MPLS	www.ciscolive.com/c/dam/r/ciscolive/us/docs/2020/pdf/DGTL-BRKMPL-1100.pdf#page=48 www.geeksforgeeks.org/multi-protocol-label-switching-mpls/

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 2026.

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

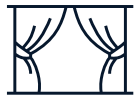

Visit the Cisco Showcase for related demos

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

Contact us at: swargo@cisco.com
jamatela@cisco.com

Thank you

cisco Live !


CISCO