The bridge to possible

Enterprise Campus Design

Multilayer Architectures and Design Principles

Marcin Hamróz, Principal Architect
Jarosław Gawron, Principal Engineer

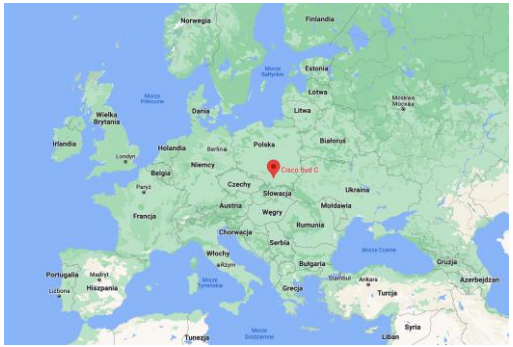

Cisco Krakow

Marcin Hamroz

Principal Architect

- At Cisco since 2012
- Based out of Cisco Krakow
- Focused on Software Defined Access
- CCIE R&S / SP
- Father of three
- Passionate about aviation and football

Jaroslaw (Jaro) Gawron

Principal Engineer

- In TAC from 2012
- Based out of Cisco Krakow
- Focused on Software Defined Access & Catalyst Platforms
- CCIE R&S / SP
- Father of three
- Fan of StarTrek and sailing

The goal of this session:

- Present the universal principles of Enterprise campus design
- Explain the most fundament aspect of the hierarchal approach for L2 and L3 networks (back to basics)
- Focus mainly on the wired campus

This is session is NOT:

- Covering SD-Access/ DNAC/ EVPN/ Cloud
- Product specific

Cisco Webex App

Questions?

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

How

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

Webex spaces will be moderated until February 24, 2023.

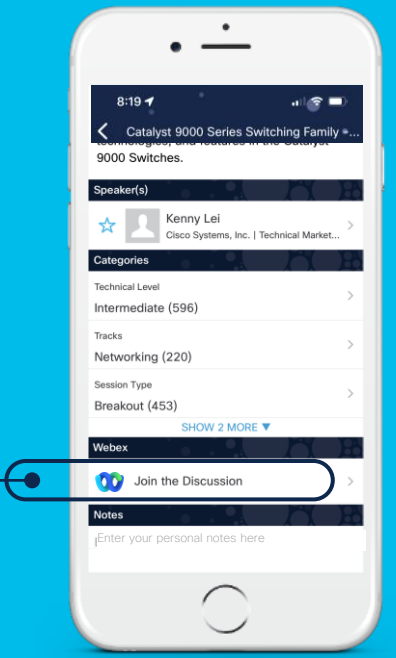

Agenda

- Introduction
- Campus Vision & Strategy
- Multilayer Campus Design Principles
- Foundation services
- Campus Design Best Practices
- Conclusion

Campus Vision & Strategy

Our Vision and Strategy

Vision

Change the way the world works, lives, plays, and learns

Strategy

Help Customers connect, secure and automate to accelerate their digital agility in a cloud-first world

Today's Network Must Drive Digital Transformation

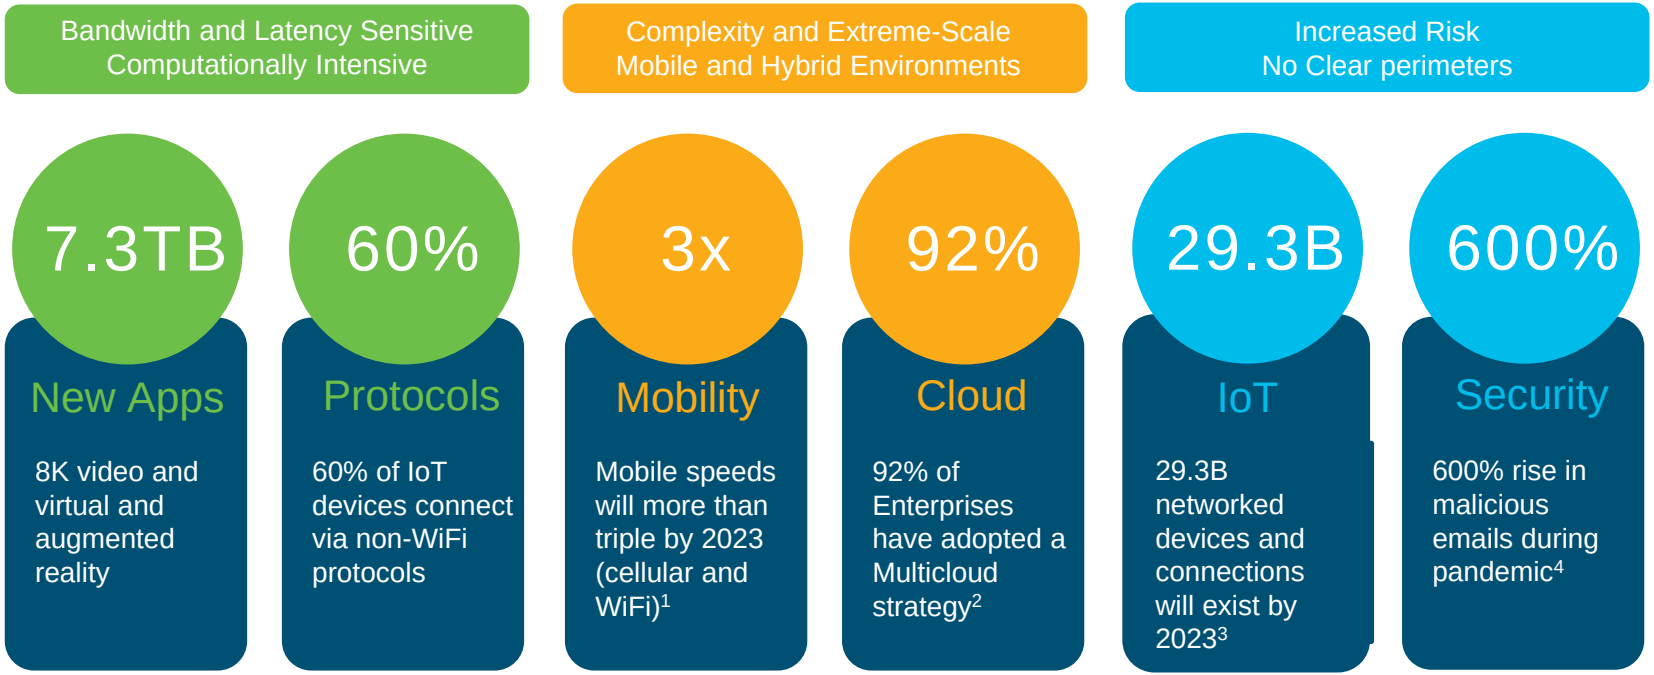

1 2020 Cisco Annual Internet Report
2 2018 MultiCloud in the New Normal
3 2020 Cisco Annual Internet Report
4 McAfee Report/Business Insider

Business Impact

Inefficiency

Up to 80% of network changes performed manually

Growth of Shadow IT Services

Complexity

3X spend on network operations

Slow and Error Prone Operations

Security

6 months to detect breach

Unconstrained Attack Surface

Cisco's Enterprise SDN Strategy

Policy and Intent to Unlock the Power of your Distributed System

Unlock the Power that Exists in the Network through **Abstraction, Automation, and Policy Enforcement**

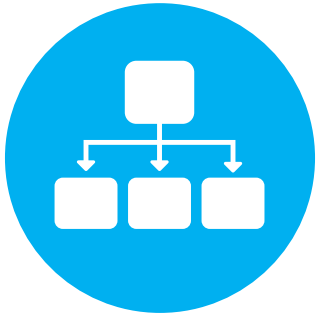

Leverage the Power of Existing **Distributed Systems**

Enable Network Wide Fidelity to an Expressed Intent **(Policy)**

Cisco's Intent Based Networking Solutions

SD-WAN
Segment your network and secure user access from the edge to the cloud

Cloud Edge
Securely connect and protect workloads moving into the cloud and between clouds

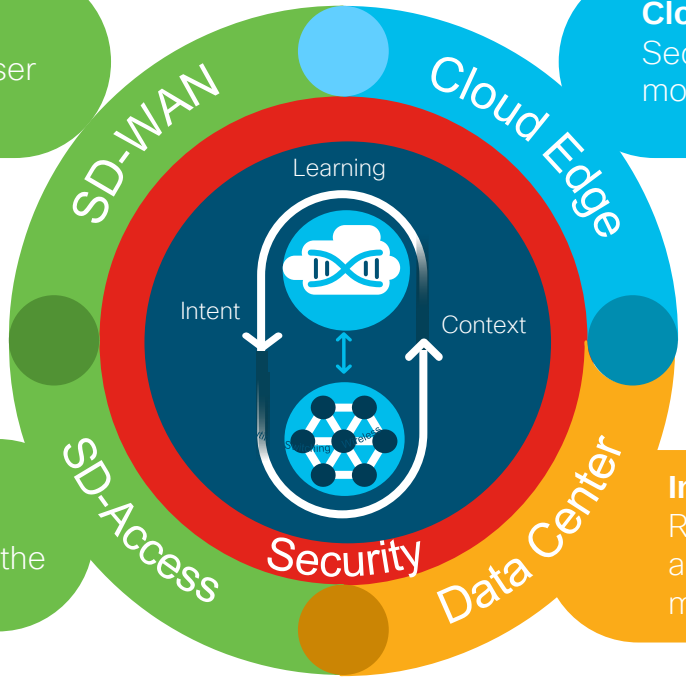

SD-Access
Optimize and secure application performance over any connection to the cloud.

Intent Based DataCenter
Run any traditional or cloud native application across any environment to meet evolving Developer needs

Built on Cisco Digital Network Architecture

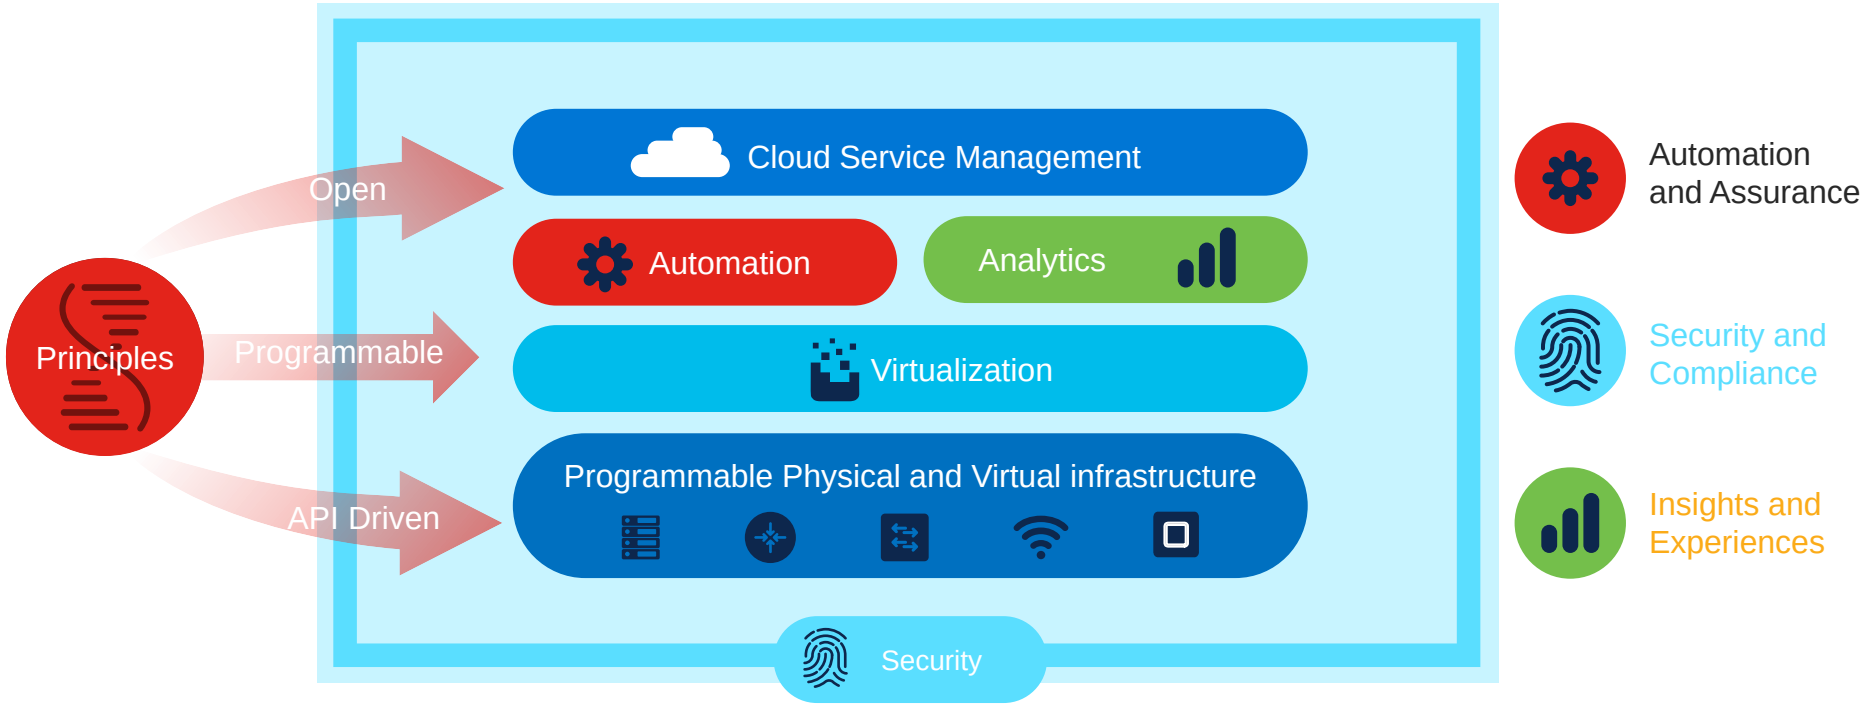

Multilayer Campus Design Principles

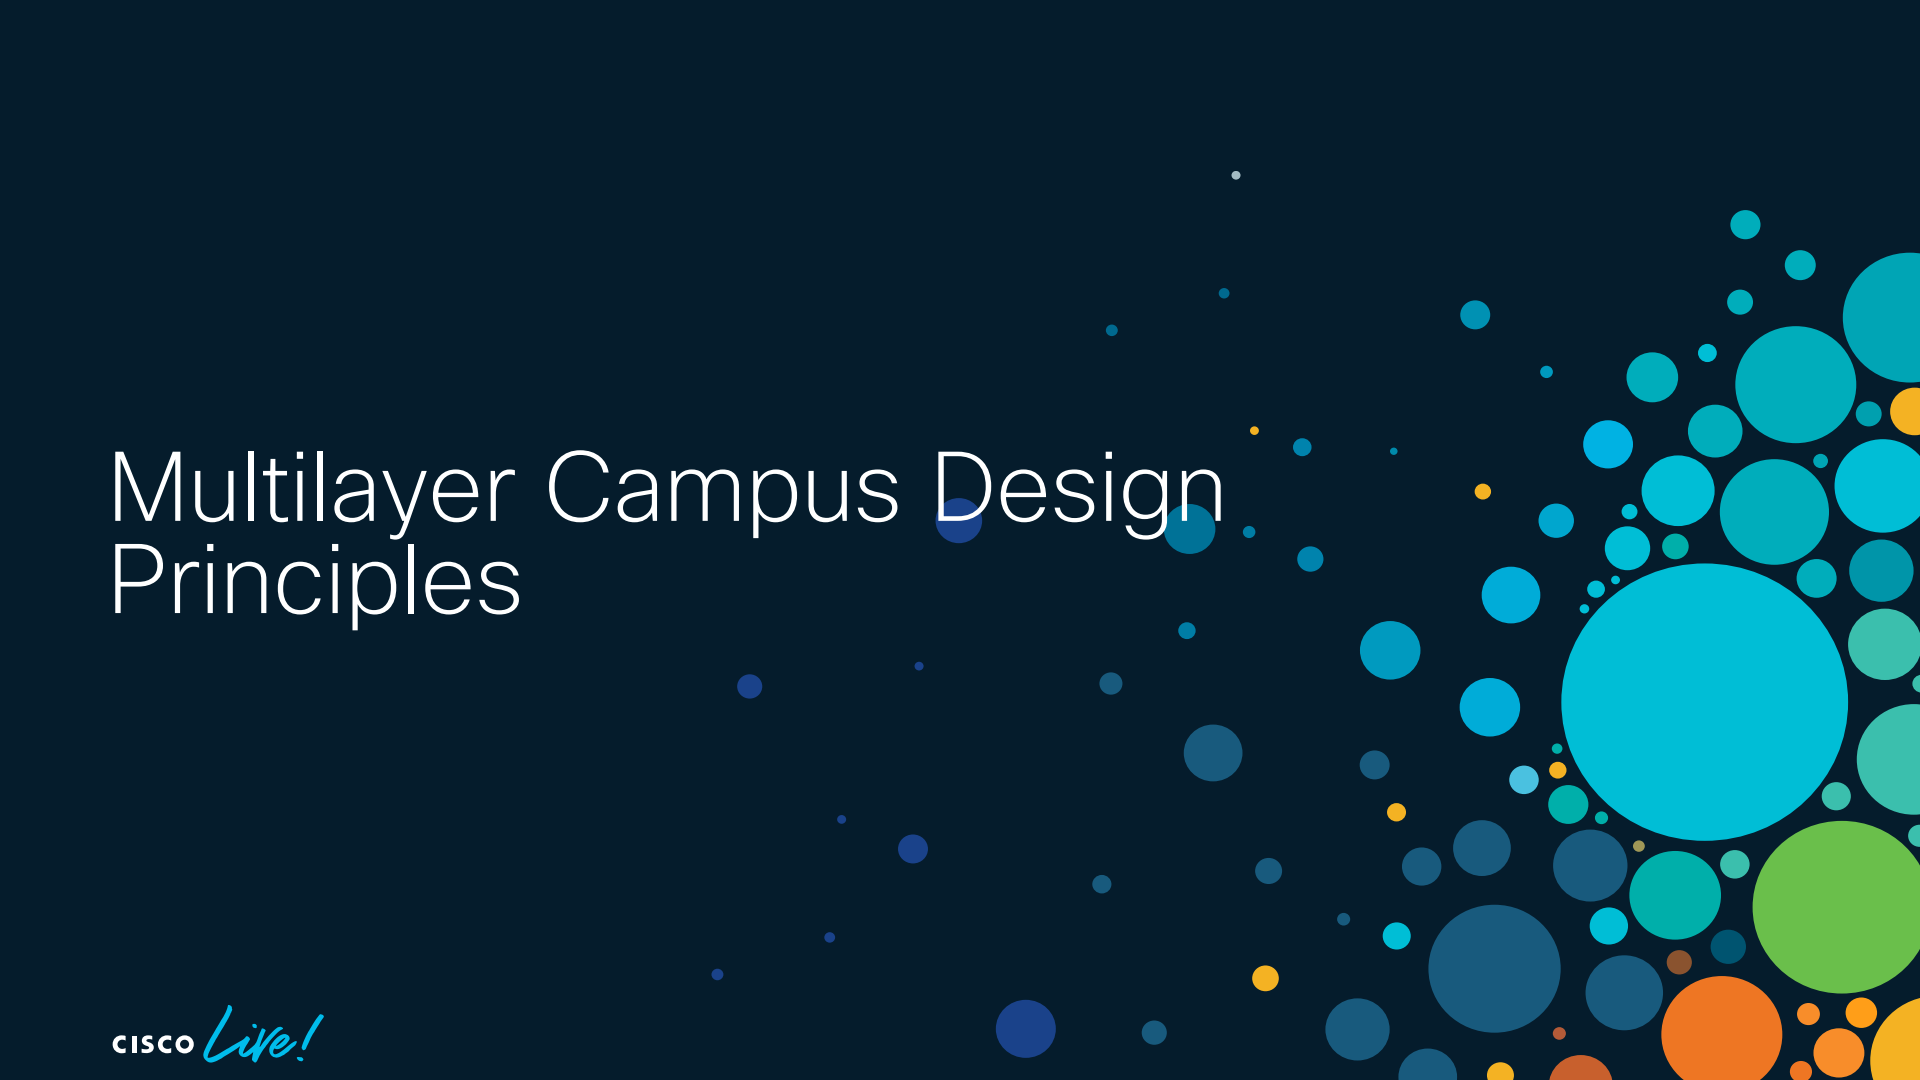

Building your own house…

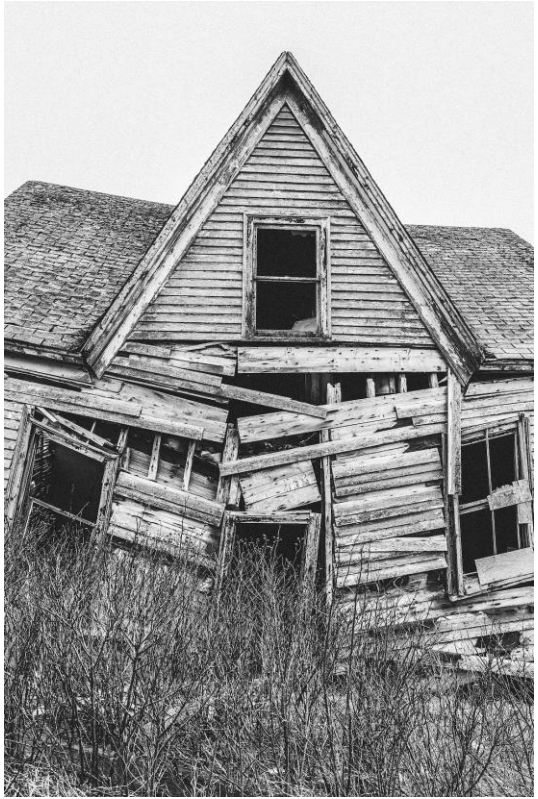

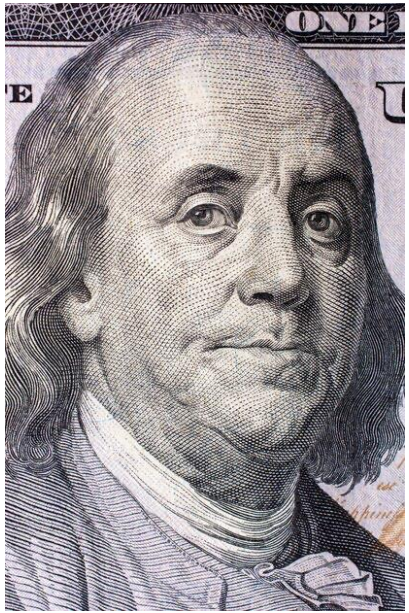

“.. If you fail to plan - you plan to fail”

Benjamin Franklin

High-Availability Campus Design

Access

Distribution

Core

Distribution

Access

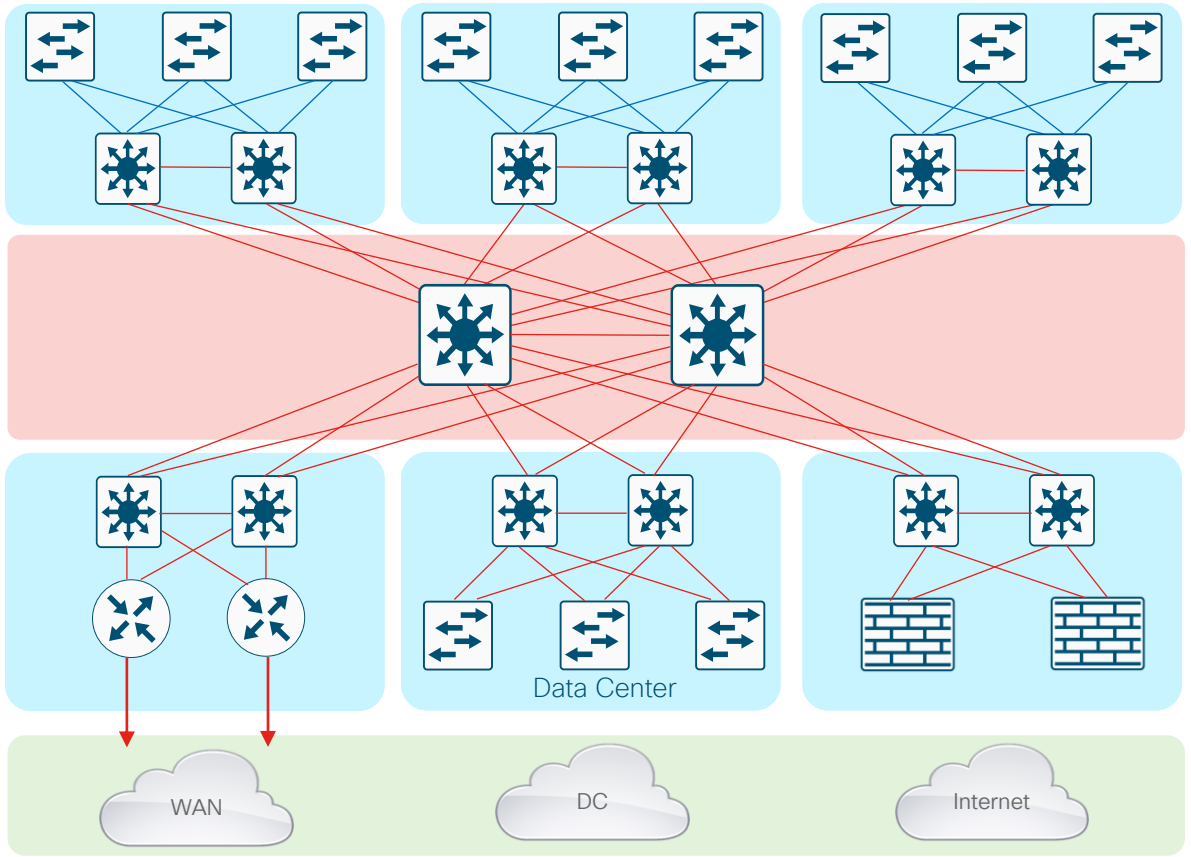

High-Availability Campus Design

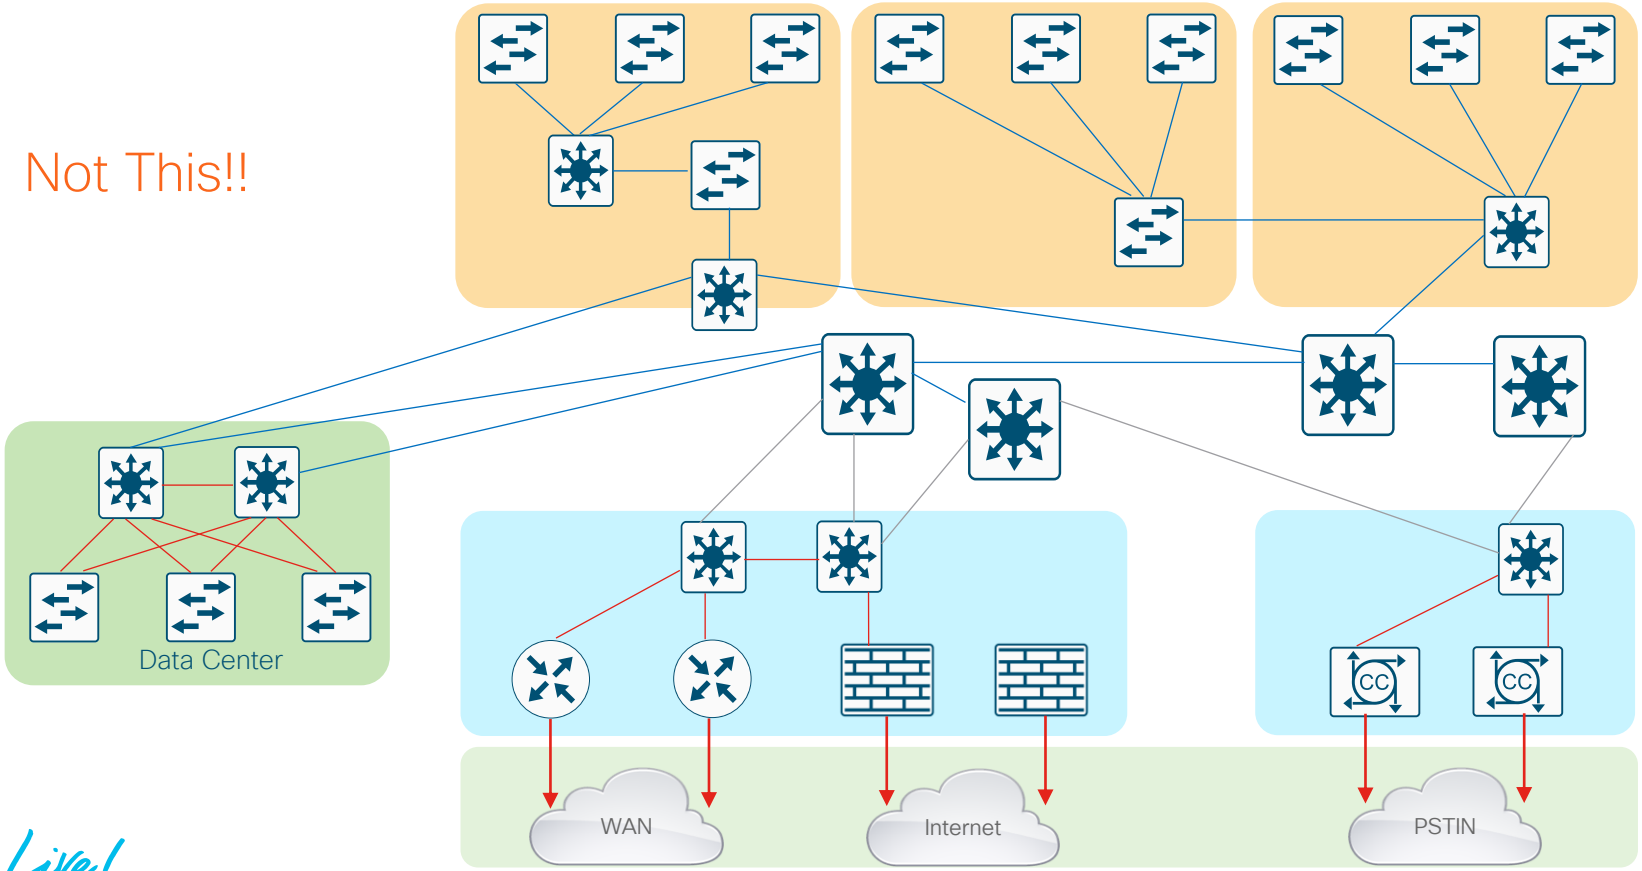

Hierarchical Network Design

Without a Rock Solid Foundation the Rest Doesn't Matter

Access

Distribution

Core

Distribution

Access

- 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

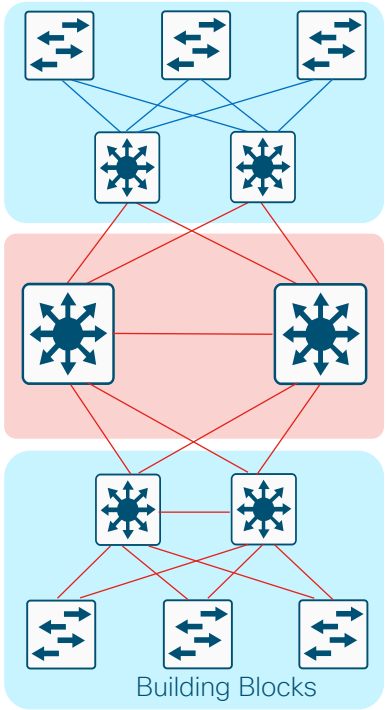

Access Layer

Feature Rich Environment

- It's not just about connectivity
- Layer 2/Layer 3 feature-rich environment; convergence, HA, security, QoS, IP multicast, etc.
- Intelligent network services: QoS, trust boundary, broadcast suppression, IGMP snooping
- Intelligent network services: PVST+, Rapid PVST+, EIGRP, OSPF, DTP, PAgP/LACP, UDLD, FlexLink, etc.
- Cisco Catalyst® integrated security features IBNS (802.1x), (CISF): port security, DHCP snooping, DAI, IPSG, etc.
- Automatic phone discovery, conditional trust boundary, power over Ethernet, auxiliary VLAN, etc.
- Spanning tree toolkit: PortFast, UplinkFast, BackboneFast, LoopGuard, BPDU Guard, BPDU Filter, RootGuard, etc.

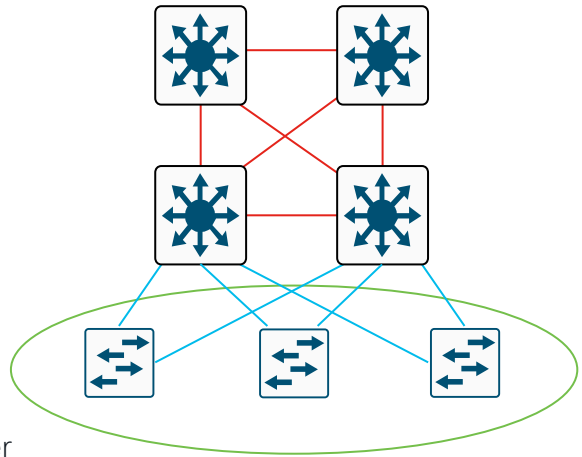

Core

Distribution

Access

Distribution Layer

Policy, Convergence, QoS and High Availability

- Availability, load balancing, QoS and provisioning are the important considerations at this layer
- Aggregates wiring closets (access layer) and uplinks to core
- Protects core from high-density peering and problems in access layer
- Route summarization, fast convergence, redundant path load sharing
- HSRP or GLBP to provide first-hop redundancy

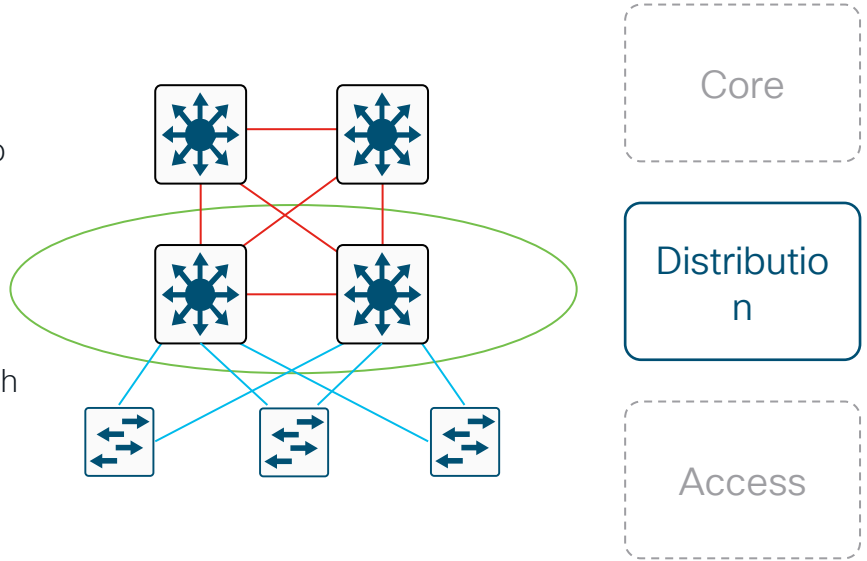

Core Layer

Scalability, High Availability, and Fast Convergence

- Backbone for the network—connects network building blocks
- Performance and stability vs. complexity— less is more in the core
- Aggregation point for distribution layer
- Separate core layer helps in scalability during future growth
- Keep the design technology-independent

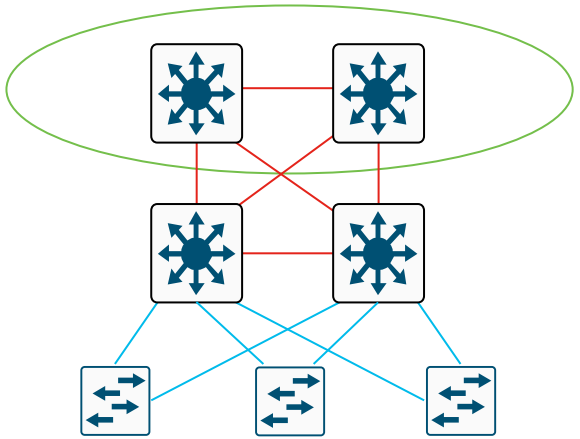

Core

Distribution

Access

Do I need a Core Layer?

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

- No Core
- Fully-meshed distribution layers
- Physical cabling requirement
- Routing complexity

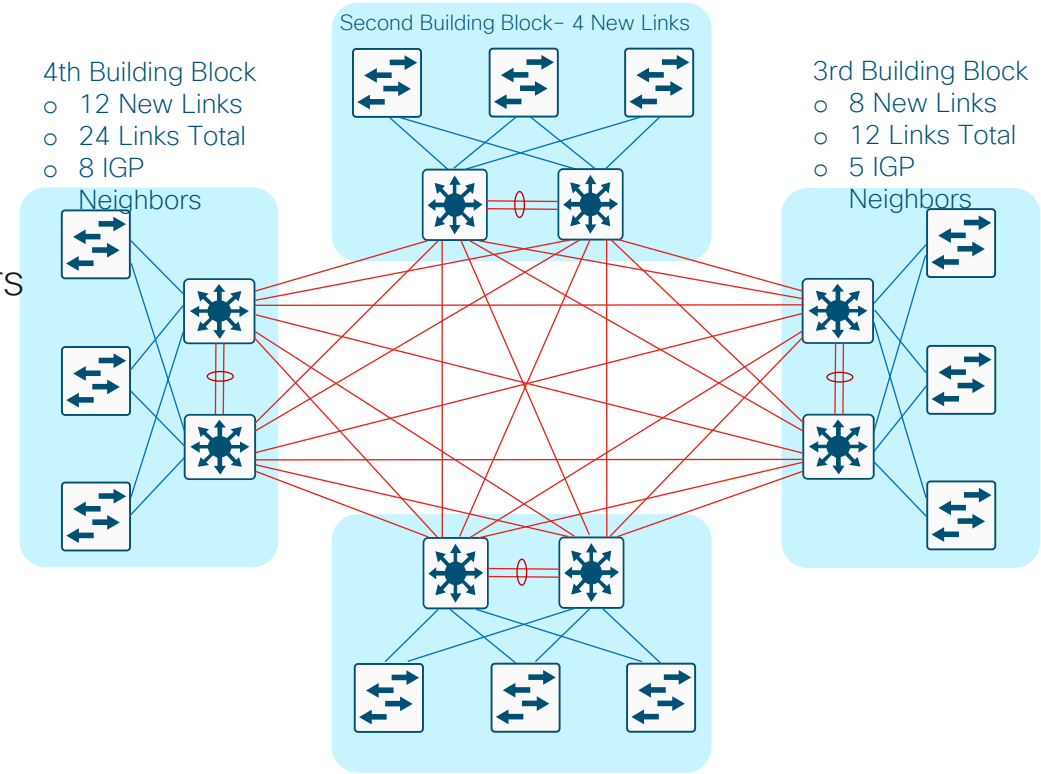

Do I need a Core Layer?

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

- Dedicated Core Switches
- Easier to add a module
- Fewer links in the core
- Easier bandwidth upgrade
- Routing protocol peering reduced
- Equal cost Layer 3 links for best convergence

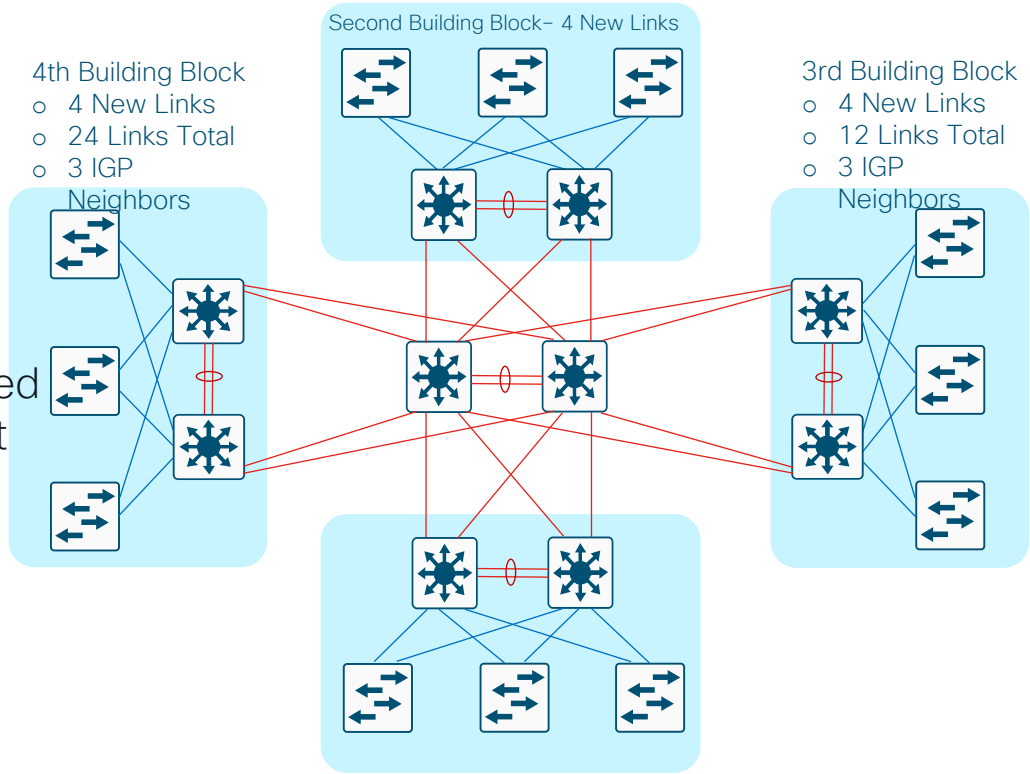

Design Alternatives Come Within a Building (or Distribution) Block

Access

Distribution

Core

Distribution

Access

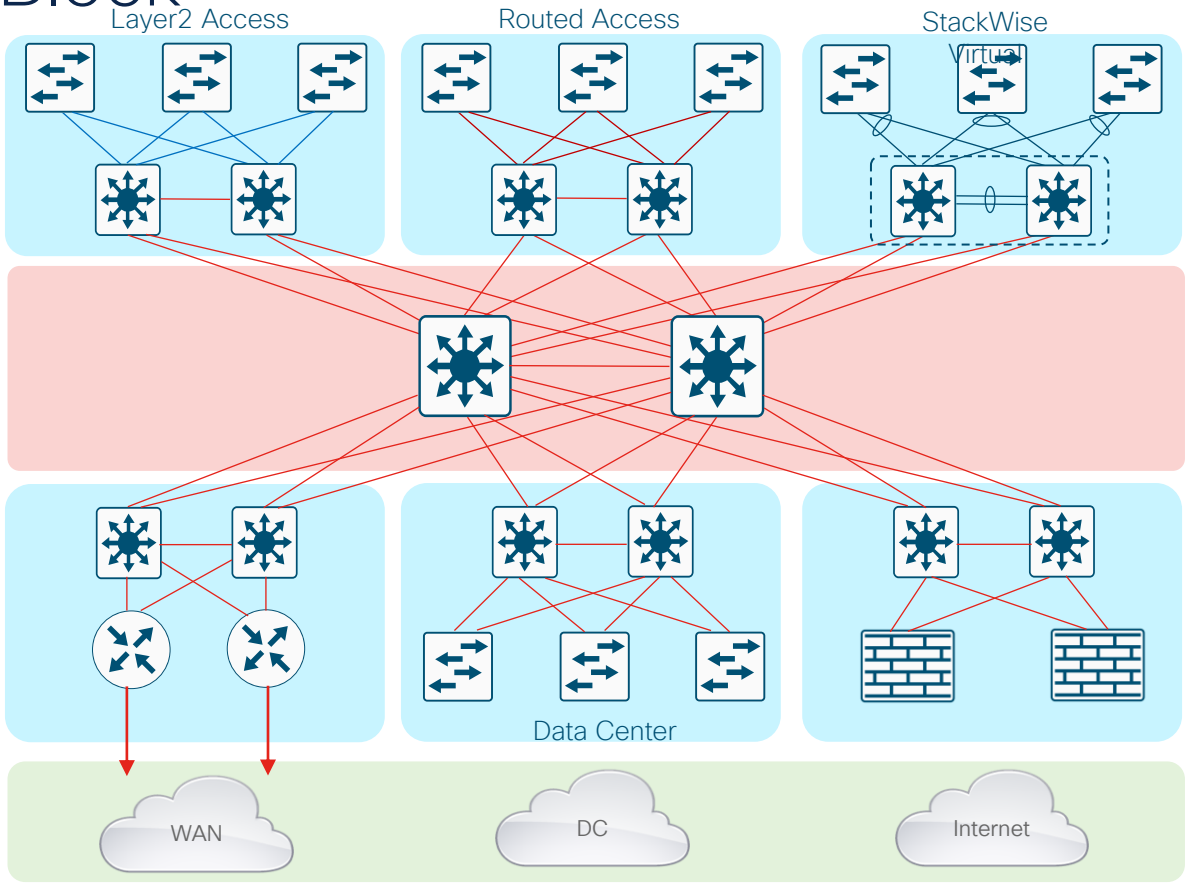

Layer 2 Distribution Interconnection

Layer 2 Access—No VLANs Span Access Layer

Core

Distribution

Access

- Summarize routes towards core
- STP Root and HSRP primary tuning or
- GLBP to load balance on uplinks
- Set trunk mode on/no-negotiate
- Set port host on access layer ports:
  - Disable trunking
  - Disable Ether Channel
  - Enable PortFast
- RootGuard or BPDU-Guard
- Use security features

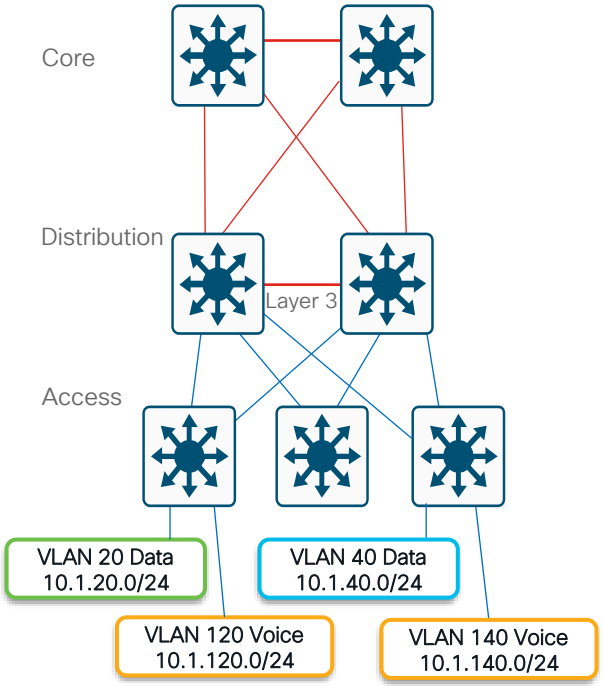

Layer 3 Distribution Interconnection

Layer 2 Access – Some VLANs Span Access Layer

Core

Distribution

Access

- Summarize routes towards core
- STP Root and HSRP primary or GLBP and STP port cost tuning to load balance on uplinks
- Set trunk mode on/no-negotiate
- RootGuard on downlinks
- LoopGuard on uplinks
- Set port host on access layer ports:
  - Disable trunking
  - Disable Ether Channel
  - Enable PortFast
- RootGuard or BPDU-Guard
- Use security features

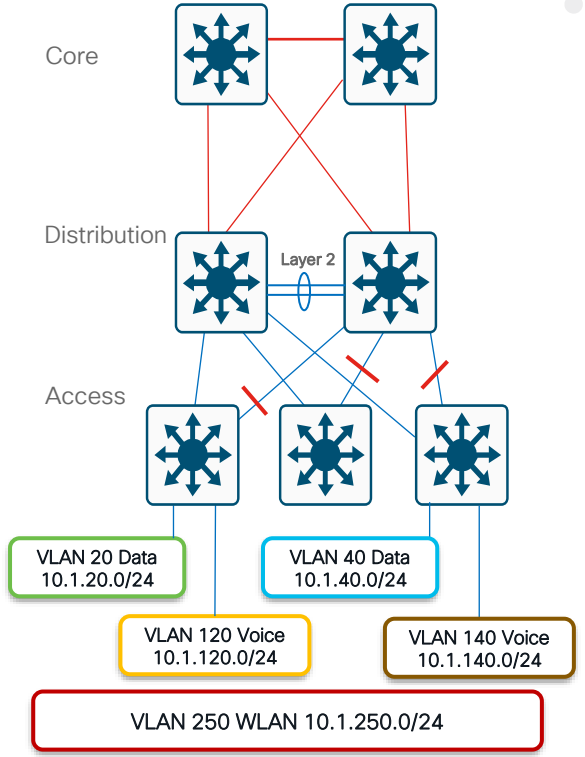

StackWise Virtual and Virtual Stacking

L2 without a STP Liability

Core

Distribution

Access

- Summarize routes towards core
- Limit redundant IGP peering
- Set trunk mode on/no-negotiate
- MUST Ether Channel else blocked ports
- Set port host on access Layer ports:
  - Disable trunking
  - Disable Ether Channel
  - Enable PortFast
- RootGuard or BPDU-Guard
- Use security features

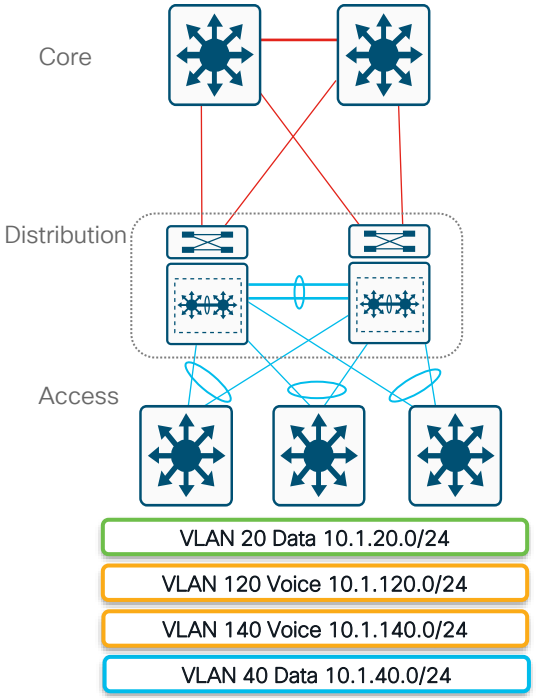

Routed Access and Virtual Switching System

Evolutions of and Improvements to Existing Designs

Advantages:

- Ease of implementation, less to get right
  - No matching of STP/HSRP/GLBP priority
  - No L2/L3 Multicast topology inconsistencies
- Single Control Plane and well-known toolset
  - traceroute, show ip route, show ip eigrp neighbor, etc.
- Catalyst 9k platform fully supports L3 switching
- EIGRP converges in < 200 msec
- OSPF with sub-second tuning converges in < 200 msec
- RPVST+ convergence times dependent on GLBP / HSRP tuning

Considerations:

- Do you have any L2 VLAN adjacency requirements between access switches
- IP addressing – Do you have enough address space and the allocation plan to support a routed access design

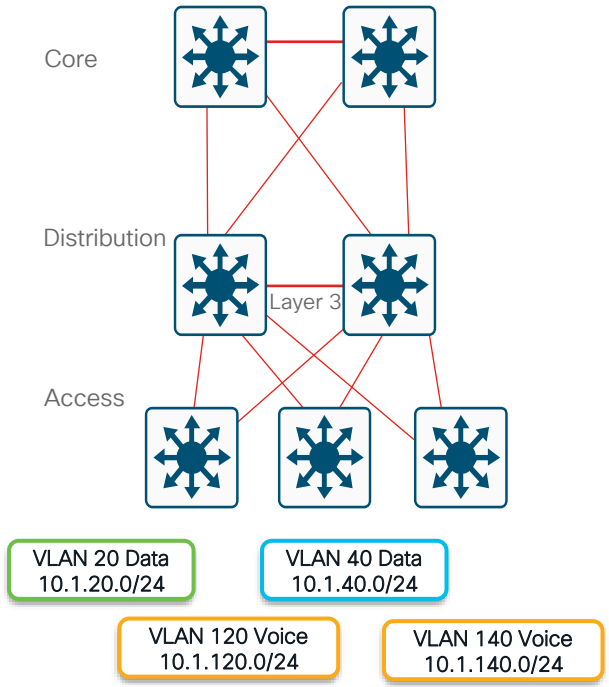

Campus Fabric – The Foundation for SDA

Architecture for the Digital Enterprise

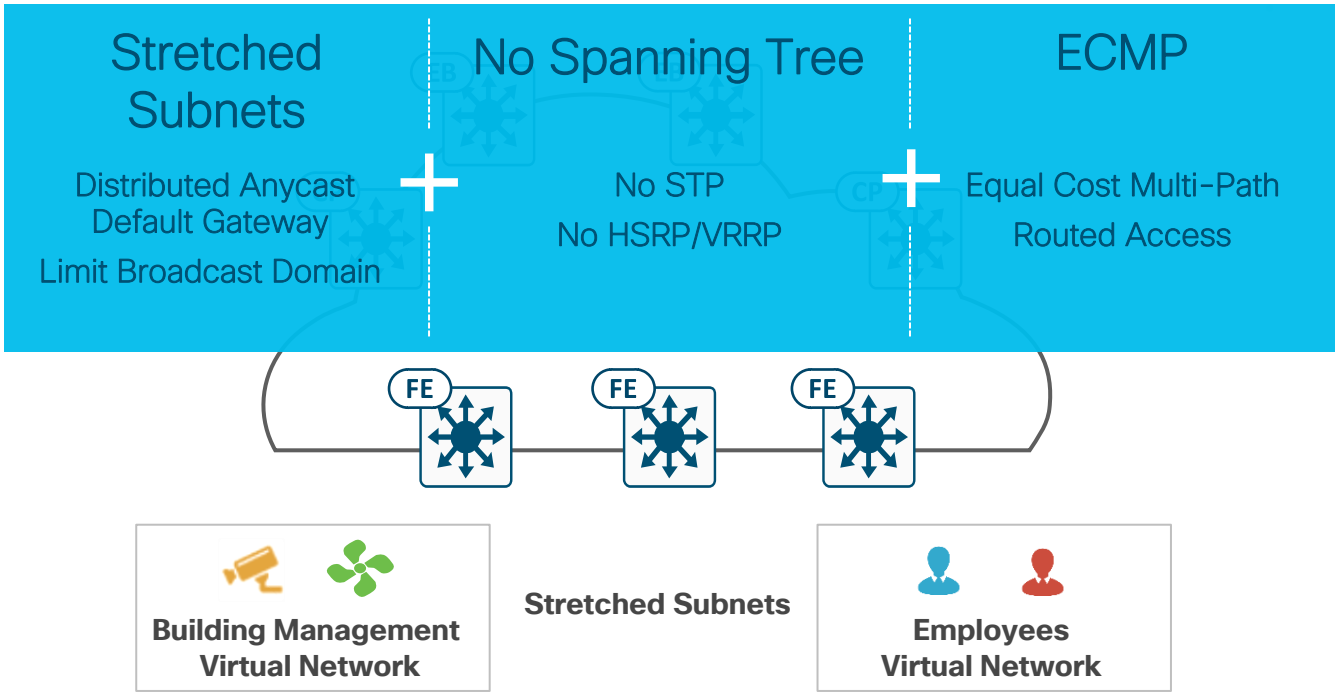

Campus Fabric – The Foundation for SDA

Architecture for the Digital Enterprise

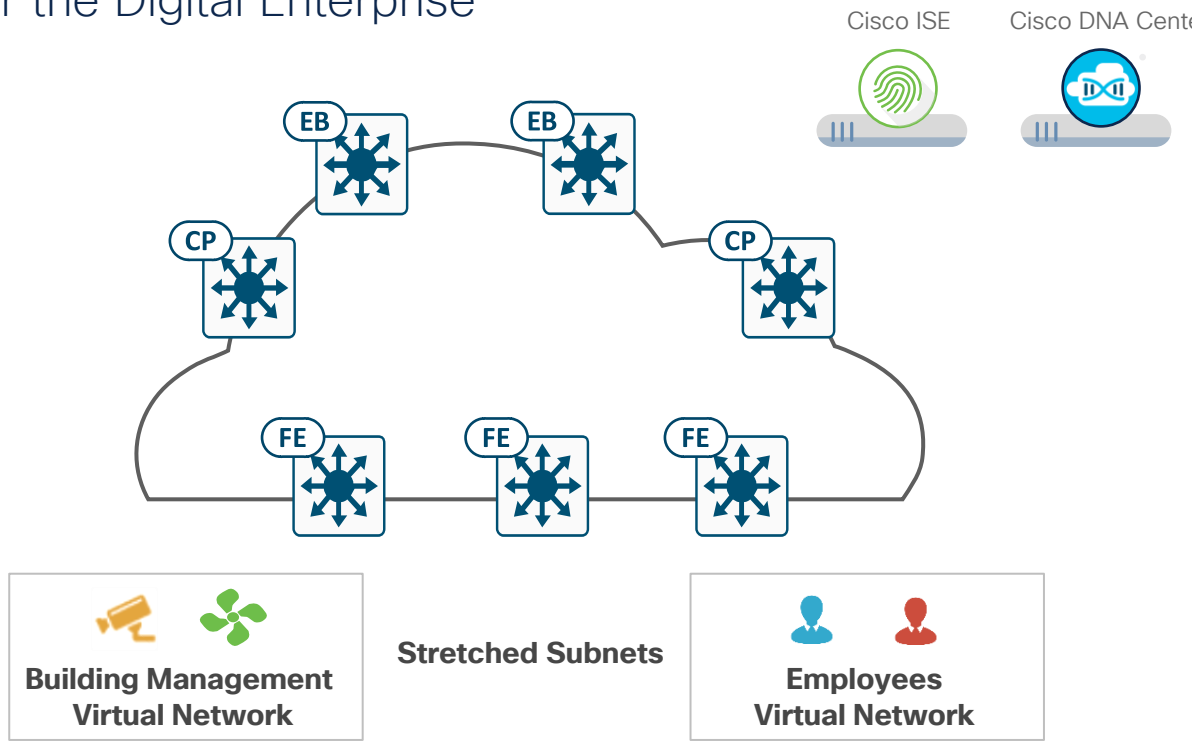

Foundation services

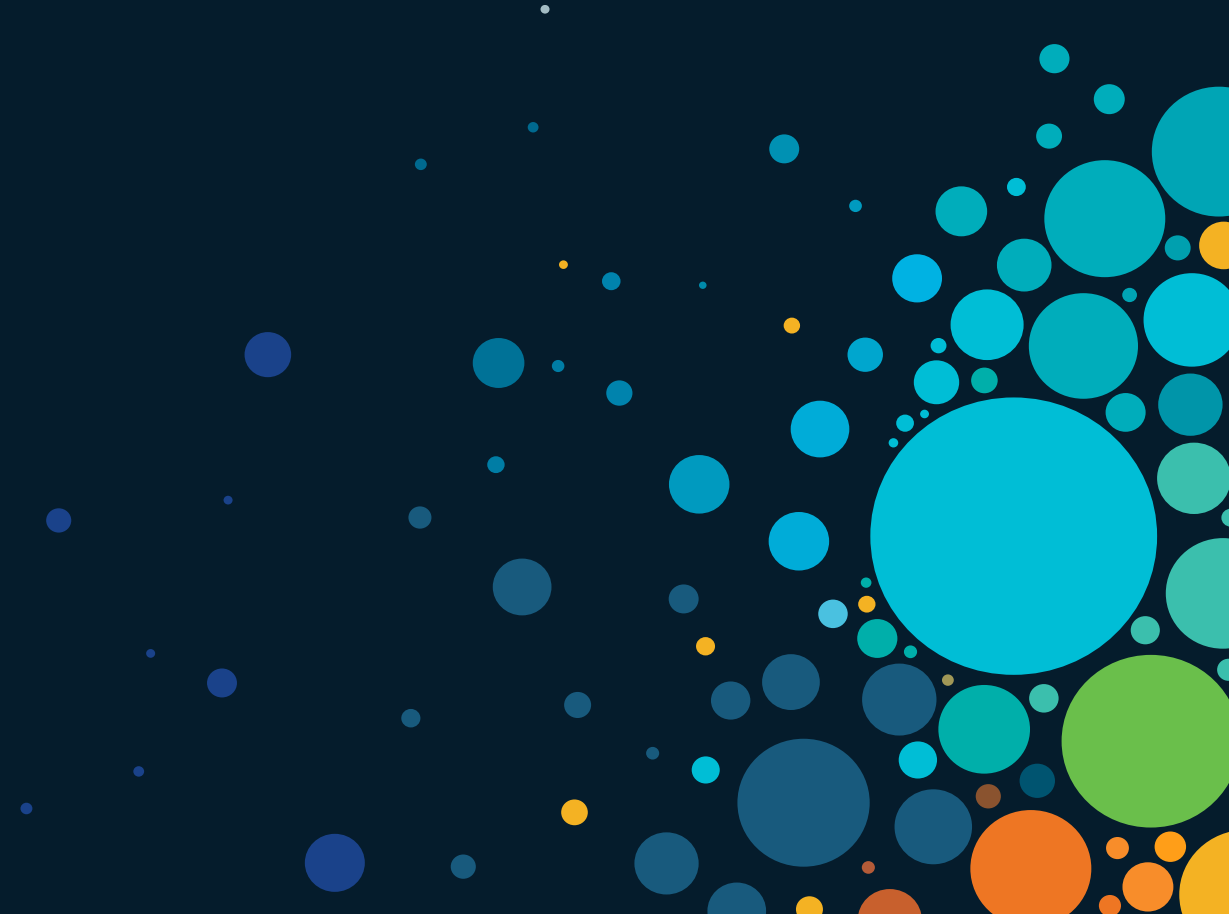

Foundation Services

- Layer 1 physical things
- Layer 2 redundancy
  - STP
  - Trunks
  - UDLD
- Layer 3 routing protocols
  - Ether Channels
  - BFD
  - FHRP

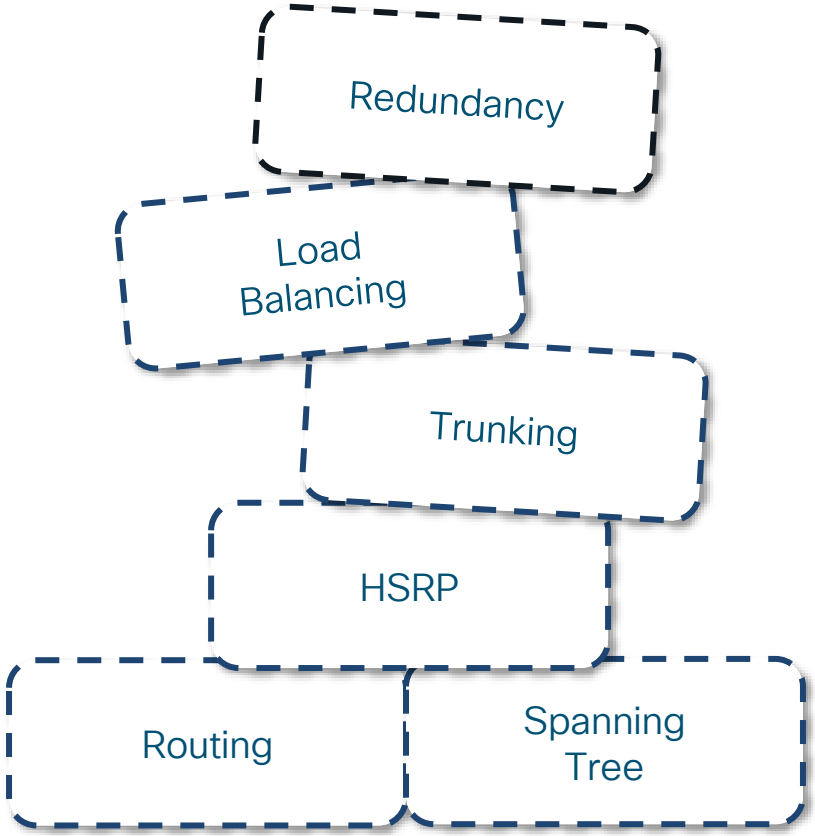

Best Practices - Layer 1 Physical Things

- Review Link Debounce and Carrier-Delay
- Use point-to-point interconnections - no L2 aggregation points between nodes
- Use configuration on the physical interface not VLAN/SVI when possible

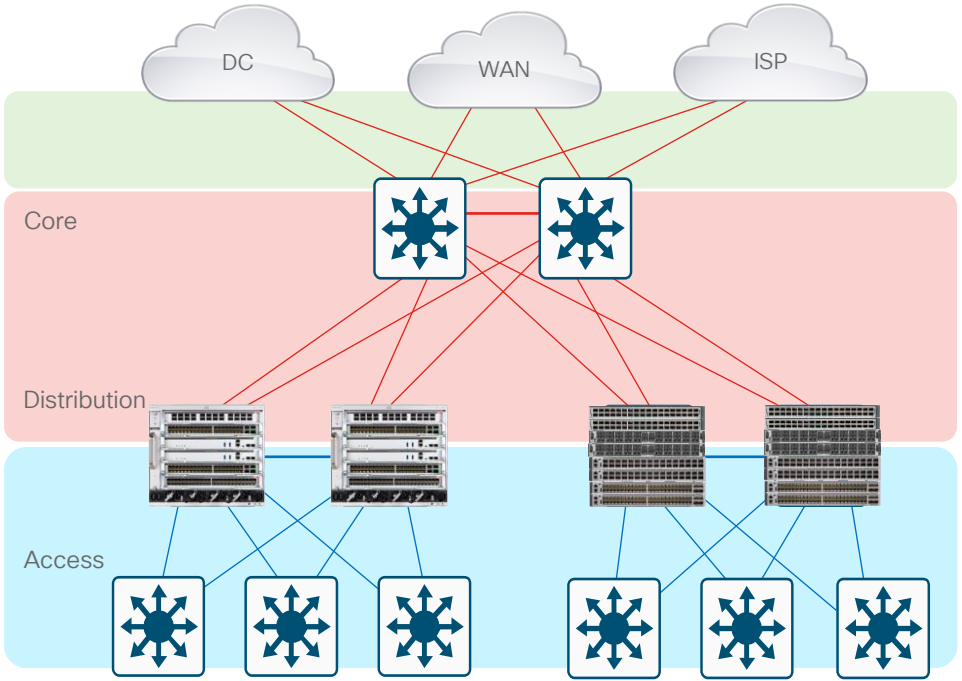

Link Debounce and Carrier-Delay

- When tuning the campus for optimal convergence, it is important to review the status of the link debounce and carrier delay configuration
- By default GigE and 10GigE+ interfaces operate with a 10 msec debounce timer which provides for optimal link failure detection
- In the current Cisco IOS levels, the default behavior for Catalyst switches is to use a default value of 0 msec on all Ethernet interfaces for the carrier-delay.
- It is still recommended as best practice to hard code the carrier-delay value on critical interfaces with a value of 0 msec to ensure the desired behavior.

Can be adjusted on **Cat9500 & Cat9600**

```
C9500-32QC-1-4#show interfaces debounce

Port          Debounce time    Value(ms)
Fo1/0/1       disable
Fo1/0/2       disable
Fo1/0/3       disable
Fo1/0/4       disable
Fo1/0/5       disable
Fo1/0/6       disable
```

```
interface GigabitEthernet1/1
 description Uplink to Distribution 1
 dampening
 ip address 10.120.0.205 255.255.255.254
 ip pim sparse-mode
 ip ospf dead-interval minimal hello- multiplier 4
 ip ospf priority 0
 logging event link-status
 load-interval 30
 carrier-delay msec 0
<snip>
```

Redundancy and Protocol Interaction

Layer 2 and 3 – Why Use Routed Interfaces

Configuring L3 routed interfaces provides for faster convergence than an L2 switch port with an associated L3 SVI

1. Link Down
2. Interface Down
3. 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
```

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

~ 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
```

Best Practices - Spanning Tree Configuration

- Only span VLAN across multiple access layer switches when you have to!
- Use rapid RSTP for best convergence
- Required to protect against user side loops
- Required to protect against operational accidents (misconfiguration or hardware failure)
- Take advantage of the spanning tree toolkit

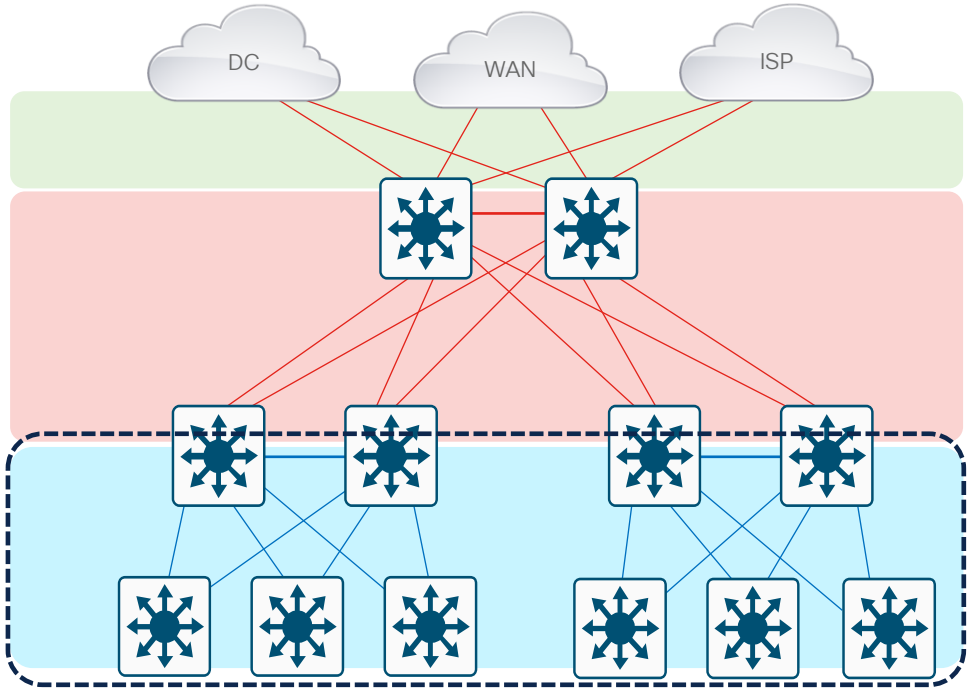

Multilayer Network Design

Layer 2 Access with Layer 3 Distribution

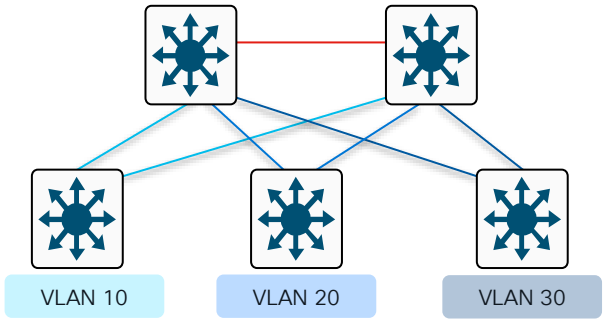

- Each access switch has unique VLANs
- No Layer 2 loops
- Layer 3 link between distribution
- No blocked links

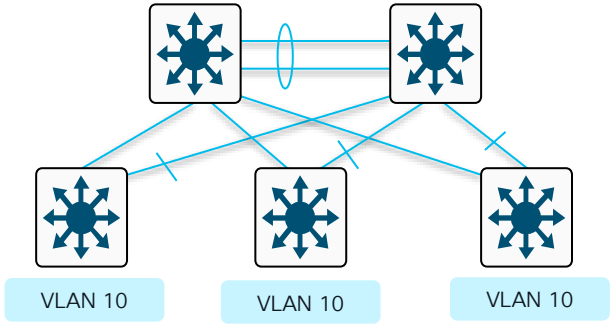

- At least some VLANs span multiple access switches
- Layer 2 loops
- Layer 2 and 3 running over link between distribution
- Blocked links

Optimizing L2 Convergence

PVST+, Rapid PVST+ or MST

- Rapid-PVST+ greatly improves the restoration times for any VLAN that requires a topology convergence due to link UP
- Rapid-PVST+ also greatly improves convergence time over backbone fast for any indirect link failures
- PVST+ (802.1d)
  - Traditional spanning tree implementation
- Rapid PVST+ (802.1w)
  - Scales to large size (~10,000 logical ports)
  - Easy to implement, proven, scales
- MST (802.1s)
  - Permits very large scale STP implementations (~30,000 logical ports)

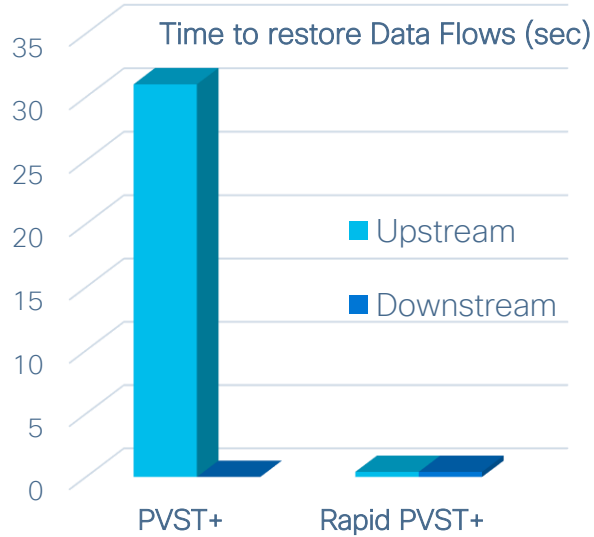

Layer 2 Hardening

Spanning Tree Should Behave the Way You Expect

- Place the root where you want it
  Root primary/secondary macro
- The root bridge should stay where you put it
  - RootGuard
  - LoopGuard
  - UplinkFast
  - UDLD
- Only end-station traffic should be seen on an edge port
  - BPDU Guard
  - RootGuard
  - PortFast
  - Port-security

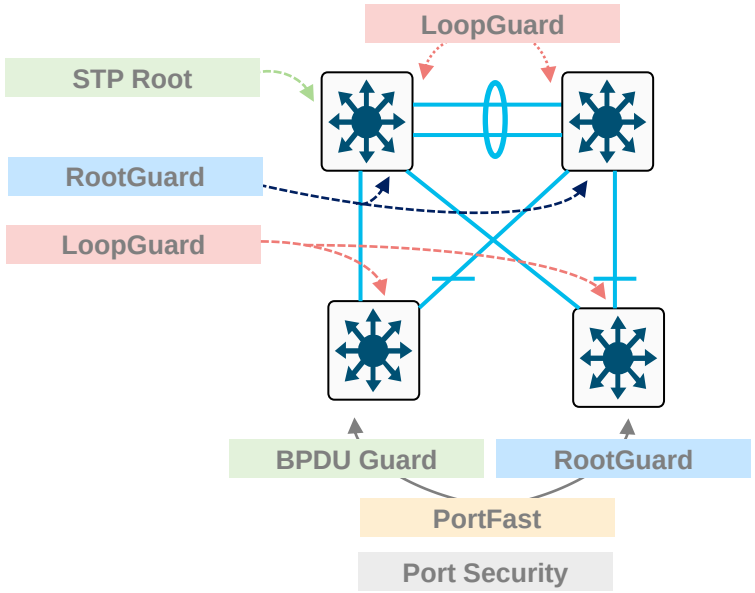

Best Practices – Trunk Configuration

- 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
- Manually prune all VLANS except those needed
- Disable on host ports

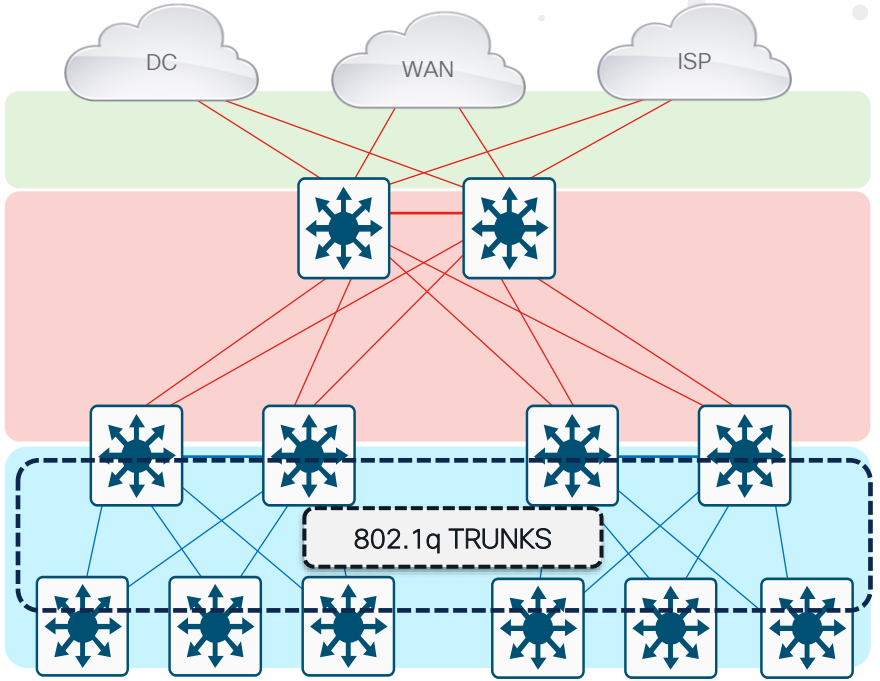

Optimizing Convergence: Trunk Tuning

Trunk Auto/Desirable Takes Some Time

- DTP negotiation tuning improves link up convergence time
  - IOS(config-if)# switchport mode trunk
  - IOS(config-if)# switchport nonegotiate

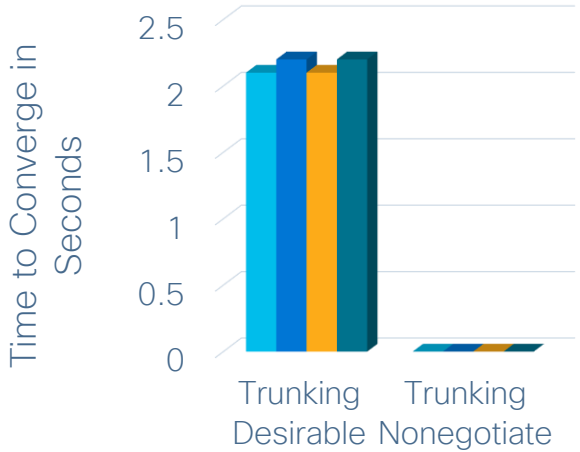

Best practices – UDLD Configuration

- Typically deployed on any fiber optic interconnection
- Use UDLD aggressive mode for most aggressive protection
- Turn on in global configuration to avoid operational error/misses

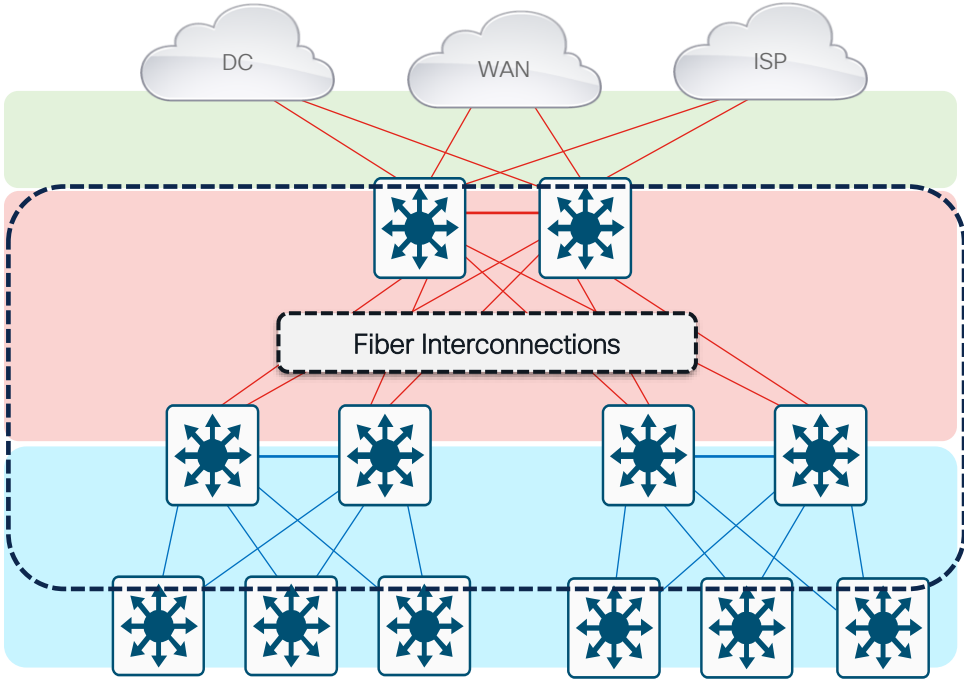

UDLD Aggressive and UDLD Normal

- Timers are the same—15-second hellos by default
- Aggressive Mode—after aging on a previously bi-directional link—tries eight times (once per second) to reestablish connection then err-disables port
- UDLD—Normal Mode—only err-disable the end where UDLD detected other end just sees the link go down
- UDLD—Aggressive—err-disable both ends of the connection due to err-disable when aging and re-establishment of UDLD communication fails

Best Practices – Ether Channel Configuration

- Typically deployed in distribution to core, and core to core interconnections
- Used to provide link redundancy—while reducing peering complexity
- Tune L3/L4 load balancing hash to achieve maximum utilization of channel members
- Deploy in powers of two (two, four, or eight)
- 802.3ad LACP for interop if you need it

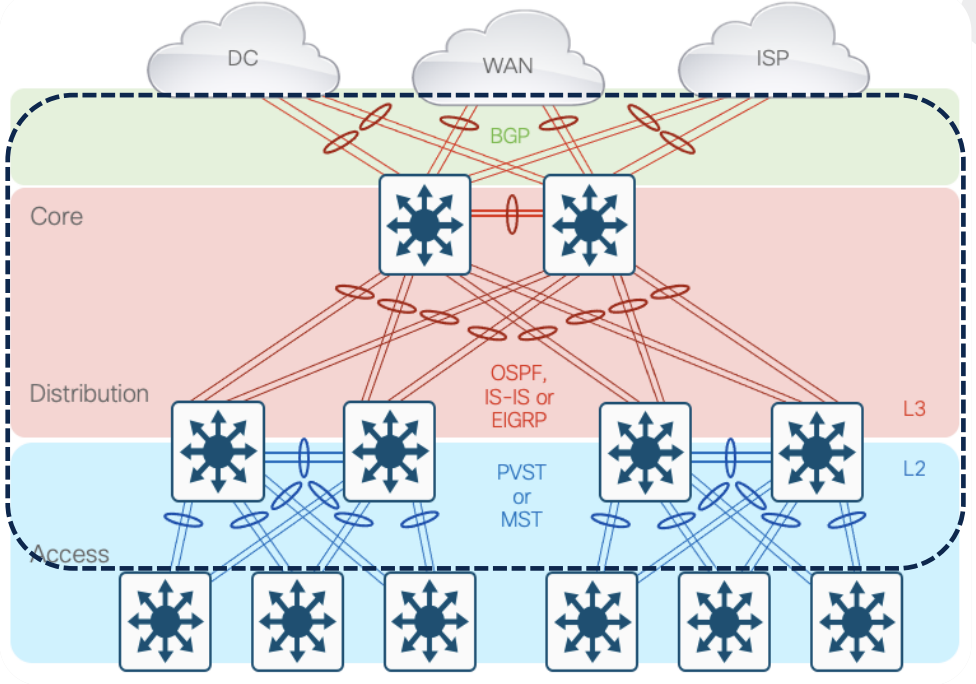

Ether Channel load balancing

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 information, or you can add an additional level of load balancing to the process by adding the L4 TCP/IP port information as an 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
```

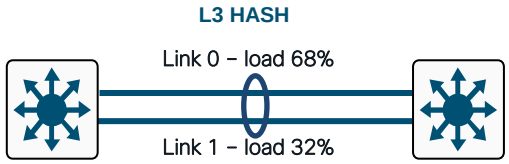

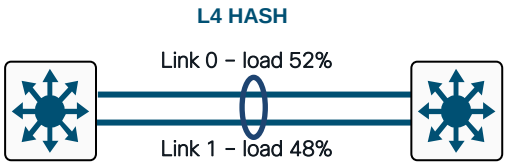

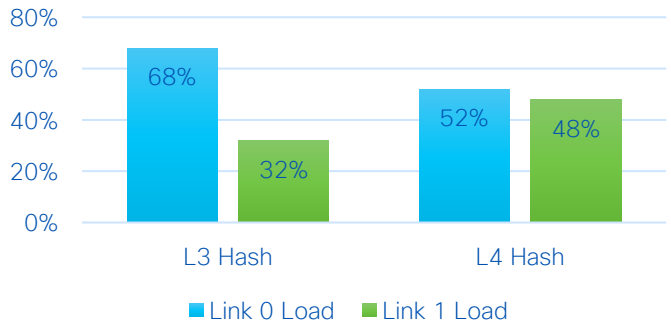

EtherChannels

Reduce Complexity/Peer Relationships

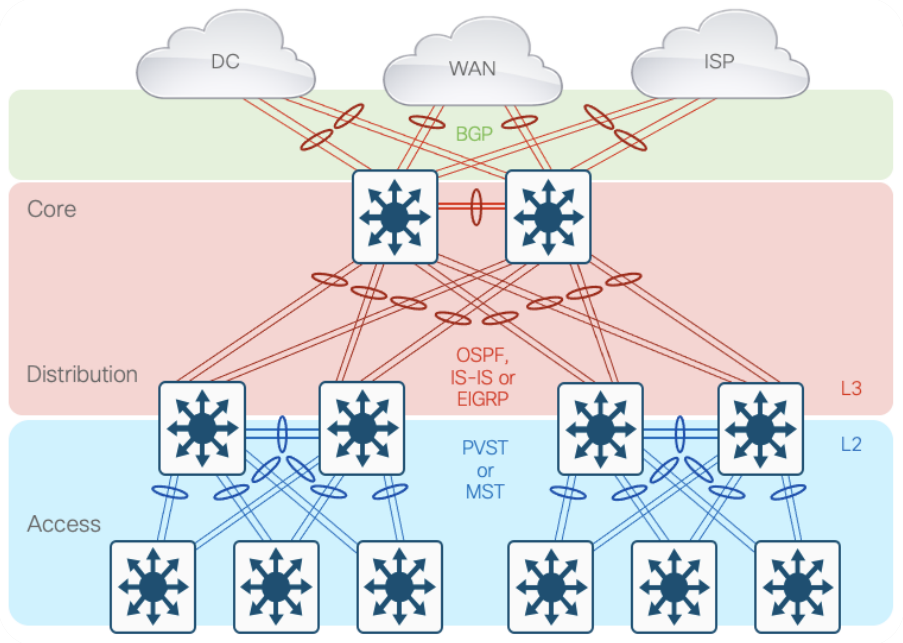

- More links = more routing peer relationships and associated overhead
- EtherChannels allow you to reduce peers by creating single logical interface to peer over
- On single link failure in a bundle
  - OSPF running on a Cisco IOS-based switch will reduce link cost and reroute traffic
  - EIGRP may not change link cost and may overload remaining links

EtherChannels

1G/10G/20G/40G/100G How do you aggregate it ?

Typical 4:1 Data Over-Subscription

Typical 20:1 Data Over-Subscription

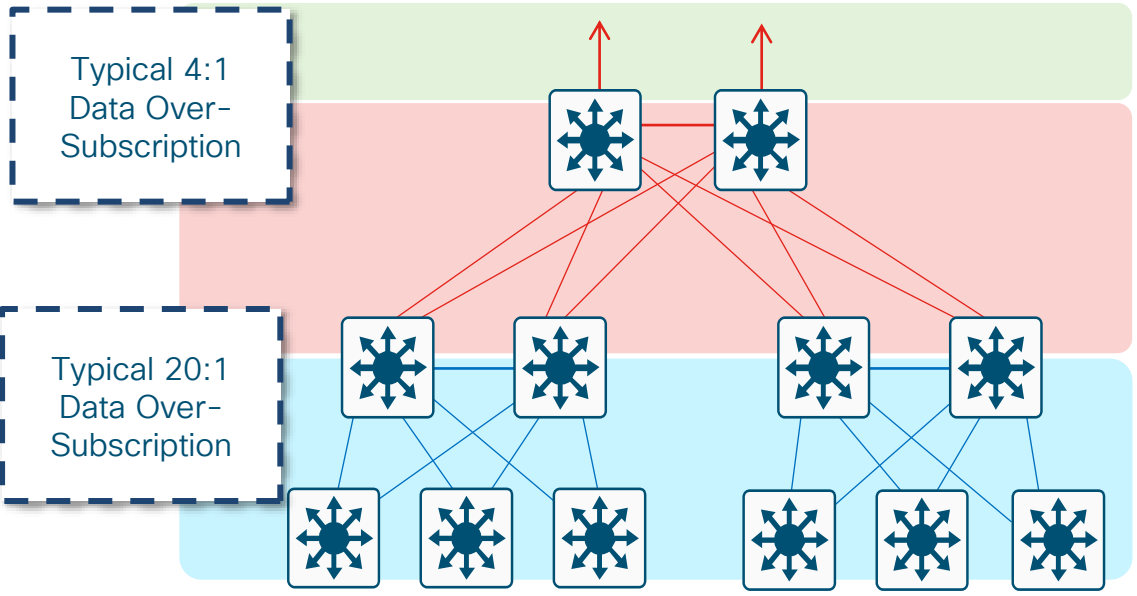

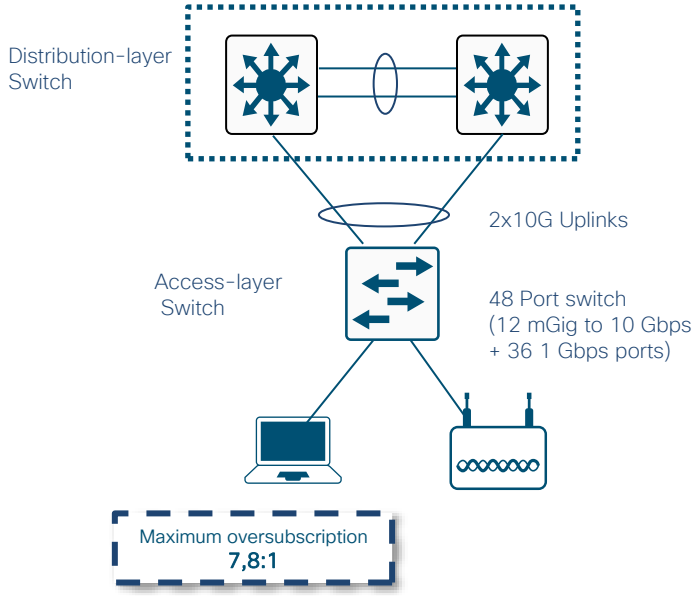

Best Practices

Layer 3 Routing Protocols

- Typically deployed in distribution to core, and core-to-core interconnections
- Used to quickly reroute around failed node/links while providing load balancing over redundant paths
- Build triangles not squares for deterministic convergence
- Only peer on links that you intend to use as transit

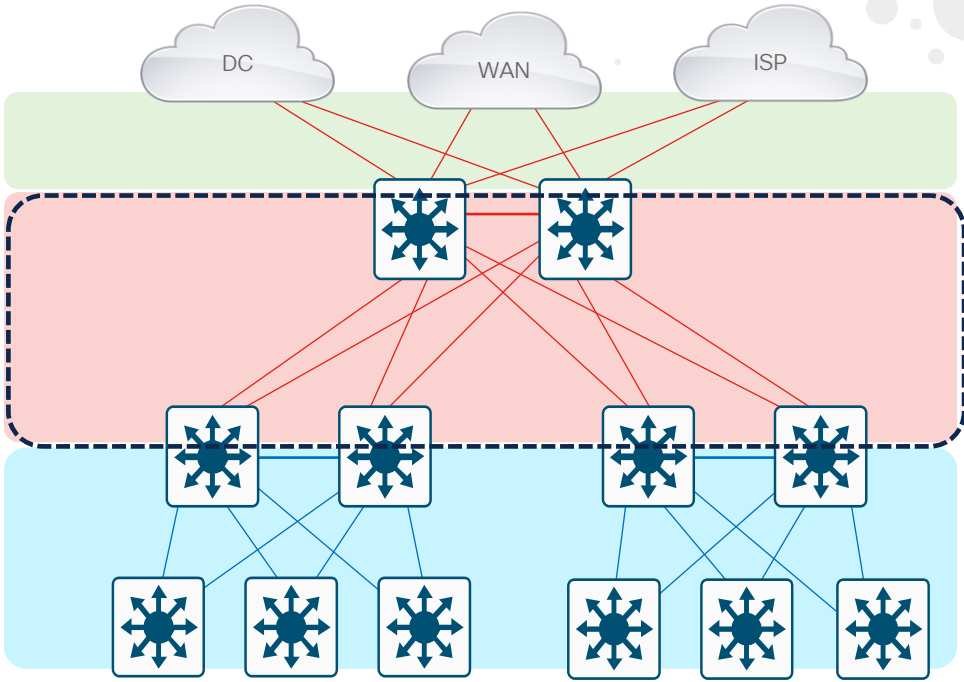

Best Practice - Build Triangles not Squares

Deterministic vs. Non-Deterministic

Squares: Link/Box Failure Requires Routing Protocol Convergence

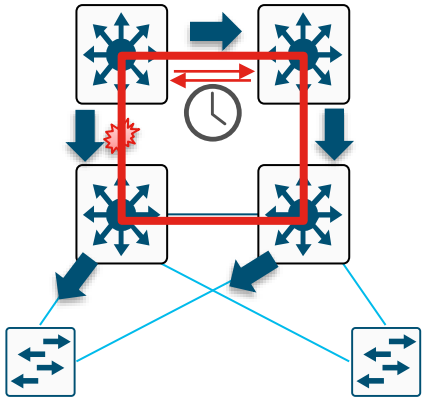

Triangles: Link/Box Failure Does not Require Routing Protocol Convergence

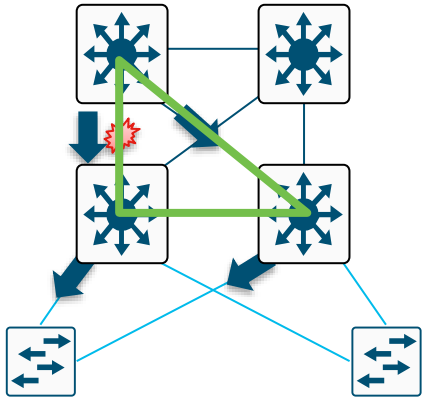

- Layer 3 redundant equal cost links support fast convergence
- Hardware based—fast recovery to remaining path
- Convergence is extremely fast (dual equal-cost paths: no need for OSPF or EIGRP to recalculate a new path)

Best Practice – Passive Interfaces for IGP

Limit IGP Peering Through the Access Layer

- Limit unnecessary peering using passive interface:
  - Four VLANs per wiring closet
  - 12 adjacencies total
  - Memory and CPU requirements increase with no real benefit
  - Creates overhead for IGP

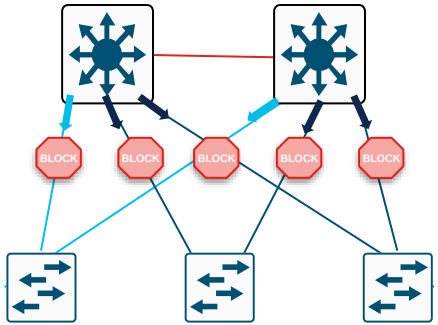


```
OSPF Example:
Router(config)#router ospf 1
Router(config-router)#passive-interfaceVlan 99

Router(config)#router ospf 1
Router(config-router)#passive-interface default
Router(config-router)#no passive-interface Vlan 99
```

```
EIGRP Example:
Router(config)#router eigrp 1
Router(config-router)#passive-interfaceVlan 99

Router(config)#router eigrp 1
Router(config-router)#passive-interface default
Router(config-router)#no passive-interface Vlan 99
```

Why You Want to Summarize at the Distribution

Limit EIGRP Queries and OSPF LSA Propagation

- 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 this reroute

EIGRP Example:

```
interface Port-channel1
 description to Core#1
 ip address 10.122.0.34 255.255.255.252
 ip hello-interval eigrp 100 1
 ip hold-time eigrp 100 3
```

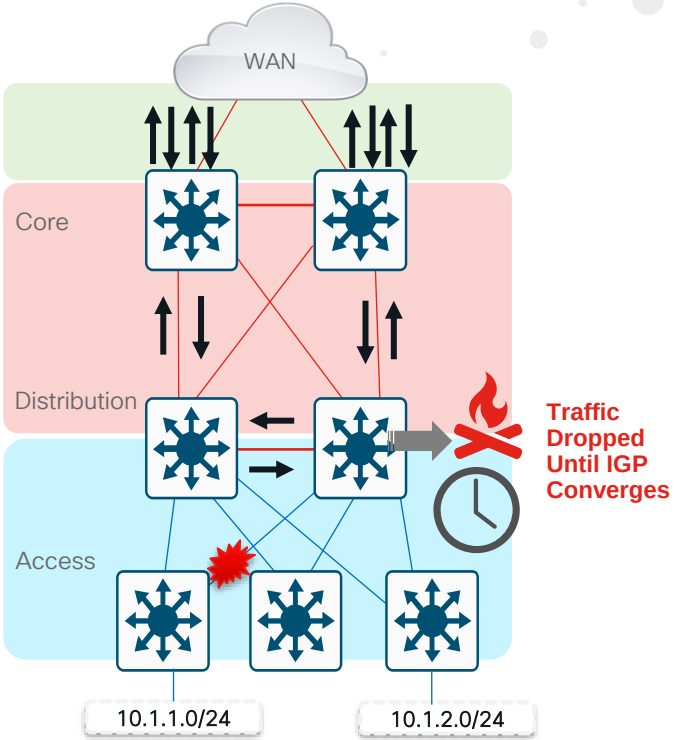

Why You Want to Summarize at the Distribution

Limit EIGRP Queries and OSPF LSA Propagation

- 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 this reroute

EIGRP Example:

```
interface Port-channel1
 description to Core#1
 ip address 10.122.0.34 255.255.255.252
 ip hello-interval eigrp 100 1
 ip hold-time eigrp 100 3
 ip summary-address eigrp 100 10.1.0.0 255.255.0.0 5
```

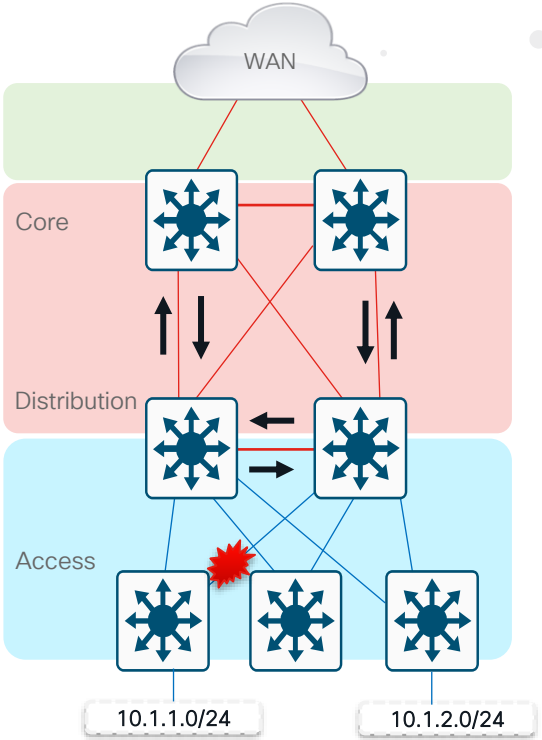

Bidirectional Forwarding Detection (BFD)

- Detect faults between 2 routers
  - Fast (reaction time in milliseconds)
  - Let the upper routing protocols (ISIS, BGP, OSFP, Static) that a link is down faster than the DEAD timer of that RP realize it
  - Works on directly connected routers, as well as routers separated by a L2 cloud (Metro Ethernet, MPLS,VPLS, Pseudowire, …)
  - 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

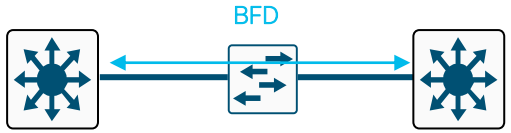


```
interface Gig1/0/1
 ip address 1.1.1.1 255.255.255.0
 bfd interval 300 min_rx 300 multiplier 3
 ip ospf 1 area 0

router ospf 1
 bfd all-interfaces
```

First Hop Redundancy with HSRP

- A group of routers function as one virtual router by sharing one virtual IP address and one virtual MAC address
- One (active) router performs packet forwarding for local hosts
- The rest of the routers provide hot standby in case the active router fails
- Standby routers stay idle as far as packet forwarding from the client side is concerned

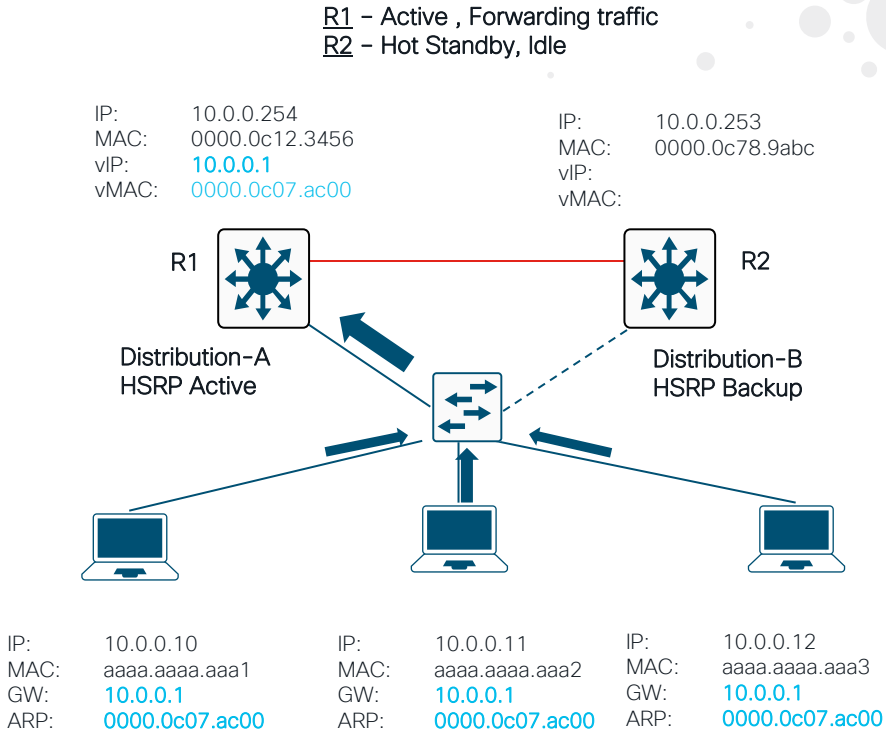

First Hop Redundancy with Load Balancing

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

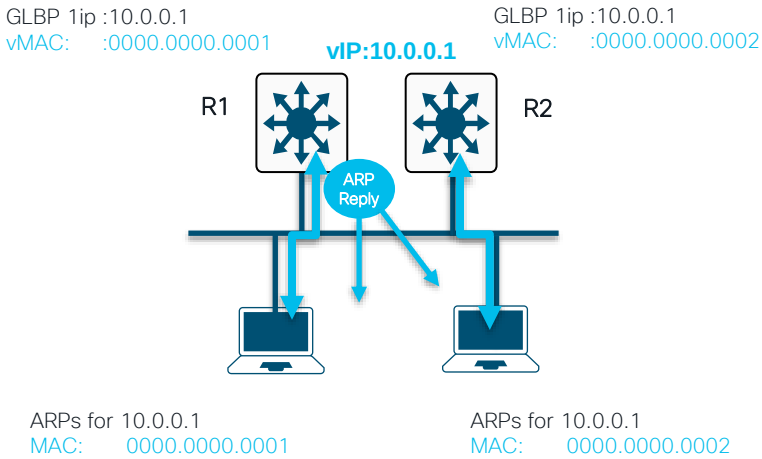

Optimizing Convergence: VRRP, HSRP, GLBP

Mean, Max, and Min—Are There Differences?

- HSRP has sub-second timers; however all flows go through same HSRP peer so there is no difference between mean, max, and min
- GLBP has sub-second timers and distributes the load amongst the GLBP peers; so 50% of the clients are not affected by an uplink failure

If You Span VLANS, Tuning Required

By Default, Half the Traffic Will Take a Two-Hop L2 Path

- Distribution switches act as default gateway
- Blocked uplink caused traffic to take less than optimal path

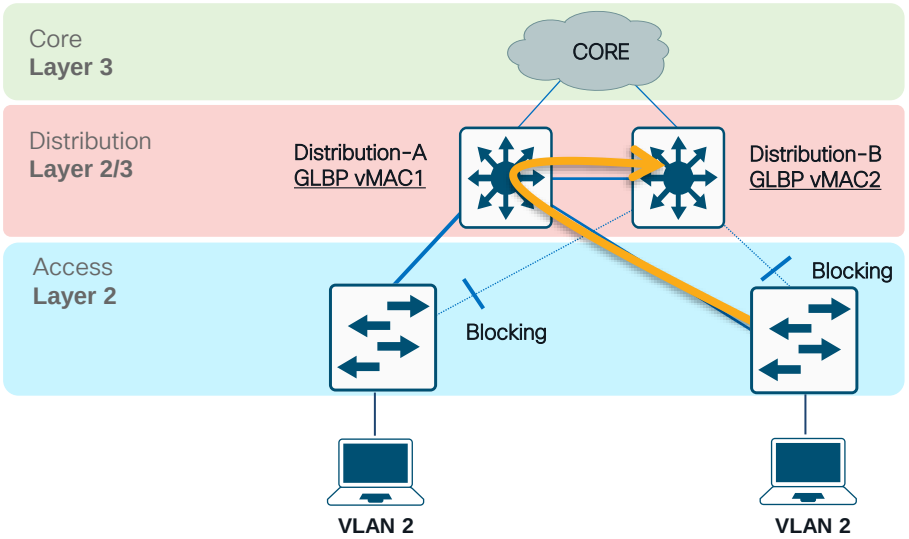

Campus Best Design Practices

Daisy Chaining Access Layer Switches

Avoid Potential Black Holes

- Return Path Traffic Has a 50/50 Chance of Being ‘Black Holed’

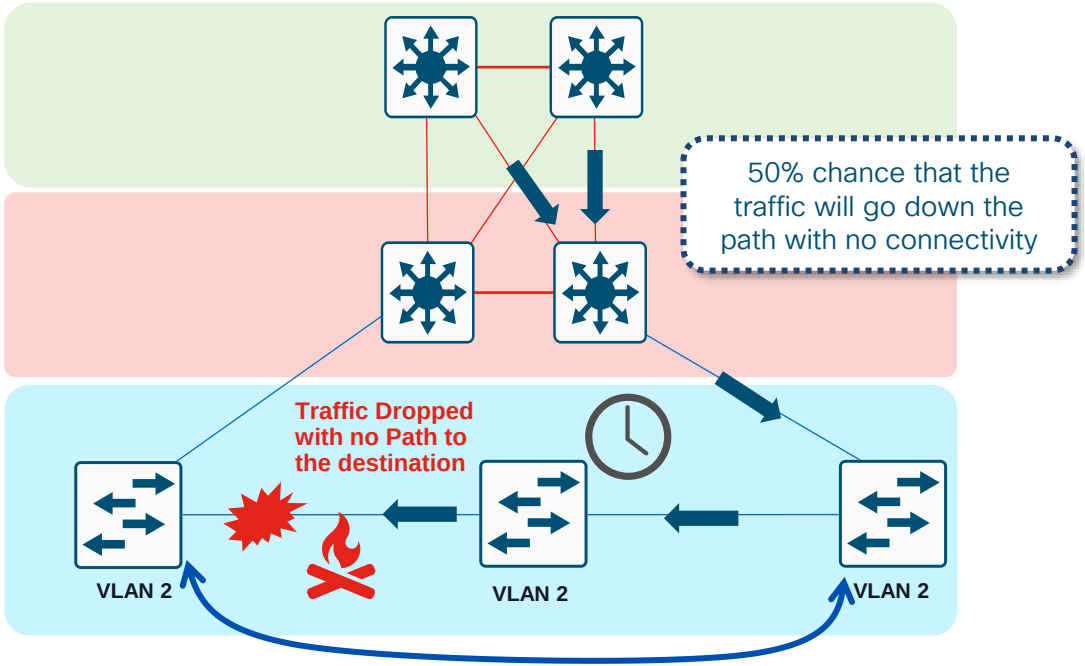

Daisy Chaining Access Layer Switches

Cisco StackWise technology

- Allows up to a maximum 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-standb**y. Remaining switches become stack **members**
- Multichassis EtherChannel **(MEC)** and cross-stack EtherChannel extend traditional EtherChannel by allowing Ethernet ports to be aggregated towards different physical chassis

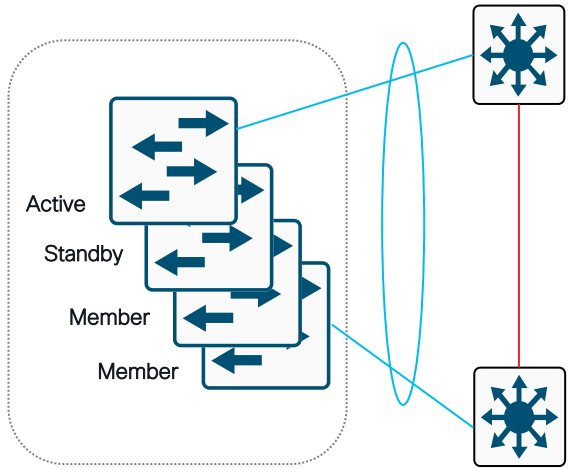

Single Unified Virtual Stack System

Cisco StackWise technology

- Catalyst 9200 Series StackWise-160/80
  - Catalyst 9200 Series switches enable 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 Series StackWise-480/360
  - Catalyst 9300 Series switches enable 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 Series StackWise-1T
  - Catalyst 9300 Series switches enable stacking of up to 8 switches and 448 ports

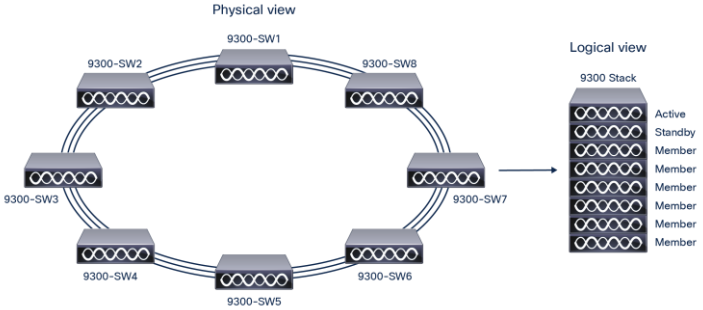

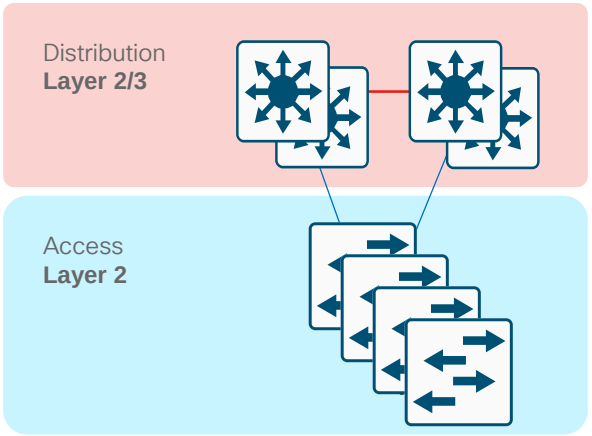

StackWise Virtual Technology

- StackWise Virtual technology combines **two** Catalyst 9000 Series switches into a single logical network entity from the network control plane and management perspectives.
- To neighboring devices a StackWise Virtual domain appears as a **single logical switch or router**
- All **control plane** functions are centrally managed by the **active switch**. From the **data-plane and traffic-forwarding** perspectives, **both switches actively** forward traffic.
- To facilitate this information exchange, a dedicated link – the **StackWise Virtual link (SVL)** – is used to transfer both data and control traffic between the peer switches. The SVL is formed as an EtherChannel interface of up to **eight physical port members.**

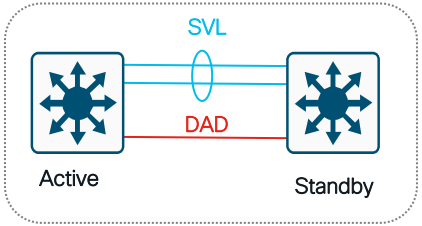

Single Unified Virtual Stack System

StackWise Virtual Technology

- Meant for Distribution and Core layer
- Formed using front panel ports
- Dual-homed connections

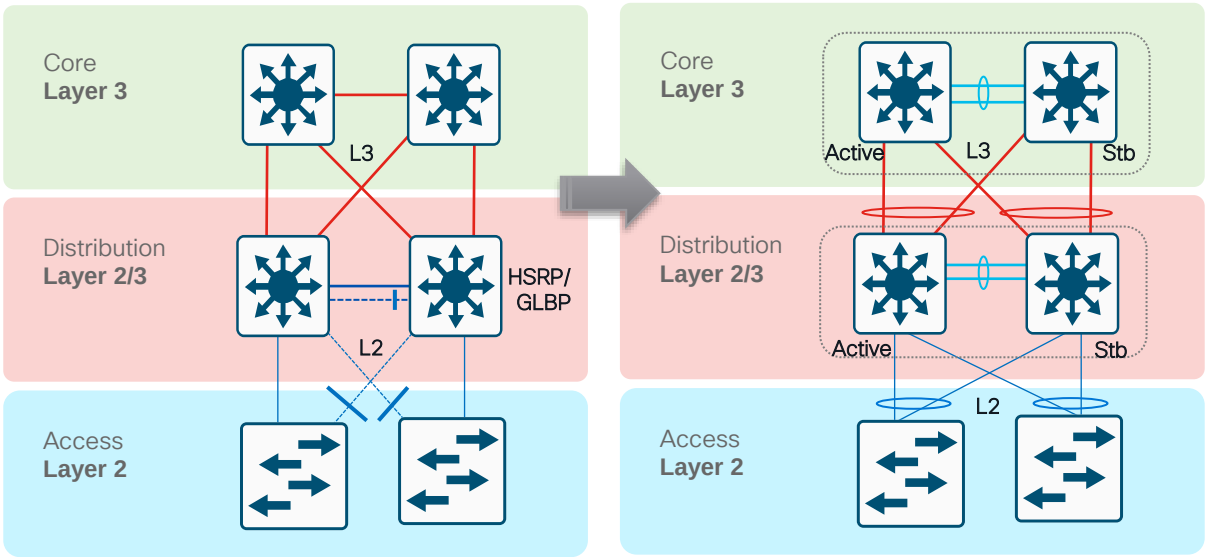

- Simplify Operations by Eliminating STP, FHRP and Multiple Touch-Points
- Double Bandwidth & Reduce Latency with Active-Active Multi-chassis EtherChannel (MEC)
- Minimizes Convergence with Sub-second Stateful and Graceful Recovery (SSO/NSF)

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

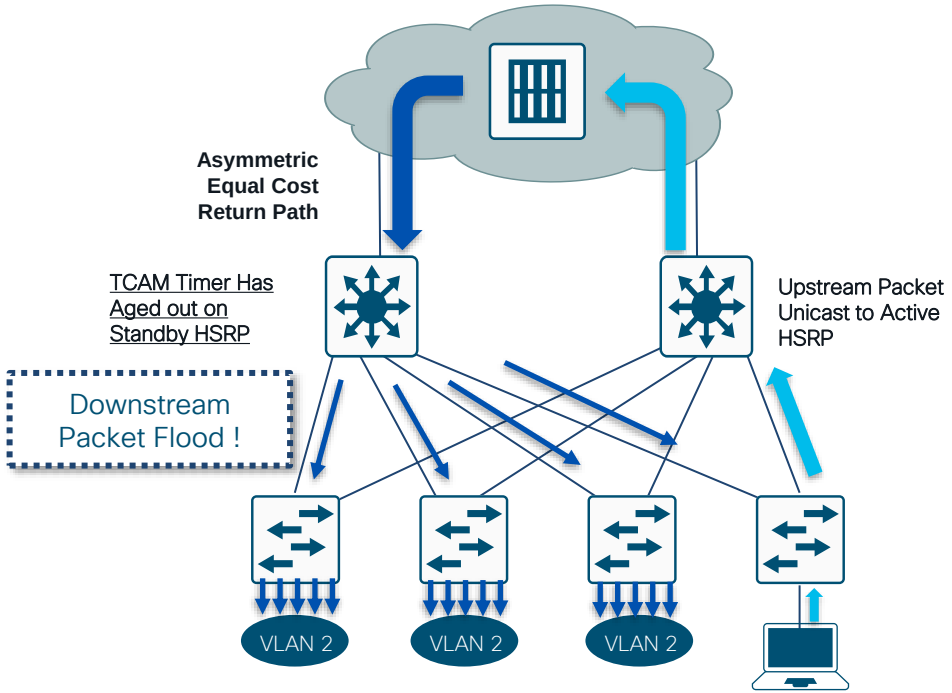

Best Practices Prevent Unicast Flooding

- 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
  - Bias routing metrics to remove equal cost routes

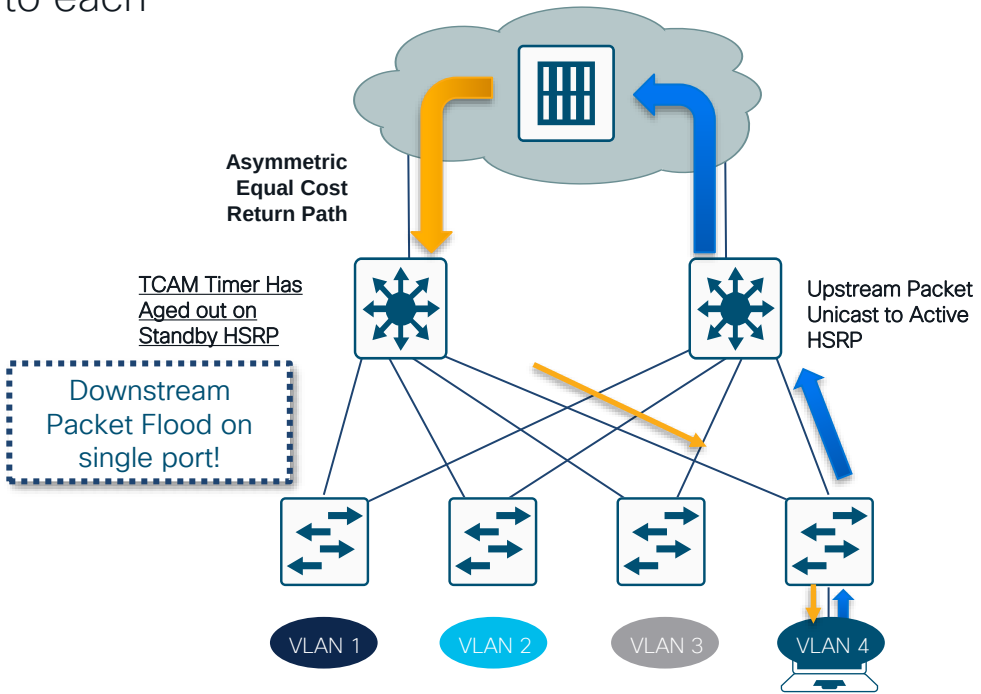

Routing in the Access

Pros:

- Improved convergence
- Simplified multicast configuration
- Dynamic traffic load balancing
- Single set of troubleshooting tools (for example, ping and traceroute)
- Ease migration towards SDA/EVPN

Cons:

- A different set of VLANs on different access switches
- Lower flexibility
- Overhead in additional IP subnetting planning

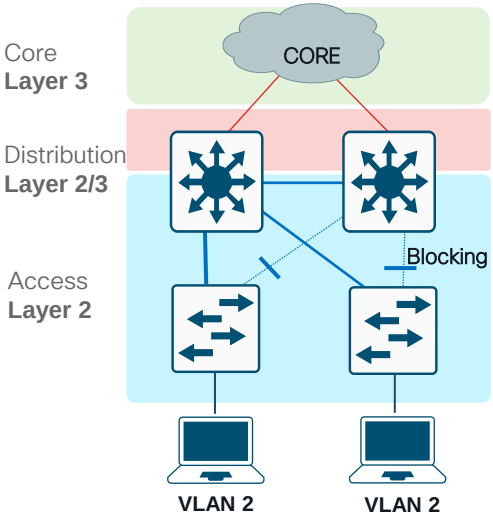

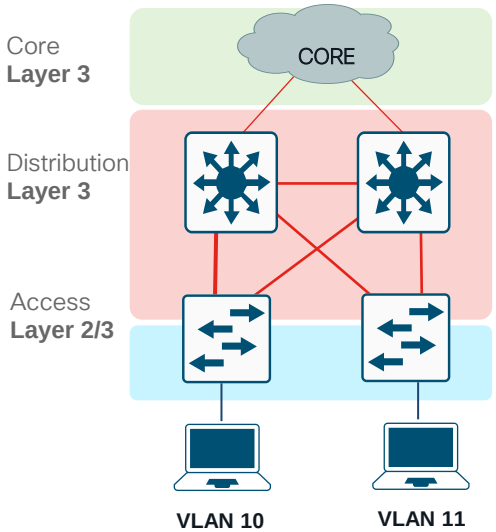

The ability to reduce convergence times to a **sub-200** msec range

Transmit Queue Congestion

The Case for Campus QoS

- The primary role of QoS in campus networks is to manage packet loss
- In campus networks, it takes only a few milliseconds of congestion to cause drops
- Rich media applications are extremely sensitive to packet drops

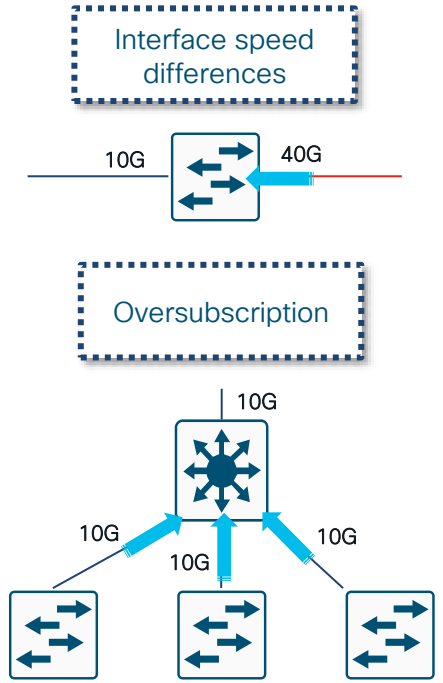

QoS End-to-End

- Prepare your strategy – what are the Critical/ Business relevant/Default applications?
- Understand QoS capabilities of used platforms
- Match the strategy against the platform capabilities
- Always build bidirectional and End-to-End policy

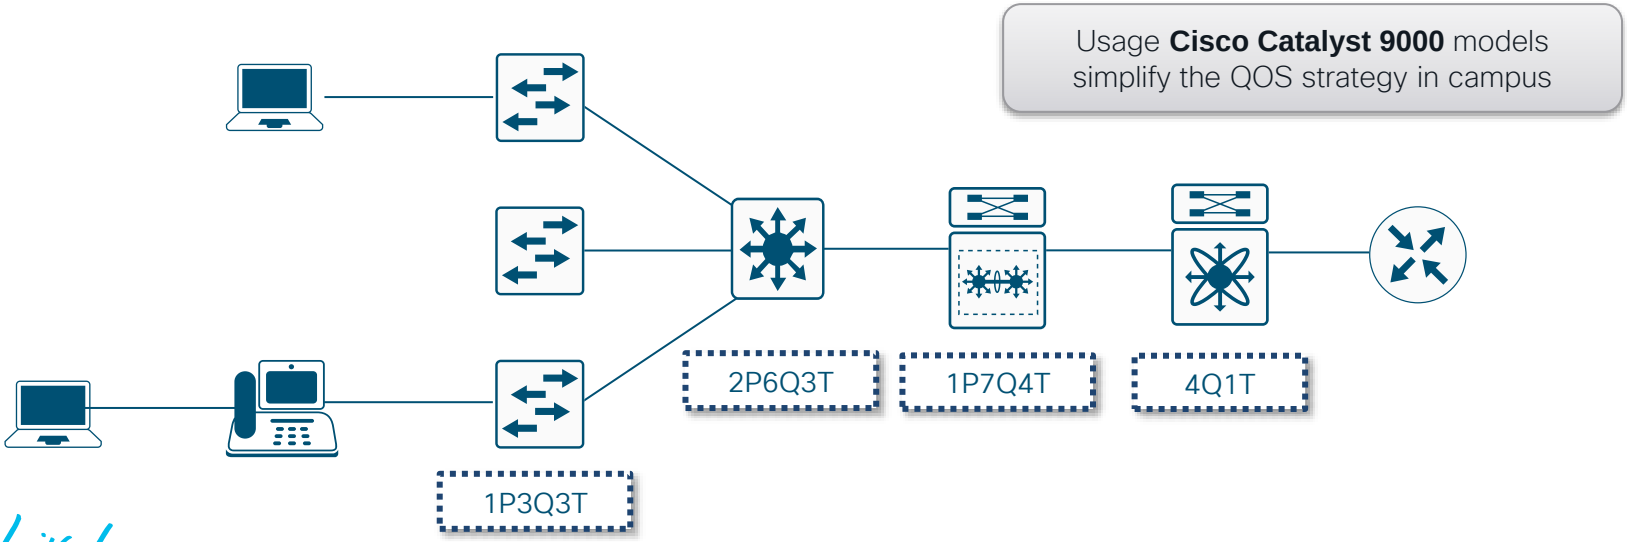

DNA-C QoS Automation with Application Policy

Network Operators express high-level business-intent to Application Policy

DNA Center

Southbound APIs translate business-intent to platform-specific configurations

Device Specific Config

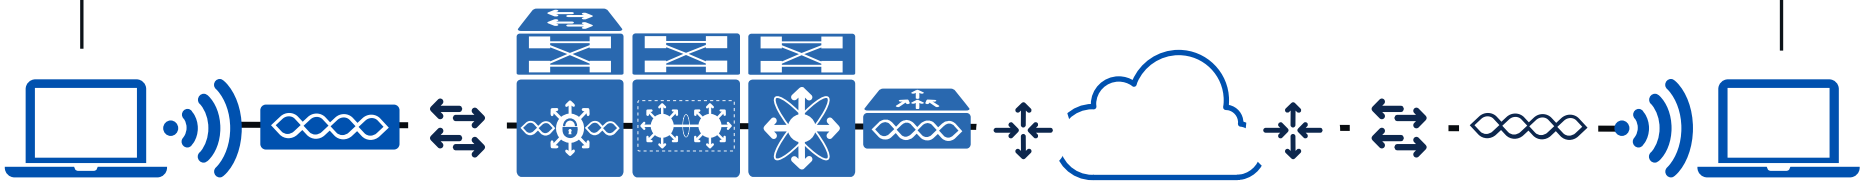

Catalyst 2960-X
1P3Q3T Egress Queuing Model

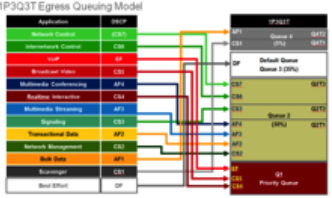

Catalyst 4500
1P7Q1T+Dynamic Buffer Limiting (DBL) Egress Queuing Model

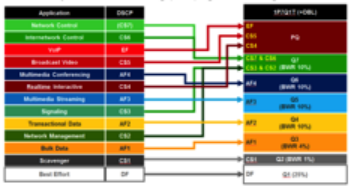

Cisco Catalyst 6500/6800
8Q4T Ingress & 1P7Q4T Egress Queuing Models (6908-10GE)

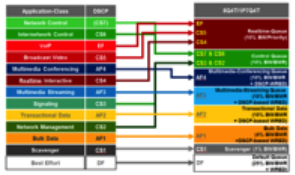

Cisco Nexus 7700 (F3)
4Q1T Ingress Queuing (DSCP-to-Queue) Model

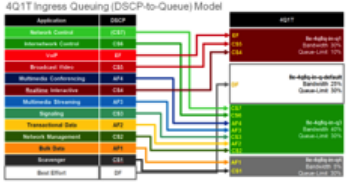

Catalyst 3650/3850
2P6Q3T with Weighted Tail Drop (WTD) Wired Port Egress Queuing Model

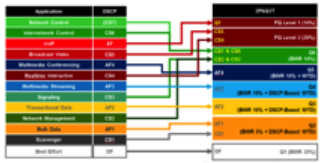

Conclusions

Without a Rock Solid Foundation – the Rest Doesn't Matter

Access

Distribution

Core

Distribution

Access

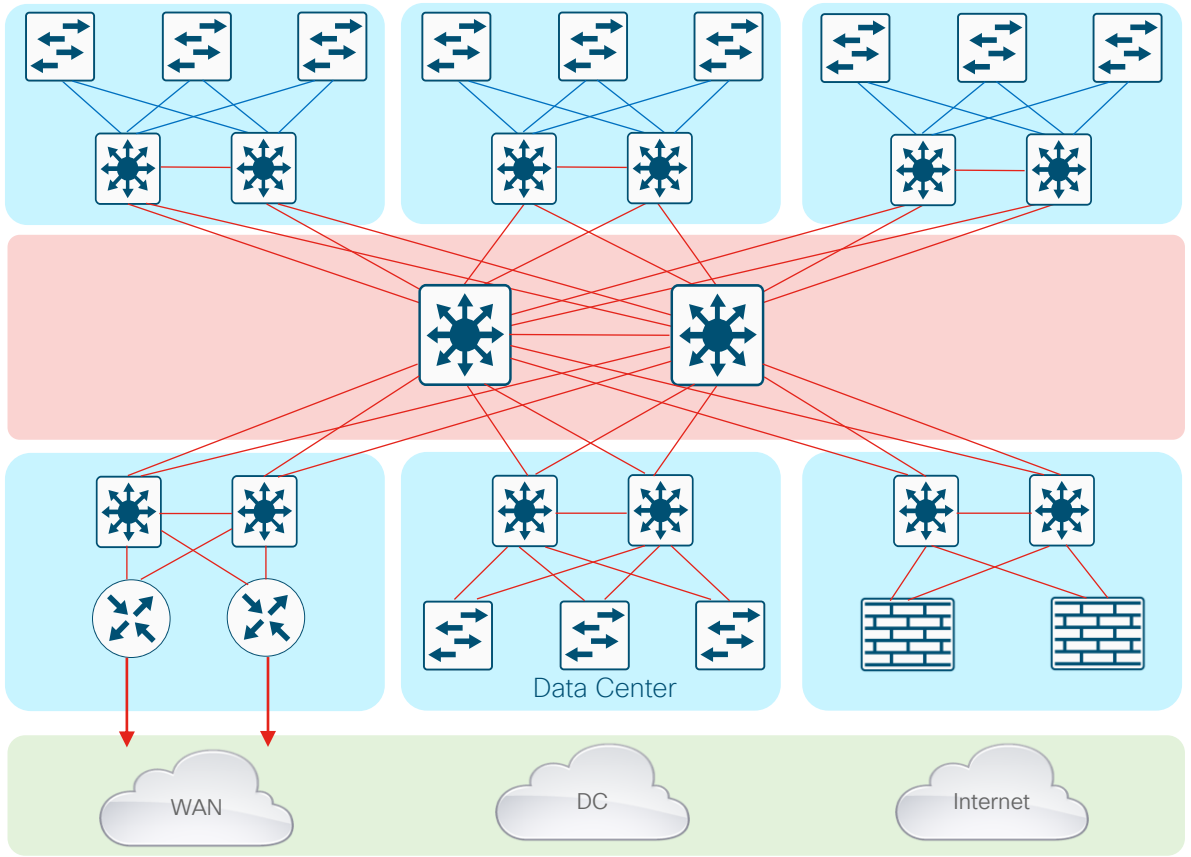

Summary

- Hierarchy – each layer has a 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

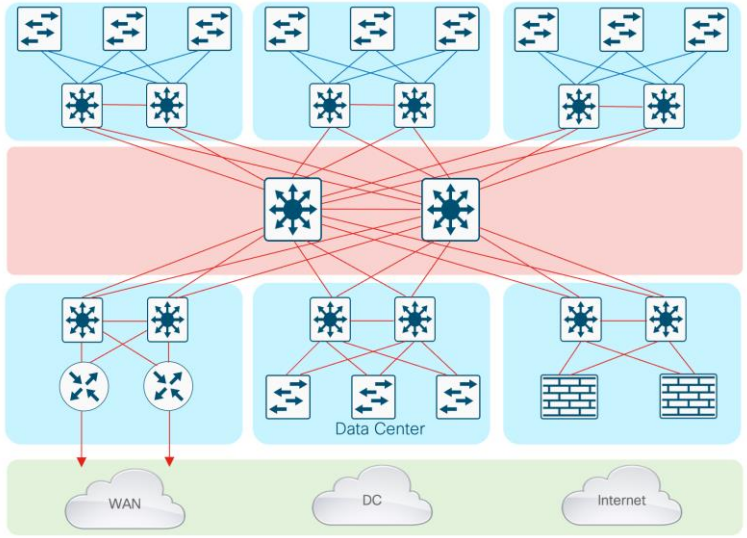

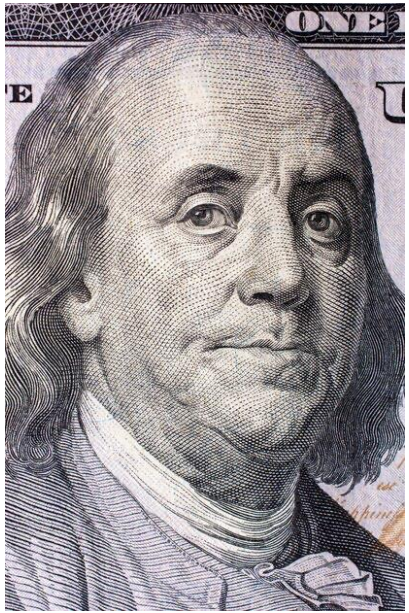

“.. If you fail to plan - you plan to fail”

Benjamin Franklin

Complete your Session Survey

- Please complete your session survey after each session. Your feedback is important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (open from Thursday) to receive your Cisco Live t-shirt.

- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Session Catalog and clicking the "Attendee Dashboard" at https://www.ciscolive.com/emea/learn/sessions/session-catalog.html

Continue Your Education

Visit the Cisco Showcase for related demos.

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

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

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


cisco
The bridge to possible
Thank you
cisco Live!


CISCO Live!
ALL IN