
cisco
You make possible

Designing for High Availability Switching and Wireless in Your Campus LAN

Dana Daum
Maren Kostede

TECCRS-2001

cisco Live!

Barcelona | January 27–31, 2020

Cisco Webex Teams

Questions?

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

How

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

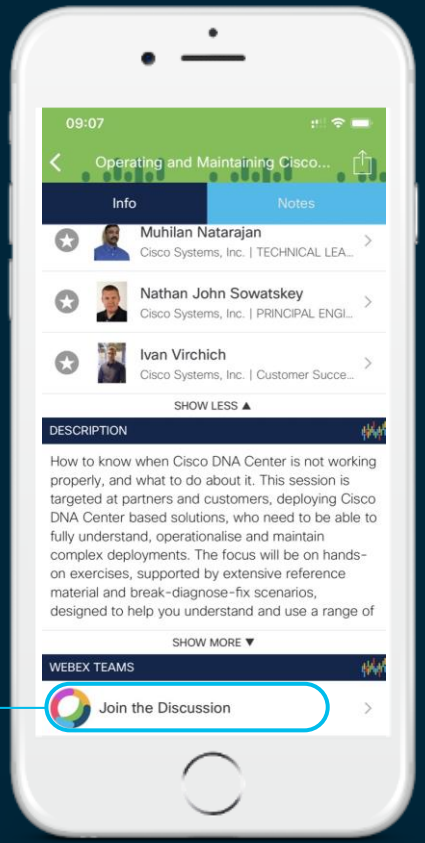

Session schedule

Dana

08:45	What is high availability?
	Campus network foundations and structured design
	Campus wired LAN design and high availability
10:45	Break
11:00	Campus wired LAN design and high availability (cont.)
11:15	Campus wireless LAN
	Summary, conclusions, Q & A
13:00	Lunch (available until 14:30)

Maren

Agenda

- What is high availability?
- Campus network foundations and structured design
- Campus wired LAN design and high availability
- Campus wireless LAN design and high availability
- Summary and conclusions

What is high availability?

What is availability?

Levels of availability

Referencing de facto industry terminology

Continuous Availability

- Designed to operate 24 hours, 7 days/week
- Goal to handle ALL unplanned faults and planned maintenance

Continuous Operations

- Designed to operate 24 hours, 7 days/week
- Supports operations during planned maintenance and handles unplanned faults

High Availability

- Designed to a specified service level
- Handles unplanned faults, typically by eliminating single points of failure

“The Nines” – Network availability and downtime

Network availability: amount of uptime of a network system over a specific time interval, measured as a percentage.

Availability	Downtime per year
90%	36 ½ days
99%	3 days, 16 hours
99.9%	8 hours, 46 minutes
99.99%	52 minutes
99.999%	5 minutes

How can we measure the predicted availability?

It's function of:

Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR)

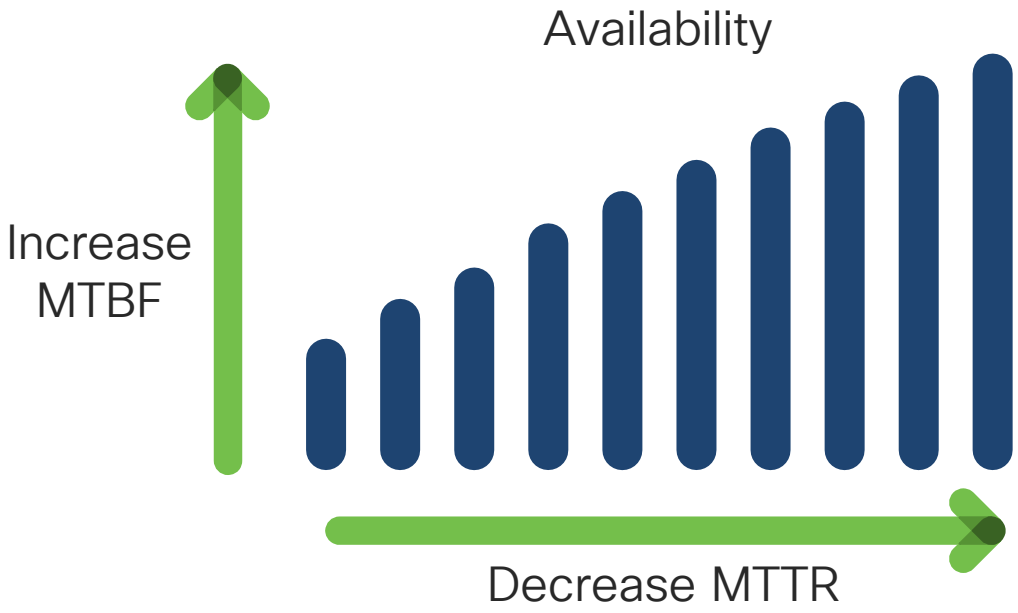

A basic predicted availability equation

Predicted Availability Equation

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

MTBF: Mean Time Between Failures

MTTR: Mean Time To Repair

Example predicted availability calculations

Component with MTBF=87,600 hours

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

24 hour depot replacement

2 hour 24 minutes predicted annual downtime

$$\text{Availability} = \frac{87{,}600}{87{,}600 + 24} = .9997\ (99.9\%)$$

4 hour depot replacement

24 minutes predicted annual downtime

$$\text{Availability} = \frac{87{,}600}{87{,}600 + 4} = .99995\ (99.99\%)$$

Warm spare (10 minute restore)

1 minute predicted annual downtime

$$\text{Availability} = \frac{87{,}600}{87{,}600 + .16666} = .999998\ (99.999\%)$$

The redundancy effect for a system

- Single components functioning in series
- System predicted availability: 99.98% (~10 min./year predicted downtime)

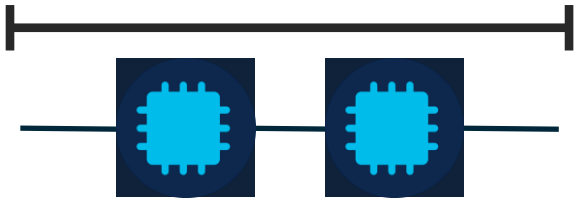

Supervisor 99.999% Availability (~5 min/yr. downtime)

Linecard 99.999% Availability (~5 min/yr. downtime)

- Redundant components functioning in parallel
- System predicted availability: 99.999999% (~½ second/year predicted downtime)

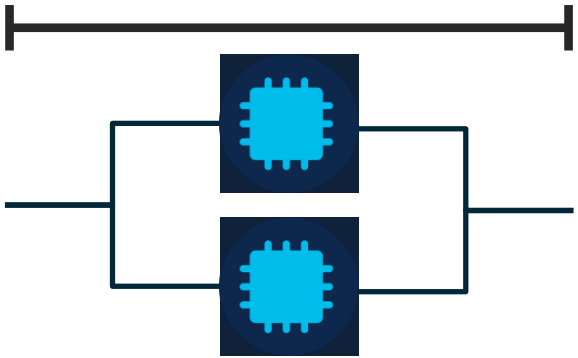

Linecards
99.999% Availability (~5 min/yr. downtime)

Example of predicted availability rating (Catalyst 6800XL non-redundant)

Reference

Catalyst 6800XL

Part	MTBF (hours)	MTTR	Combined MTBF Hrs.	Combined Availability	Predicted Annual Downtime
Chassis C6807-XL	638,440	4 hrs.	638,440	99.99937348%	--
C6807-XL-FAN	3,077,880	4 hrs.	3,077,880	99.99987004%	--
SFP-10GSR	2,294,776	4 hrs.	2,294,776	99.99982569%	--
Supervisor VS-S2T-10G	231,910	4 hrs.	231,910	99.99827522%	--
WS-X6904-40G-2T	256,490	4 hrs.	256,490	99.99844051%	--
C6800-XL-3KW-AC	3,000,000	4 hrs.	3,000,000	99.99986667%	--
System MTBF			91,987	**99.99565168%**	22.87 min.

Components combined in **series** calculation

Chassis X Fan Tray X Power Supply X Line Card X Supervisor Module X SFP Uplink = System MTBF

Example of predicted availability rating (Catalyst 6800XL with redundancy)

Reference

Catalyst 6800XL with Redundancy

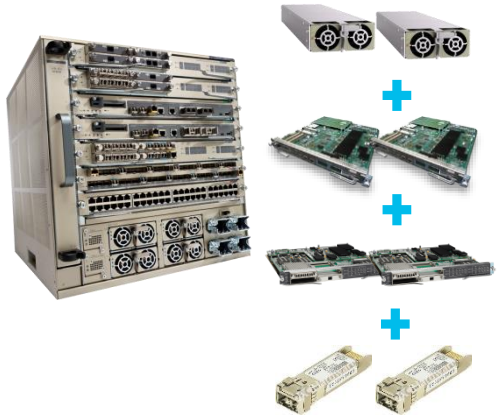

Part	MTBF Hrs.	MTTR Hrs.	Switchover time (seconds)	Combined MTBF Hrs.	Combined Availability	Predicted Annual Downtime
Chassis C6807-XL	638,444	4 Hrs.	--	638,440	99.99937348%	--
C6807-XL-FAN=	3,077,880	4 Hrs.	--	3,077,880	99.99987004%	--
SFP-10GSR	451,610	4Hrs.	.500	2,633,000,739,868	100.00000000%	--
Supervisor VS-S2T-10G	2,294,776	4 Hrs.	.500	26,891,355,961	99.99999997%	--
WS-X6904-40G-2T	402,386	4 Hrs.	.500	32,893,816,541	99.99999998%	--
C6800-XL-3KW-AC	3,000,000	4 Hrs.	0	4,500,003,000,001	100.00000000%	--
System MTBF				**528,687**	**99.99924347%**	**3.98min.**

Redundant components combined in **parallel** calculation

Chassis X Combined Power Supply X Combined Line Card X Combined Supervisor Module X Combined SFP Uplink = System MTBF

Predicted availability ratings and system choices

- Predicted availability ratings are not guarantees of component or network availability
- Ratings are based on industry standard methodologies and statistical analysis
- Useful in making design decisions through comparison of different options
- Design choices are driven by business requirements – availability is one aspect
- Platform choices often based on mix of capabilities, capacities, and compliance
  - Backplane throughput and performance; interface types and port densities
  - Scalability for future growth / investment protection
  - Software upgrade procedures, software feature support
  - Simplicity and ease of use
  - Industry certifications

Systems approach to campus network availability

- System-level resiliency
- Network-level redundancy
- Enhanced management
- Human ear notices the difference in voice within 150–200 msec (10 consecutive G.711 packet loss)
- Video loss is even more noticeable
- 200 msec typical end-to-end campus convergence target

Ultimate goal – 100% availability

Examples:

- Next-generation applications, video conferencing, unified messaging, e-business, wireless
- Mission-critical applications, databases, order entry, CRM, ERP
- Desktop applications, e-mail, file, print

An organization's applications drive requirements for high availability networking

What if video delivery is key to your organization?

1920 lines of Vertical Resolution (Widescreen Aspect Ratio is 16:9)

1080 lines of Horizontal Resolution

1080p60

1080 x 1920 lines =

2,073,600 pixels per frame

x 24 bits of color per pixel

x 60 frames per second

= 2,985,984,000 bps

or 3 Gbps Uncompressed!

Cisco (H264/H.265) codecs transmit 3-5 Mbps per 1080p60 video stream (99.8%+ *compression, ~1000:1).* Packet loss is proportionally magnified by compression ratios. Users can notice a single packet lost in 10,000.

HD video is *one hundred times more sensitive to packet loss than VoIP!*

Measure and analyze event total service downtime

- Measure all previous events
  - Note each in trouble tickets
  - Analyze trends
- Automation
  - Trouble ticketing
  - Technology/database
- Redundant network design and resiliency features
  - Required for very high availability

Fault starts | Notification time | Dispatch time (parts, SW, people) | Repair time

Failure detected | Diagnostic time | Arrival time | GO

Examples: Measuring network availability

OSI model layers	Visibility / measurements
Application layer	Custom application scripts, HTML, TCL, Python, many others
Presentation layer	
Session layer	
Transport layer	ICMP ping, IP traceroute, Bidirectional Forwarding Detection, IP SLA
Network layer	
Data link layer	UDLD, BPDU, CDP, LLDP
Physical layer	Cable testers, power meters, OTDR

Main operational challenges

Source: 2016 Cisco study

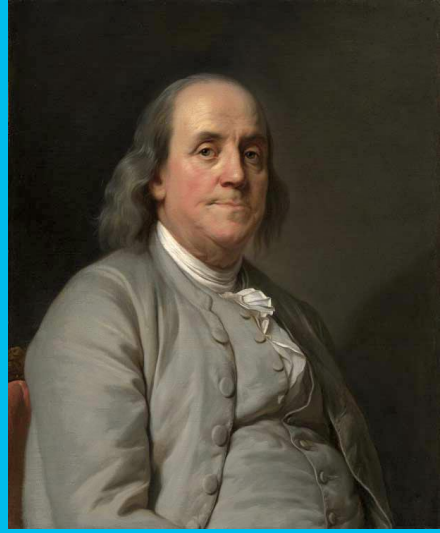

“By failing to prepare,
you are preparing to fail.”

– Ben Franklin

Planned versus unplanned outages

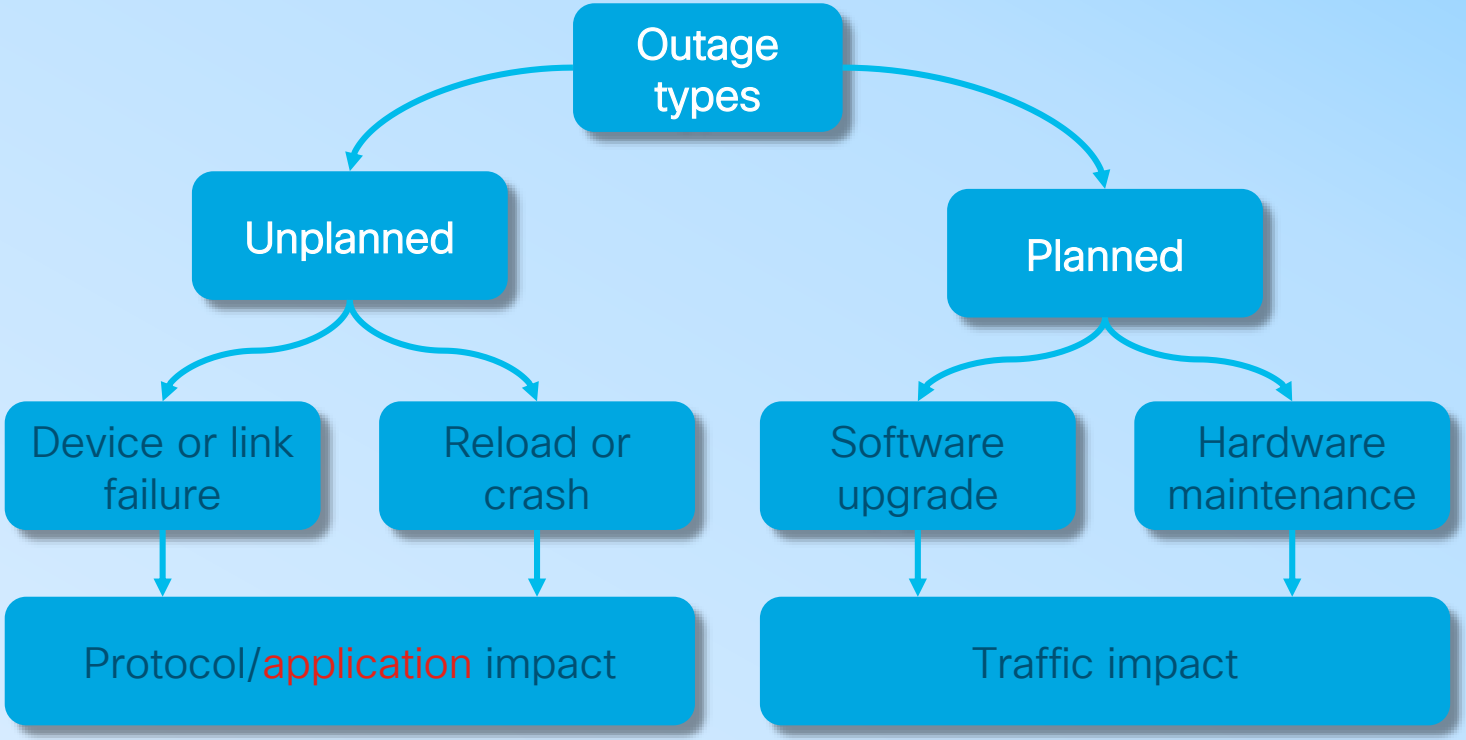

Where can outages occur ?

Unplanned outages

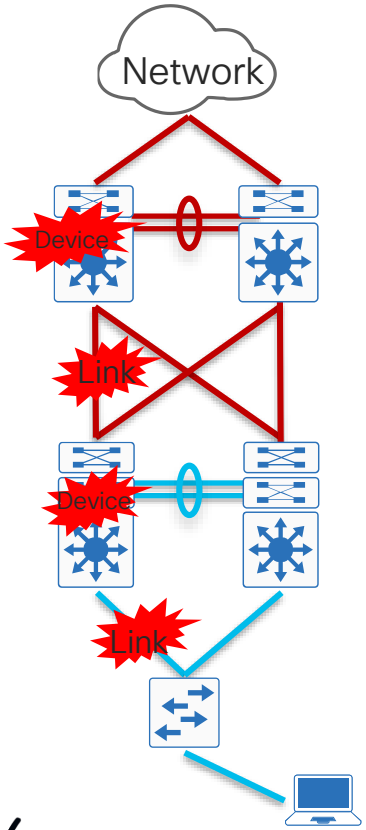

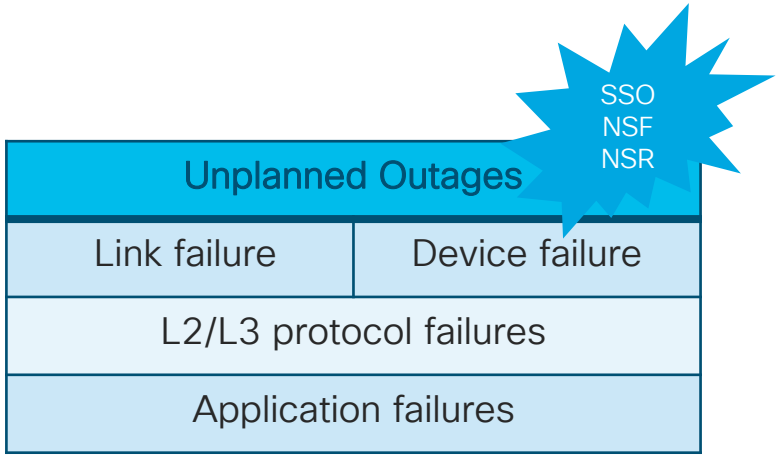

Some platforms also support Process Restart

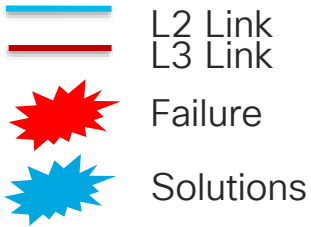

Where can outages occur ?

Planned outages

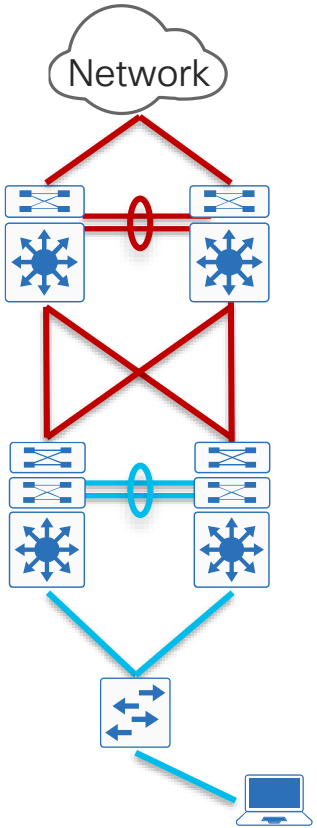

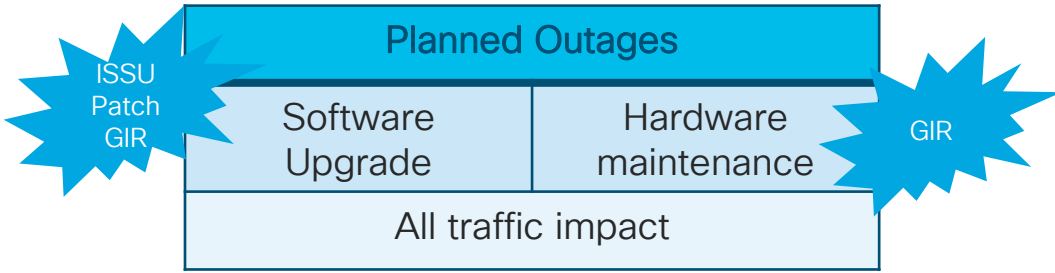

Patching can be also used, depending on the upgrade

L2 Link
L3 Link

Solutions

Reducing MTTR: Many tools in the toolbox

Preview of deeper dives

- Device resiliency:
  - Redundant components
  - Redundant chassis/stacking
  - Virtualized stacking
  - Controller HA SSO
- SSO / NSF
- ISSU / SMU / FSU / XFSU / Staggered upgrades
- NSR
- GIR

Campus network foundations and structured design

Agenda

- What is high availability?
- **Campus network foundations and structured design**
  - **Wired campus platform hardware and software features for HA**
  - Overview of campus structured design
- Campus wired LAN design and high availability
- Campus wireless LAN design and high availability
- Summary and conclusions

Cisco Catalyst 9000 Series–switching transitions

Greater flexibility from small remote site to mission critical campus core.

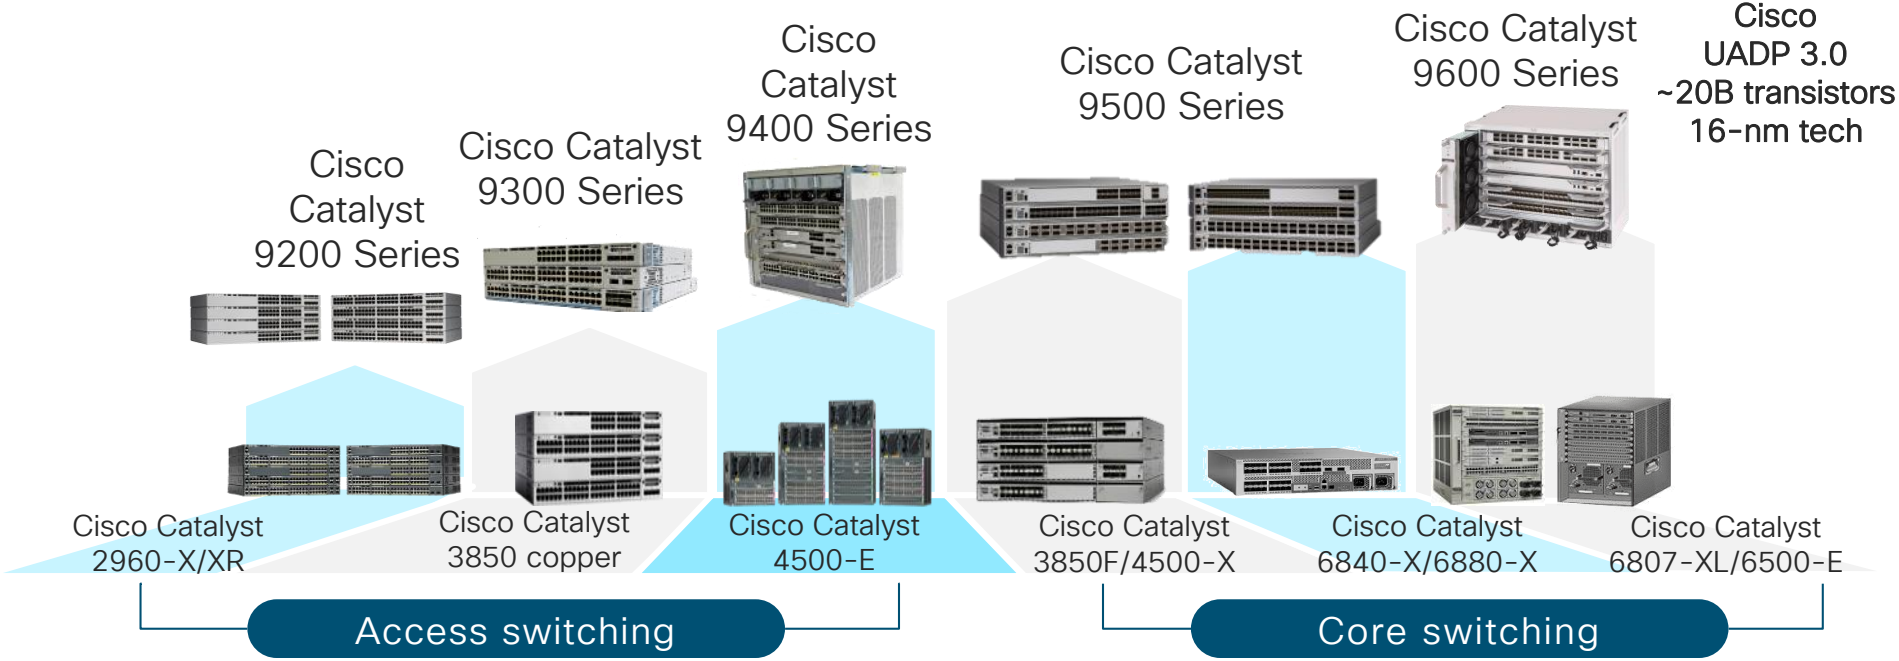

“Classic” Catalyst 2960-X stack resiliency

Reference

- Stack Master provides central control over multiple 2960 Series switches configured in a stack
- To increase resiliency in a 2960 stack of three or more switches:

Configure the Stack Master on a switch that does not have uplinks configured

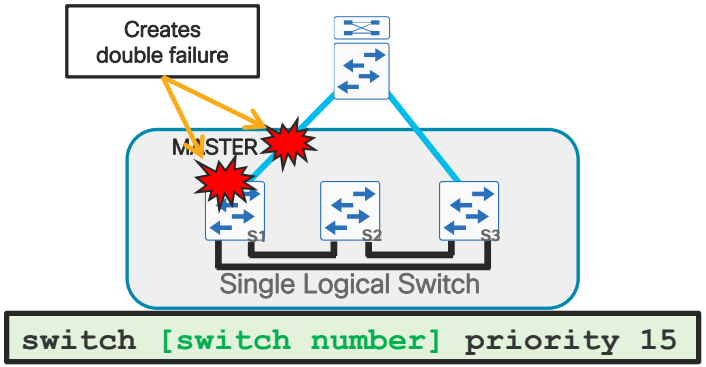


```
switch [switch number] priority 15
```

Ensure that the original Stack Master MAC address remains the stack MAC address after a failure to prevent protocol restart

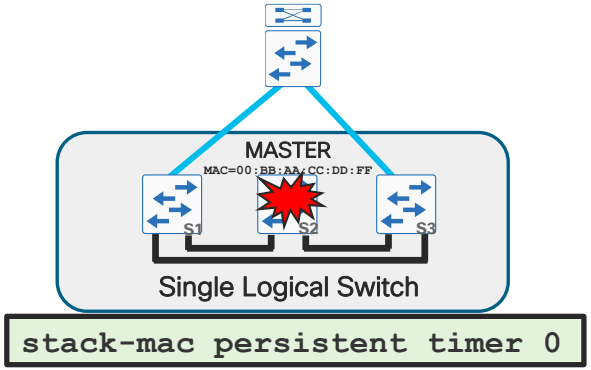


```
stack-mac persistent timer 0
```

Stateful Switchover

Catalyst 9000 Series and 3x50 stacks, also 4500, 6500, 6800 Modular

- Modular chassis with dual supervisors offers Stateful Switchover (SSO) configuration
- Redundant chassis with StackWise, StackWise Virtual, or Virtual Switching System (VSS) also provides SSO
- Traffic loss minimized for failure of active control plane

Stateful Switchover

Modular Chassis

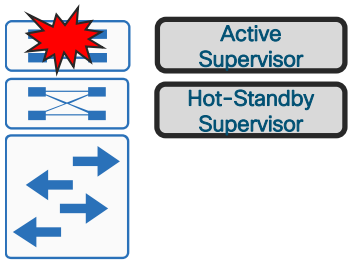

Stateful Switchover

C9300/C3x50 Stack

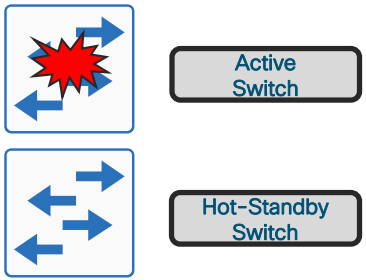

Catalyst 9300 Series

Cisco StackWise-480

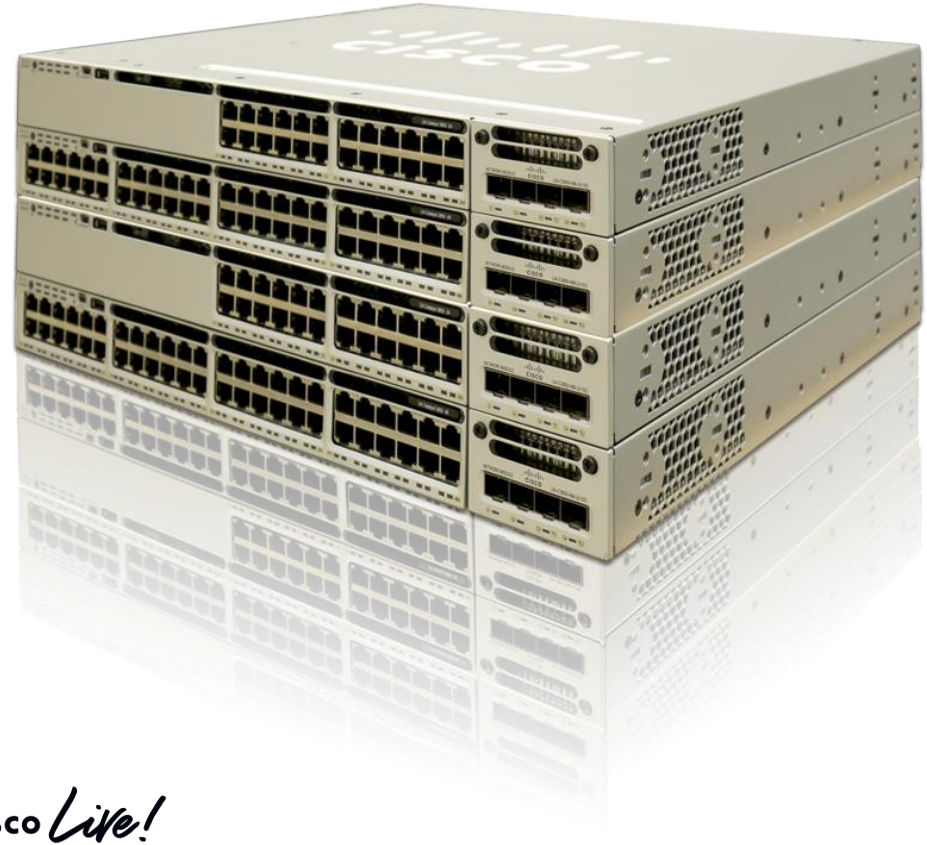

Cisco StackWise-480: Stack Ring

Example: 4x Catalyst 9300 Series switches

- 6 rings in total
- 3 rings clockwise
- 3 rings counter/anti-clockwise
- Each ring is 40Gbs
- Total Stack BW = 240Gbs
- With Spatial Reuse = 480Gbs

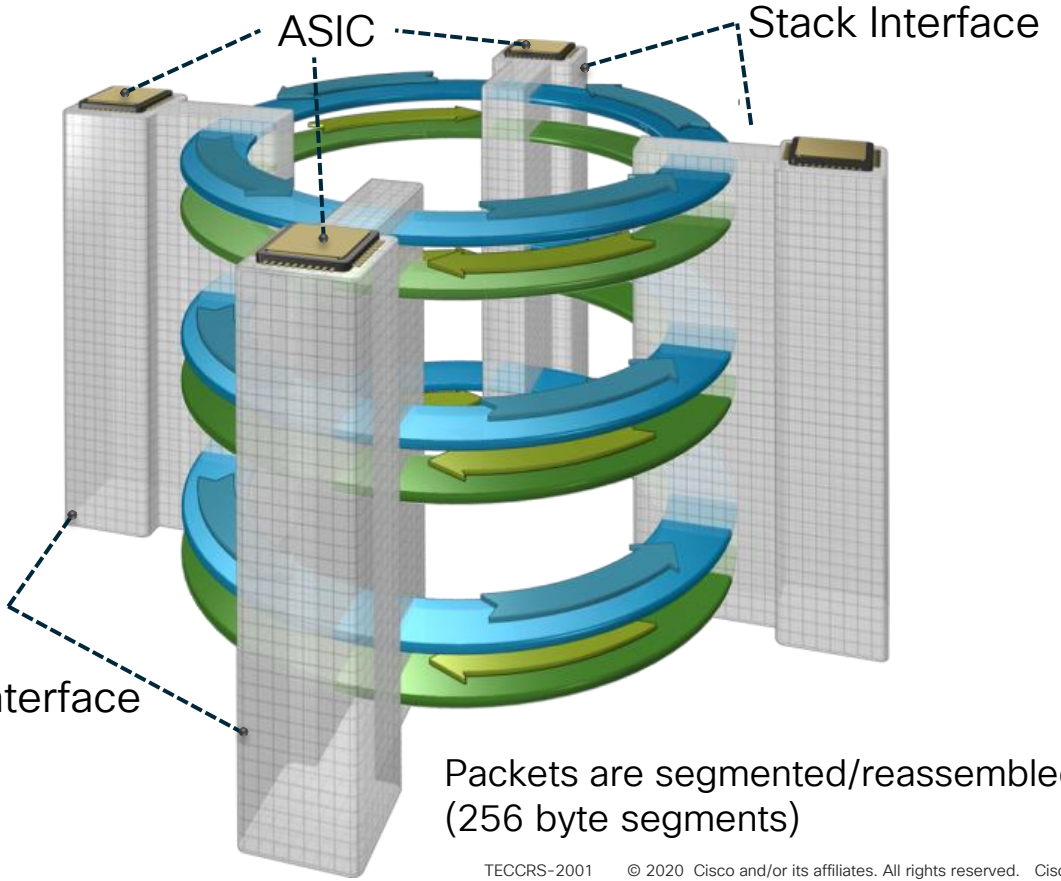

Packets are segmented/reassembled in HW (256 byte segments)

SSO and `show switch` command output

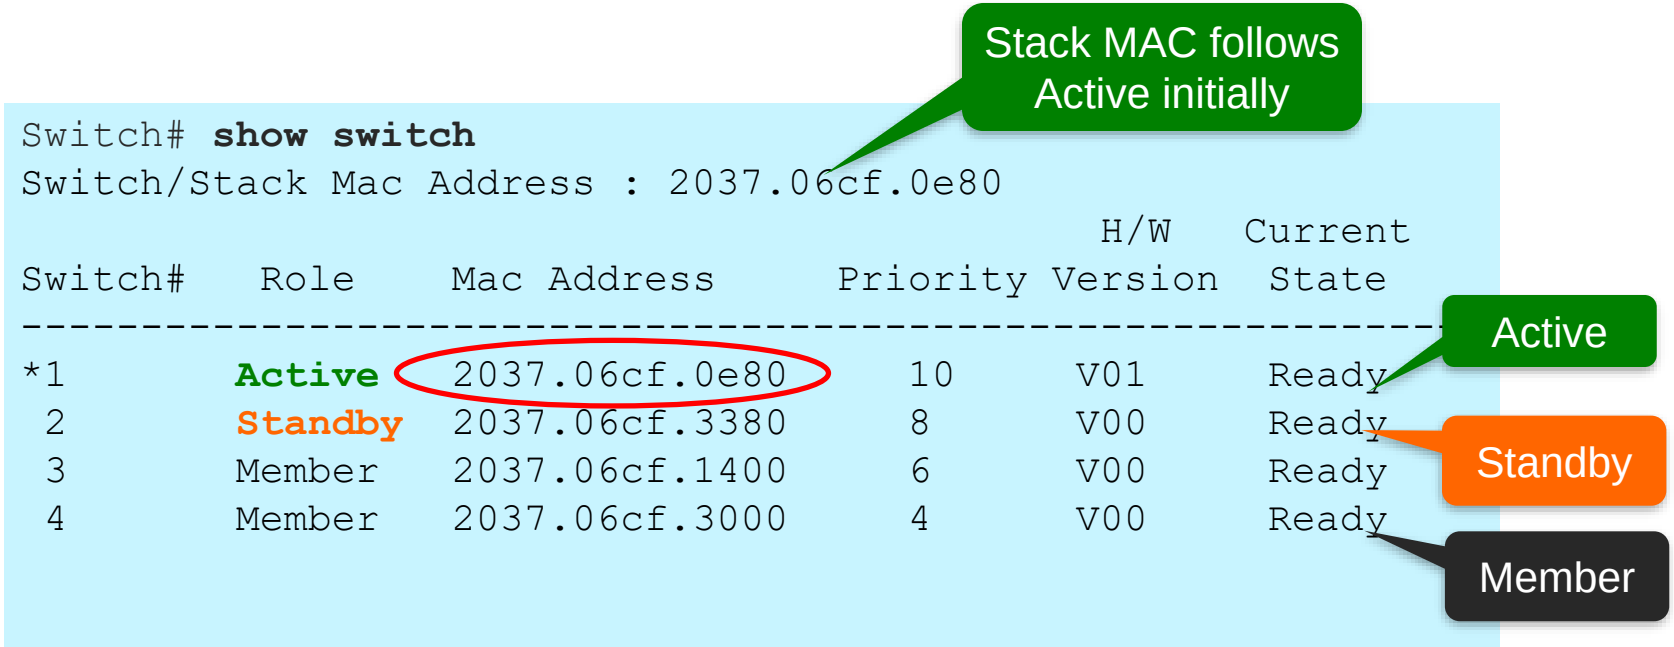

* Indicates which member is providing the “stack identity” (aka “stack MAC”)

SSO – Catalyst 9000 Series modular chassis

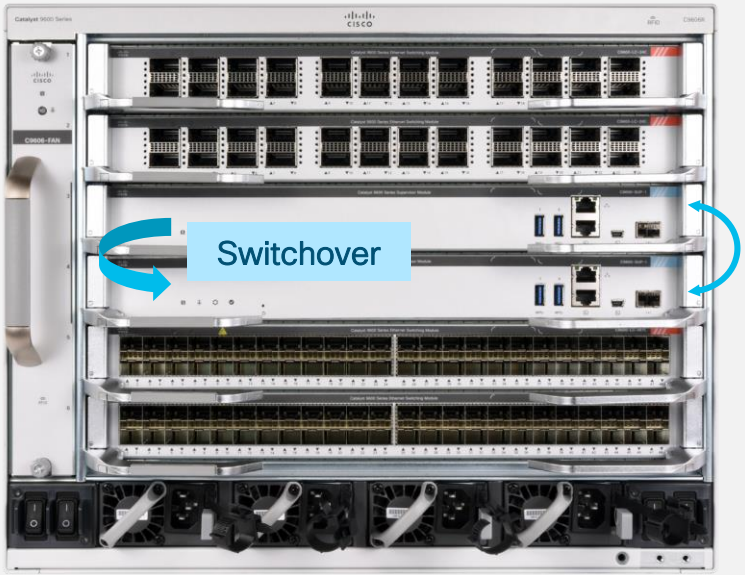

SSO is the default redundancy mode with two supervisors in the system

- The active supervisor is responsible for all control plane processing
- The active supervisor is responsible for hardware programming on both the active and standby supervisors

Supervisors and line cards: data path

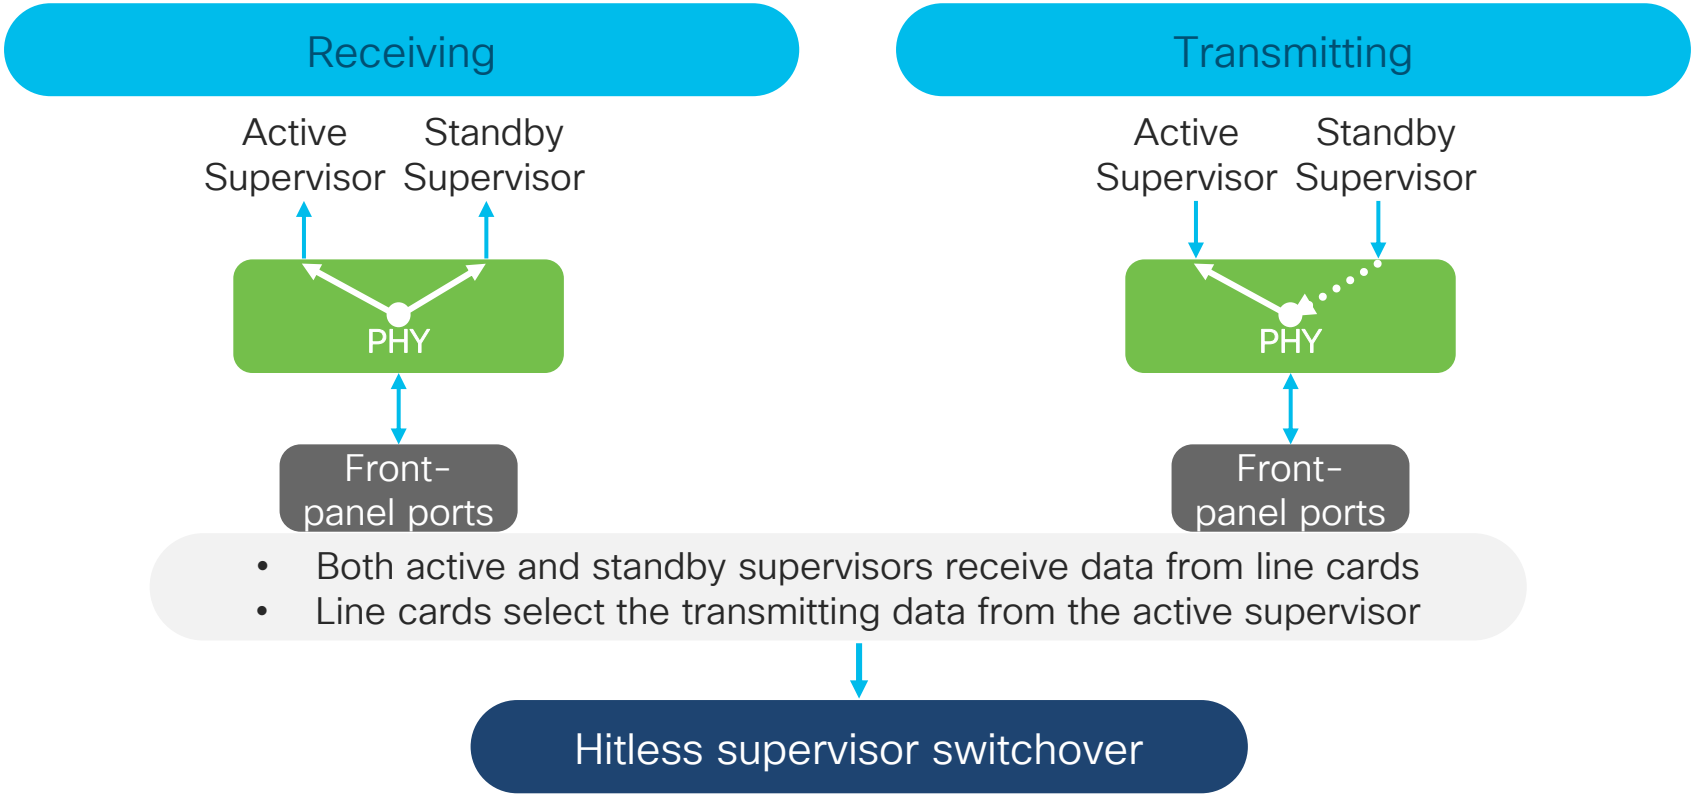

Additional SSO-capable options

Catalyst 9000 Series and Catalyst 3850 – Cisco StackWise Virtual

- Cisco StackWise Virtual: an evolution of Catalyst Virtual Switching System technology
- Fixed switch hardware architecture with distributed forwarding architecture
- StackWise Virtual Link (SVL) between two nodes (10Gb or 40Gb)
- Both StackWise Virtual members must have consistent Cisco IOS-XE and license
- Check software release notes for versions, supported platforms, and additional uplink/line card hardware

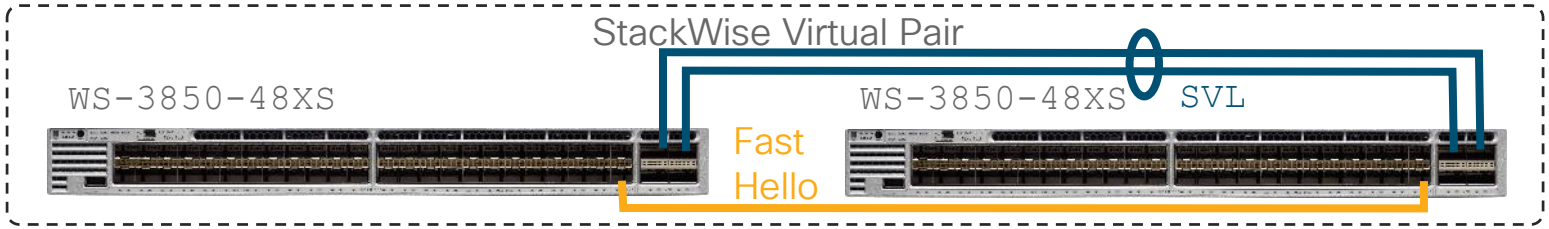

Cisco StackWise Virtual – Catalyst 9600

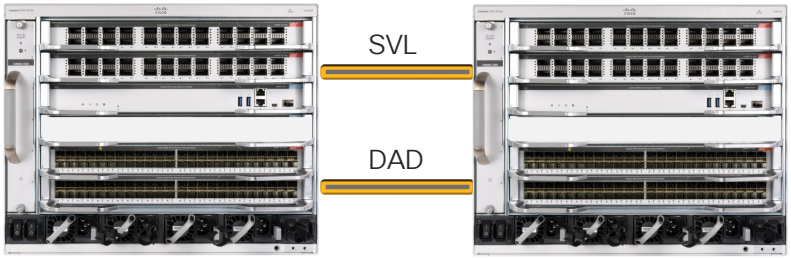

Cisco StackWise Virtual for Catalyst 9600 is supported with IOS-XE 16.12.1 or later.
Check release notes for hardware / software constraints.

- SVL: StackWise Virtual Link
  - same speed ports (10G or higher)
  - Up to 8 ports
- DAD: Dual Active Detection:
  - Fast Hello
    - Directly connected
    - Up to 4 links
  - Enhanced PAgP
    - EtherChannel with PAgP
    - Up to 4 port-channels
- In SVL mode, 2nd Supervisor is not supported in the chassis and will be powered off if inserted.

- Typically a distribution layer technology, allowing “stacking” of 2 switches
- Supports flexible distances with support of all supported cables and optics
- SVL and DAD are supported on any port with 10G or high speed, including QSA.

Quad-Supervisor RPR StackWise Virtual

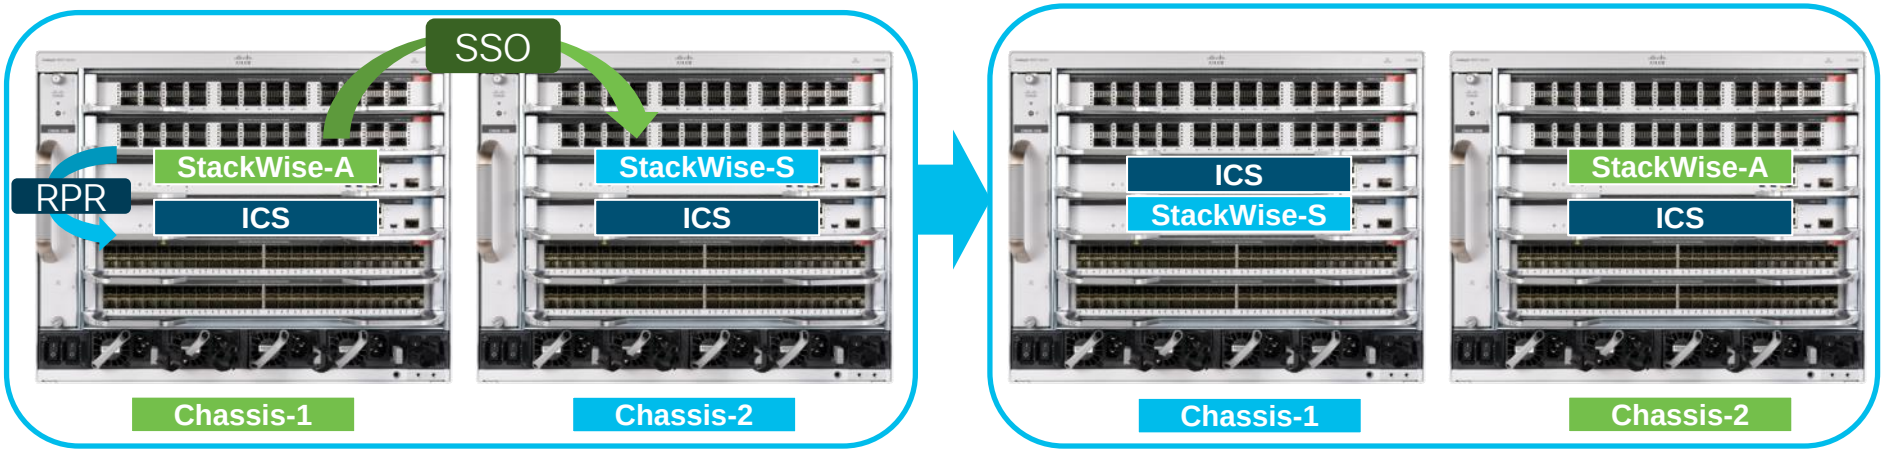

- Initially on Catalyst 9600 (Limited Availability)
- Active supervisor in chassis-2: becomes StackWise ACTIVE
- Warm standby supervisor in chassis-1: continues the boot process to become StackWise STANDBY-HOT while the line cards in chassis-1 get reset

RPR:	Route Processor Redundancy
SSO:	Stateful Switchover
StackWise-A:	StackWise Virtual ACTIVE
StackWise-S:	StackWise Virtual STANDBY-HOT
ICS:	In-chassis Warm Standby

9600 IOS-XE 17.1
Limited Availability

Non-Stop Forwarding (NSF) compliments SSO

- Non-Stop Forwarding:
  Router continues forwarding data to known routes, during routing protocol information restoration (graceful restart)
- NSF Aware (NSF Helper*) router:
  Runs NSF-compatible software, capable to assist neighbor router performing NSF restart
- NSF Capable router:
  Router configured for NSF restart, can rebuild routing information from neighbor NSF-aware router

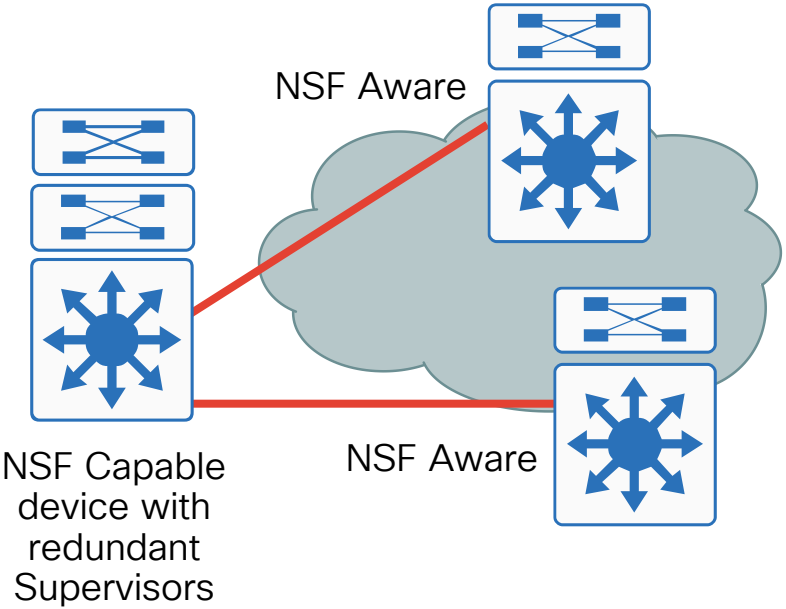

* NSF Helper - Term used in IETF terminology

NSF Interoperability

Interoperability between different Cisco devices

Reference

- The Graceful Restart extensions used in NX-OS are based on the IETF RFCs except for EIGRP, which is Cisco proprietary and can interoperate with Cisco NSF.
- This implies that routing protocols that support the GR extensions in NX-OS are compatible with versions of IOS-XE only when using the RFC based extensions

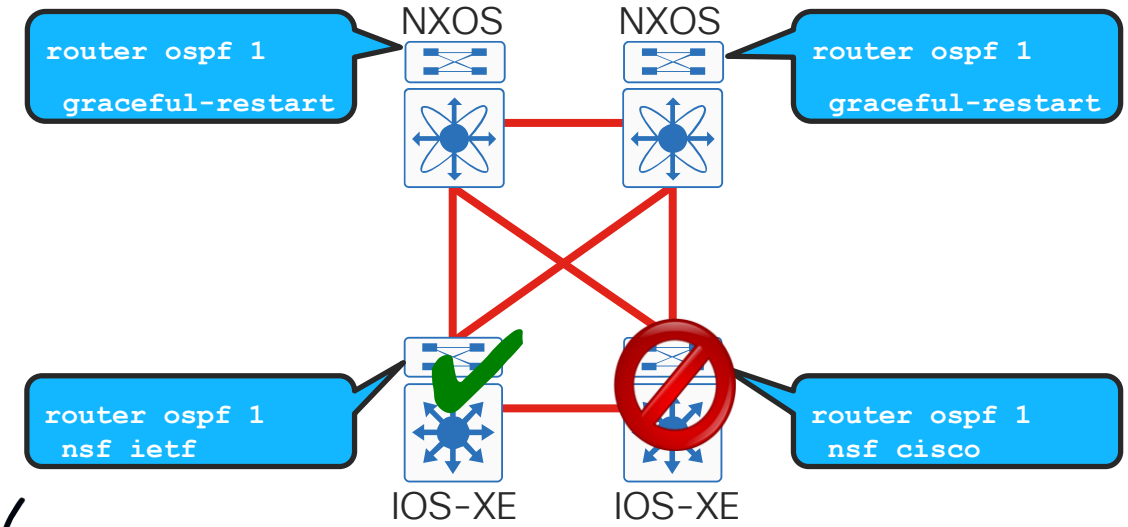

Reducing reload time on Catalyst 9300

Fast Software Upgrade (FSU) and Extended FSU (xFSU)

- FSU:
  - Mechanism to upgrade and downgrade the software image
  - Segregates updates of control plane and data plane
- xFSU: IOS-XE 17.1.1
  - Updates the **control plane** by leveraging the NSF/GR (SSO) architecture
  - Uses a flush and relearn mechanism to reduce **data plane** impact
- Single command; **install mode only**

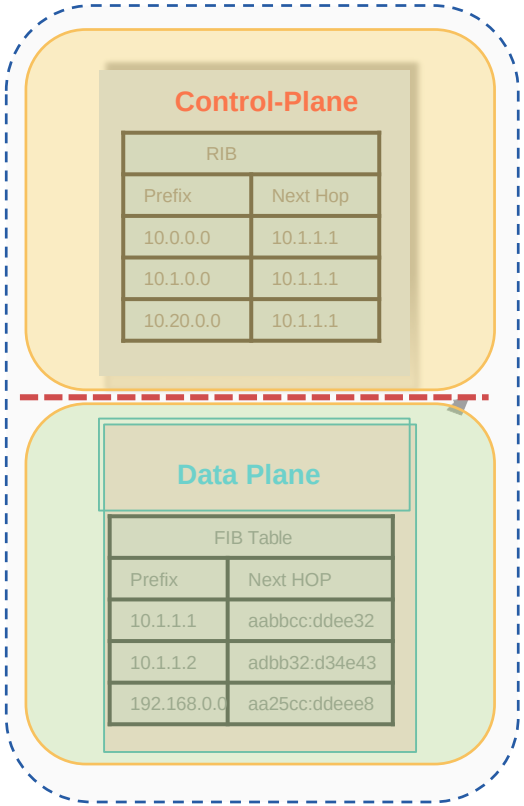

Extended Fast Software Upgrade

9300 standalone

```
#install add file image activate reloadfast commit
```

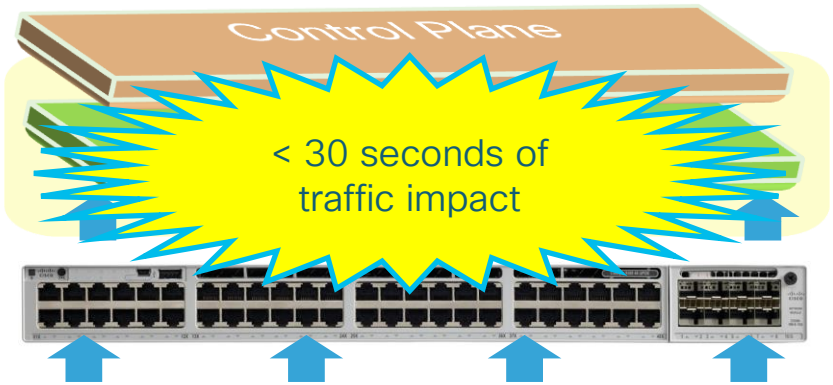

Extended Fast Software Upgrade

9300 stack

```
#install add file image activate reloadfast commit
```

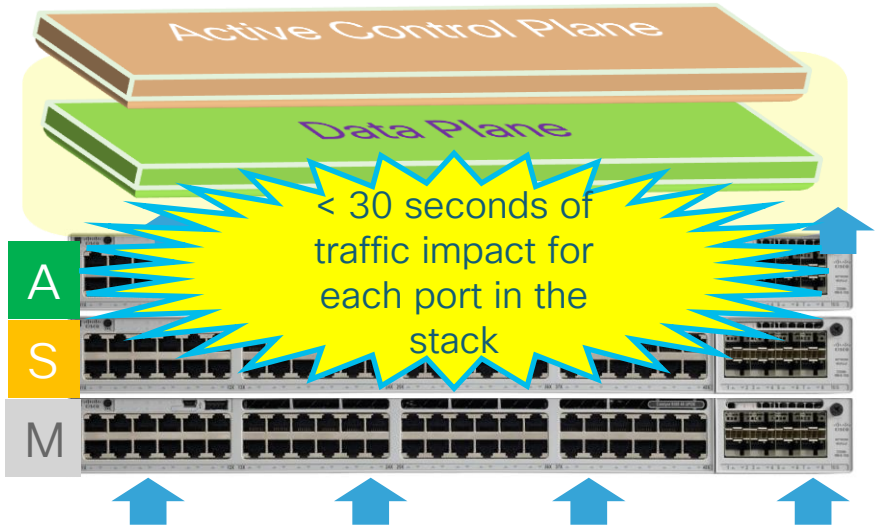

Extended Fast Software Upgrade on Stack

```
#install add file image activate reloadfast commit
```

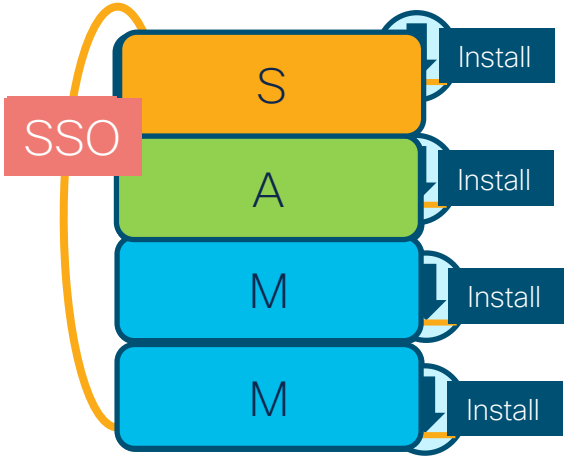

1. Install the images on all switches
2. Fast reload the standby and member switches
3. Fast reload the active switch only
4. Standby becomes the new active
5. Previous Active switch becomes the new standby

Traffic Impact during the complete upgrade is less than 30 seconds

Convergence

xFSU

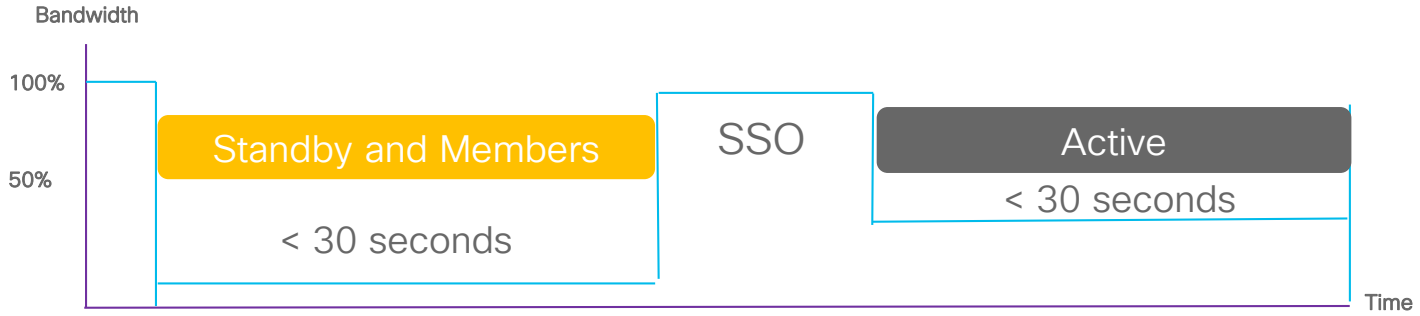

Regular Upgrade

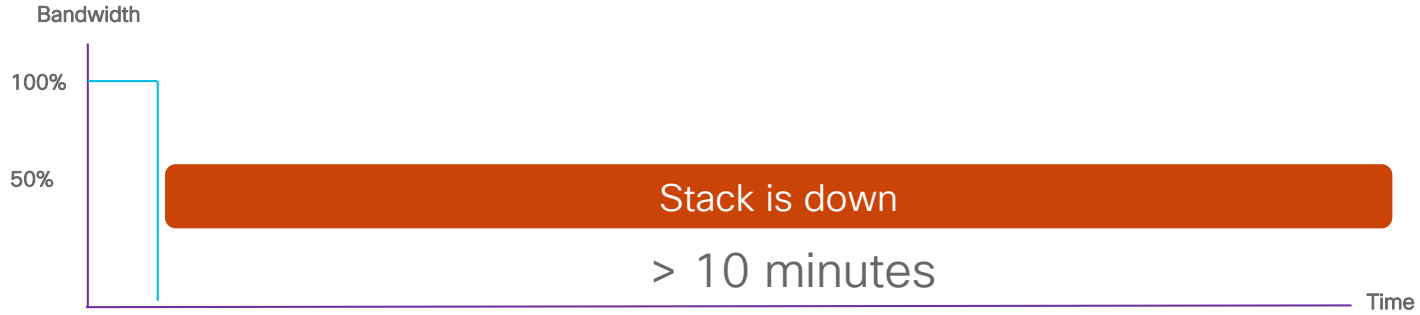

In-Service Software Upgrade (ISSU) Overview

- ISSU provides a mechanism to perform software upgrades and downgrades without taking the switch out of service
- Leverages the capabilities of NSF and SSO to allow the switch to forward traffic during Supervisor upgrade (or downgrade)
- Key technology is the ISSU infrastructure
- Allows SSO between different extended maintenance versions

Modular Catalyst
with dual Supervisors

In-Service Software Upgrade (ISSU) is also supported by NSO

Three-step process:

```
install add file [tftp|ftp|flash|disk:*.bin]
install activate issu
install commit
```

Granular control on the upgrade process with the ability to roll back

One-step process:

```
install add file[tftp|ftp|flash|disk:*.bin] activate issu commit
```

Single command to perform a complete ISSU

Cisco Catalyst 9000 Series ISSU workflow

Reference

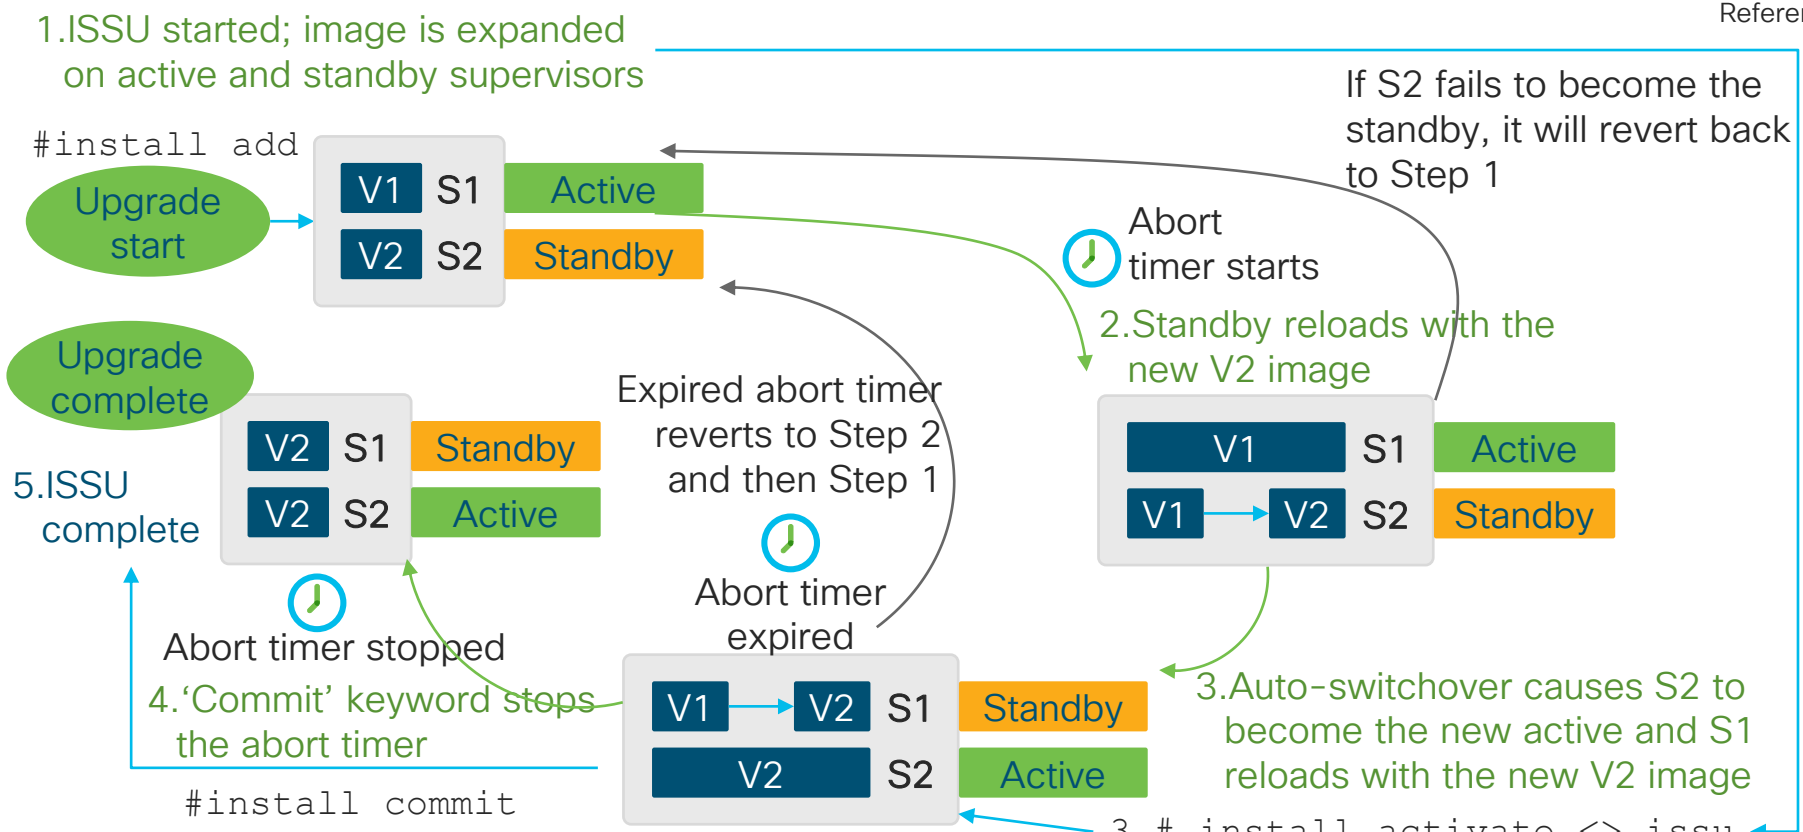

Install command line interface (CLI) commands

Supported in install mode, extended maintenance releases

Reference

Step-by-step workflow:

```
#  install add <tftp://cisco.com/image.bin>
#  install activate issu
```

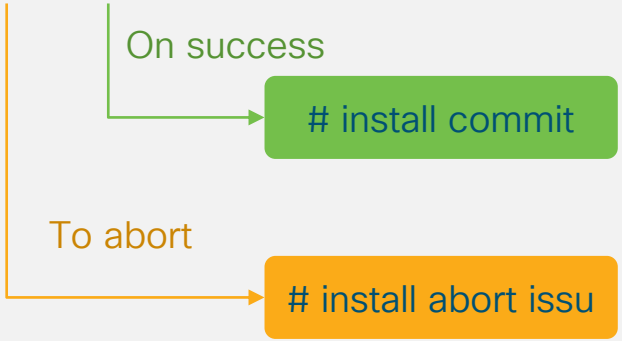

Workflow steps details:

- **`install add`** – performs the image download from the posted location
- **`install activate`** – upgrades the chassis with a new software version
- **`install commit`** – makes the changes permanent and deletes the older version of software from the chassis
- **`install abort issu`** – The operator can issue the abort command to revert the software back to the original state

Non-Stop Routing (NSR)

- Cisco IOS-XE Non-Stop Routing preserves the full state information (prefixes and related data) in the Routing Information Base across Supervisor Engine (Route Processor) switchover events.
- Avoids reconvergence with peer (versus NSF, which delays during grace time)
- Good for peer config not under your control (Example: CE attached to PE environment)
- Consumes more resources than NSF (memory, CPU)
- Device can also use NSR selectively (peering with P/PE/RR/other CE devices) to reduce resource consumption
- Available on some NX-OS and IOS-XE platforms (future for Catalyst 9000 Series)

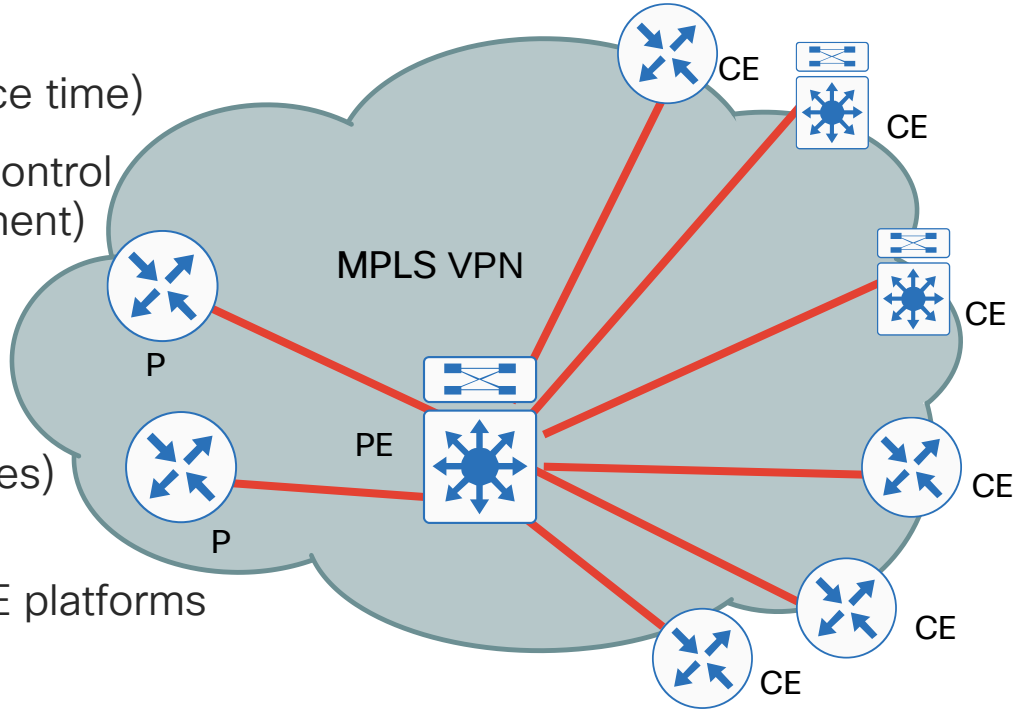

Software Maintenance Update (SMU)

- SMU: An emergency “point fix” for expedited delivery to an organization
- SMUs are:
  - Quick (deliver point fixes much faster than regular IOS-XE software release)
  - Effective (do not require a monolithic IOS-XE code upgrade)
  - Focused (target the specific area of concern in the IOS-XE code)
- Types:
  - Hot patching: no system reload required
  - Cold patching: requires reload
- SMU is like medication:
  - Addresses the issue effectively
  - In theory – no limit to the number you can take
  - In practice – you probably should be selective and minimize the amount

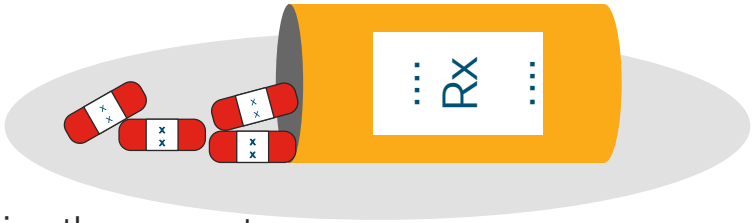

Graceful Insertion and Removal (GIR)

Planning changes

- Isolating a node with minimal impact
  - Hardware replacement
  - Software upgrades
  - Configuration changes
- Can be done manually with CLI/scripts or using GIR commands (customizable):

```
start maintenance

stop maintenance
```

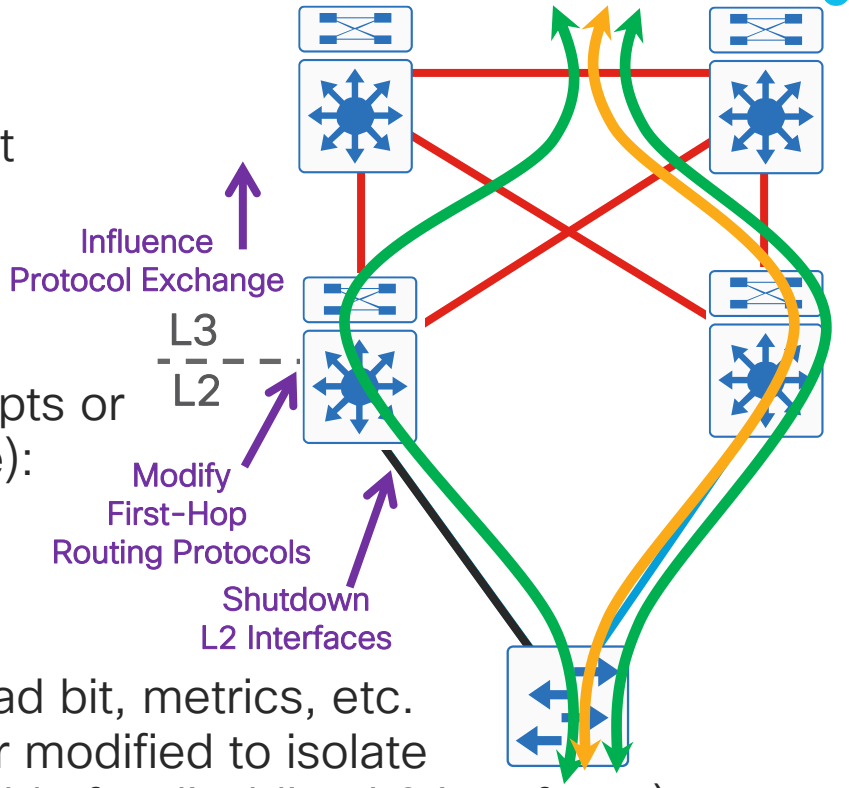

- L3 protocols influenced with overload bit, metrics, etc.
  L2 protocols (HSRP/VRRP) behavior modified to isolate
  (NX-OS has shutdown option available for disabling L2 interfaces)

*Future on Catalyst 9600

Graceful Insertion and Removal

Maintenance Profile

- Contains a sequence of CLI commands to be applied sequentially
- 2 mandatory sub-sections for Maintenance Profile:
  - Normal-Mode section: CLIs to execute when entering Normal Mode
  - Maintenance-Mode section: CLIs to execute when entering Maintenance Mode

```
switch# show maintenance profile
[Normal Mode]
router bgp 100
  no isolate
router eigrp 100
  no isolate
router ospf 100
  no isolate
router isis 100
  no isolate
[Maintenance Mode]
router bgp 100
  isolate
router eigrp 100
  isolate
router ospf 100
  isolate
router isis 100
  isolate
switch#
```

Graceful Insertion and Removal

Maintenance Profile can be generated in two possible ways


Reference

- Custom or user-defined
- User can define a new profile with any set of configuration commands in it
- User can update any existing profile (system-generated or user-defined)
- Useful for dealing with protocols not supported with “isolate” mode

- Automatically system-generated
- System generates automatically during CLI execution:
  `[no] system mode maintenance`
- `system mode …`
  generates Maintenance Mode section
- `no system mode …`
  generates Normal Mode section

Summary: Campus high availability using the Catalyst 9000 Series modular chassis

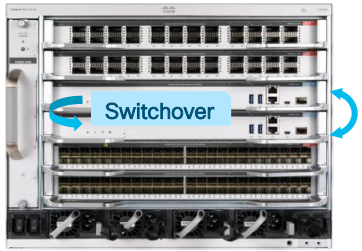

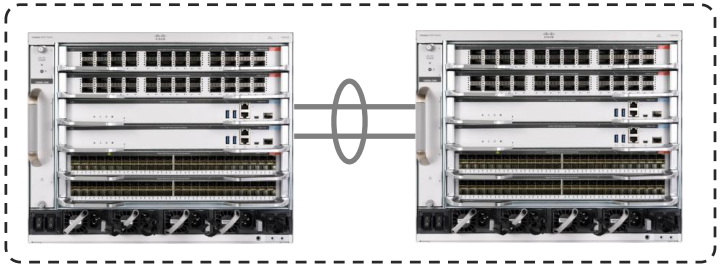

Physical redundancy	Stateful Switchover (SSO)	Non-Stop Forwarding (NSF)	In-Service Software Upgrade (ISSU)	Cisco StackWise Virtual
Redundant hardware • Power supplies • Fans (in tray) • Supervisors • Line cards	**Sub-second failover** • In chassis <5ms between Sups • Between chassis: Cisco StackWise-Virtual	**Resilient L3 topologies** • NSF support for OSPF, EIGRP, ISIS, BGP	**Minimize upgrade downtime** • SMU • ISSU • GIR (9600 future)	**Infrastructure resilience** • Multi-chassis EtherChannel (MEC) provides hardware-based failover

Summary: Using the platform features

What is the recommendation?

Option \ Situation	Critical Bug Fix & PSIRT	Hardware Upgrade	New Image Version
SMU Patching	✹	X	X
ISSU	✓	X	✹
GIR	X	✹	X
Box reload (Cold Boot)	✓	X	✓

Recommended	✹
Possible	✓
Not recommended	X

Agenda

- What is high availability?
- **Campus network foundations and structured design**
  - Wired campus platform hardware and software features for HA
  - **Overview of campus structured design**
- Campus wired LAN design and high availability
- Campus wireless LAN design and high availability
- Summary and conclusions

Hierarchical network design

High availability using hierarchy, modularity, and structure

- Hierarchical Design
  Each layer in hierarchy has a specific role
- Modular Design
  Modularity makes it easy to grow, understand, and troubleshoot
- Structured Design
  Creates small fault domains and predictable network behavior
  —clear demarcations and isolation
- Promotes load balancing and resilience

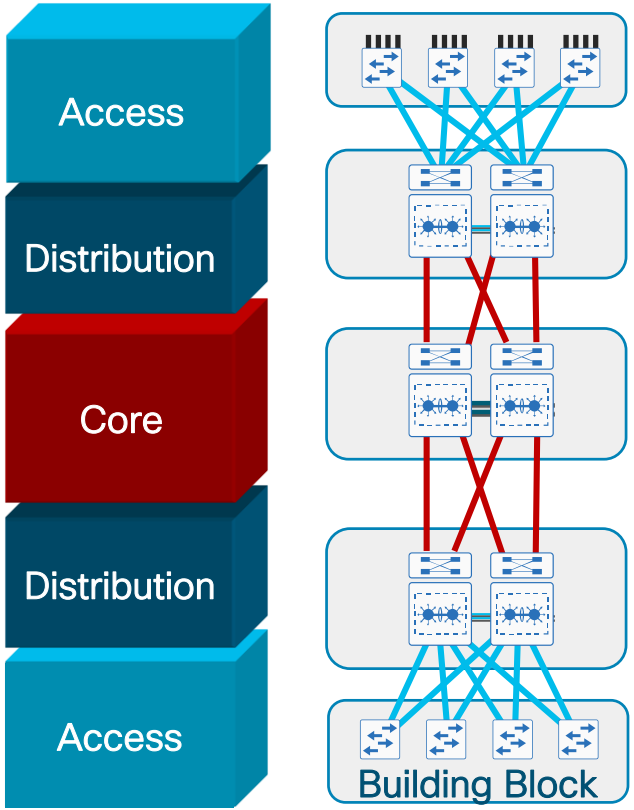

Hierarchical network design: Campus wired LAN

- Core
  - Connectivity, availability and scalability
- Distribution
  - Aggregation for wiring and traffic flows
  - Policy and network control point (FHRP, L3 summarization)
- Access
  - **Physical** – Ethernet wired 10/100/1000(802.3z)/mGig(802.3bz); 802.3af(PoE), 802.3at(PoE+), and Cisco Universal POE (UPOE)
  - **Policy enforcement** – security: 802.1x, port security, DAI, IPSG, DHCP snooping; identification: CDP/LLDP; QoS: policing, marking, queuing
  - **Traffic control** – IGMP snooping, broadcast control

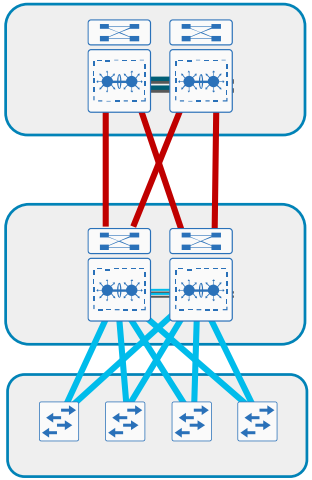

Hierarchical network design: Campus wired LAN

Do I need a core layer?

- It is a question of operational complexity and a question of scale
  - n x (n-1) scaling
  - Routing peers
  - Fiber, line cards, and port counts ($,€,£)

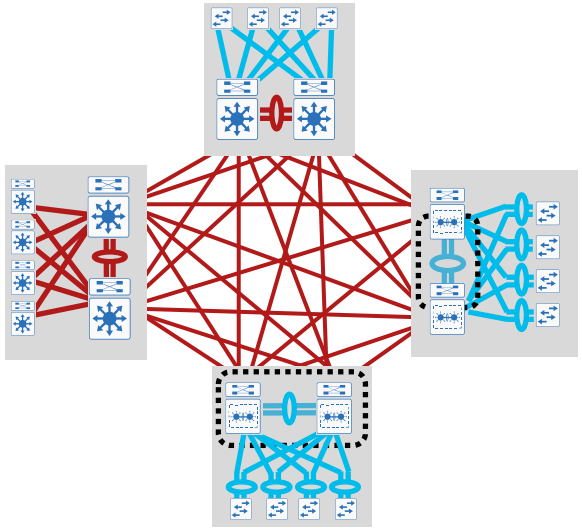

Hierarchical network design: Campus wired LAN

Do I need a core layer?

- It is a question of operational complexity and a question of scale
  - n x (n-1) scaling
  - Routing peers
  - Fiber, line cards, and port counts ($,€,£)
- Capacity planning considerations
  - Easier to track traffic flows from a block to the common core than to 'n' other blocks
- Geographic factors may also influence the design
  - Multi-building interconnections may have fiber limitations

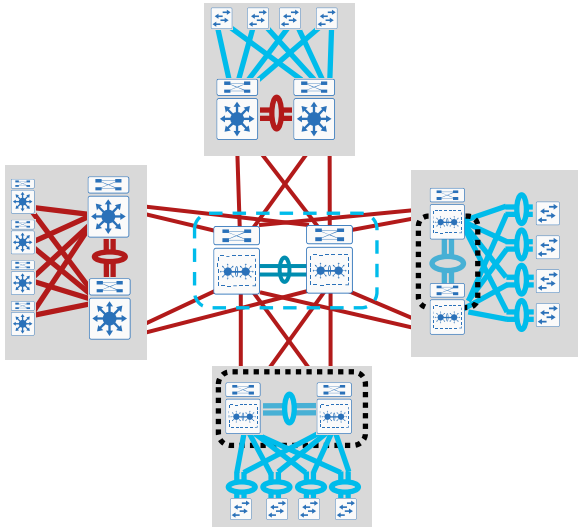

Chassis Redundancy at the Core

Depends on topology

- Redundant topologies with equal cost multi-paths (ECMP) provide sub-second convergence
- NSF/SSO provides superior availability in environments with non-redundant paths

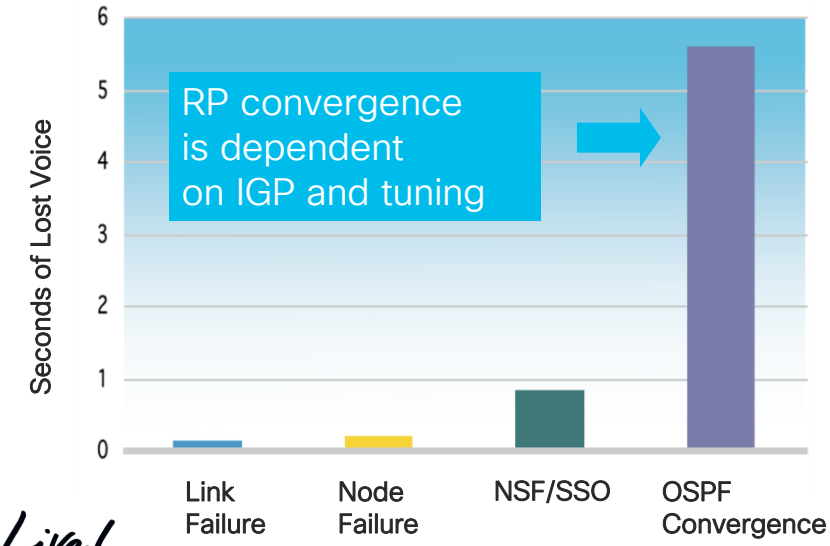

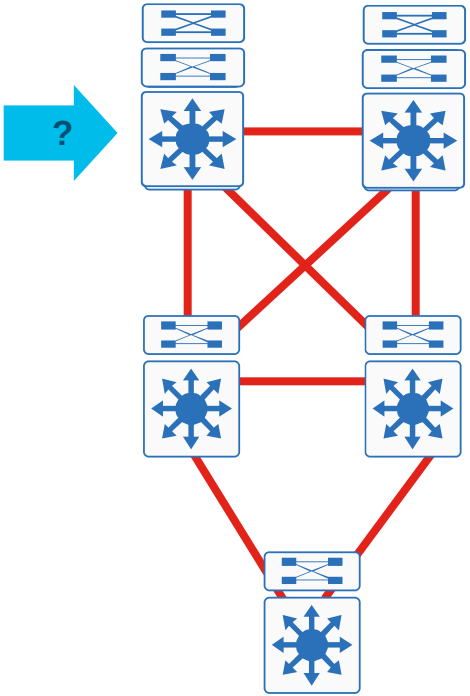

Chassis Redundacny at the Distribution

Recommended

- HSRP doesn't flap on Supervisor SSO switchover
- Reduces the need for sub-second HSRP timers

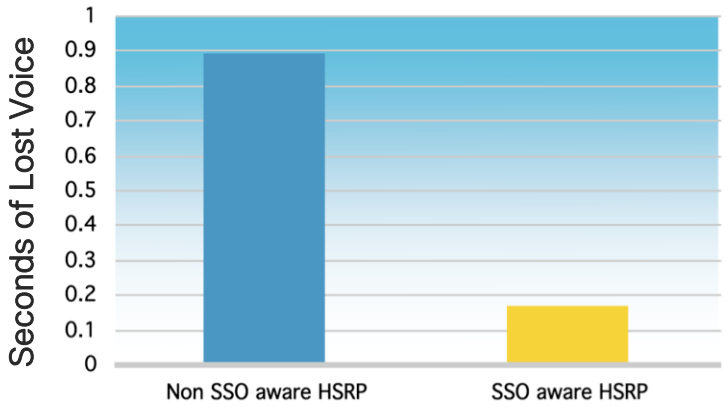

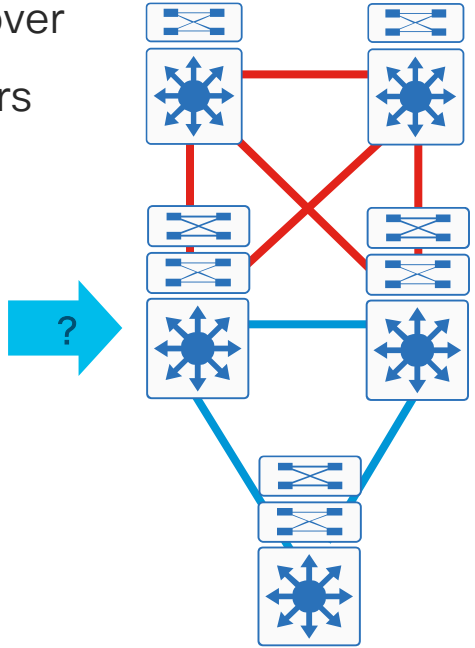

Chassis Redundancy at the Access

Recommended for highest availability

- Access switch is the single point of failure in best practices HA design
- Supervisor failure is most common cause of access switch service outages

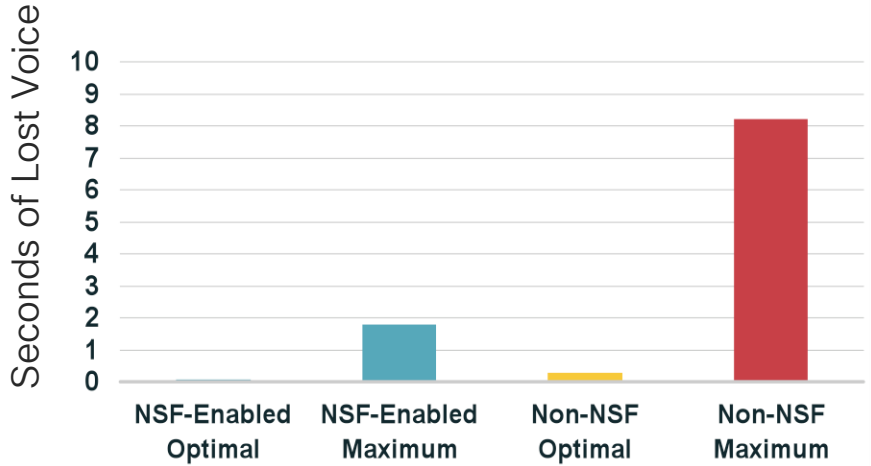

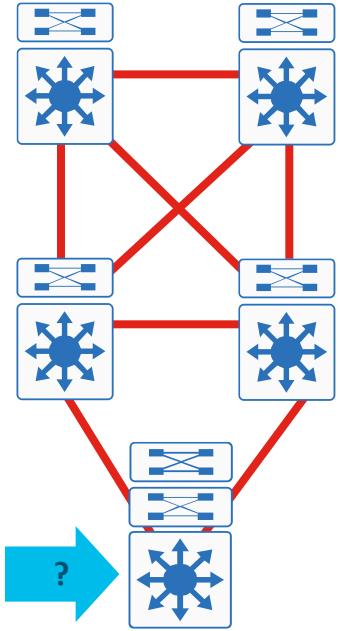

High availability design optimization of the elements

- **Optimize the interaction** of the physical redundancy with the network protocols
  - Provide the necessary amount of redundancy
  - Pick the right protocol for the requirement
  - Optimize the tuning of the protocol
- The network looks like this so that we can map the protocols onto the physical topology

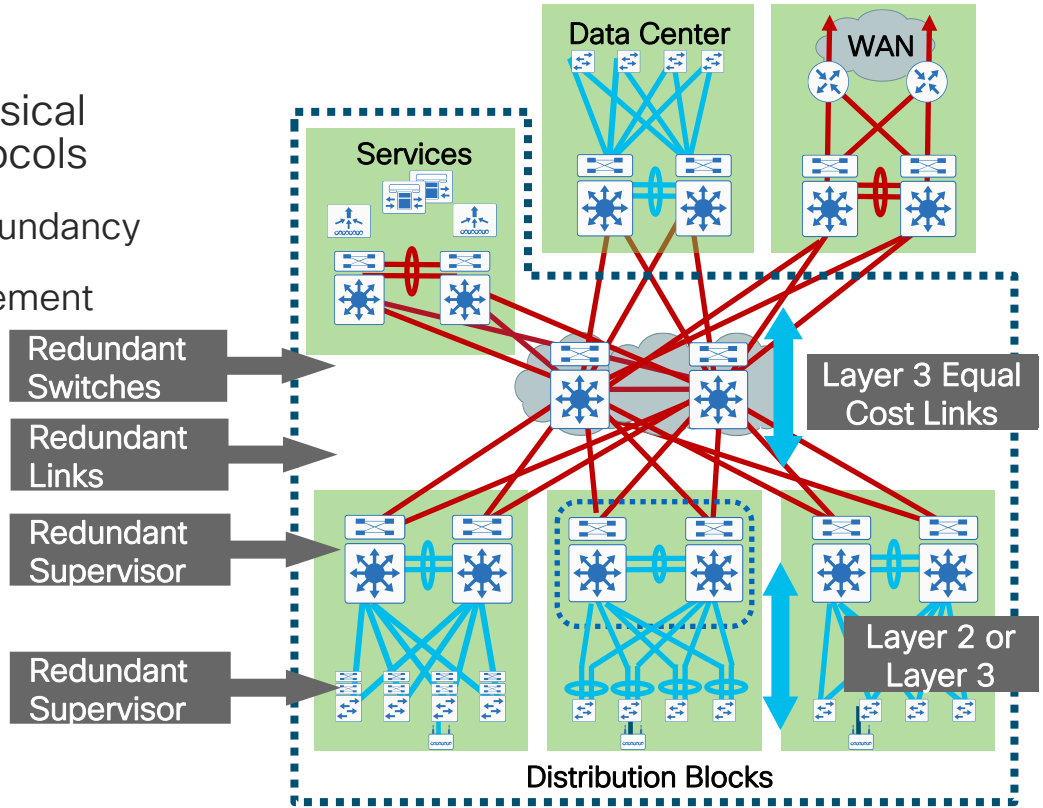

Strive to build networks that look like this.

What we are trying to avoid!

No hierarchy

Multiple single points of failure

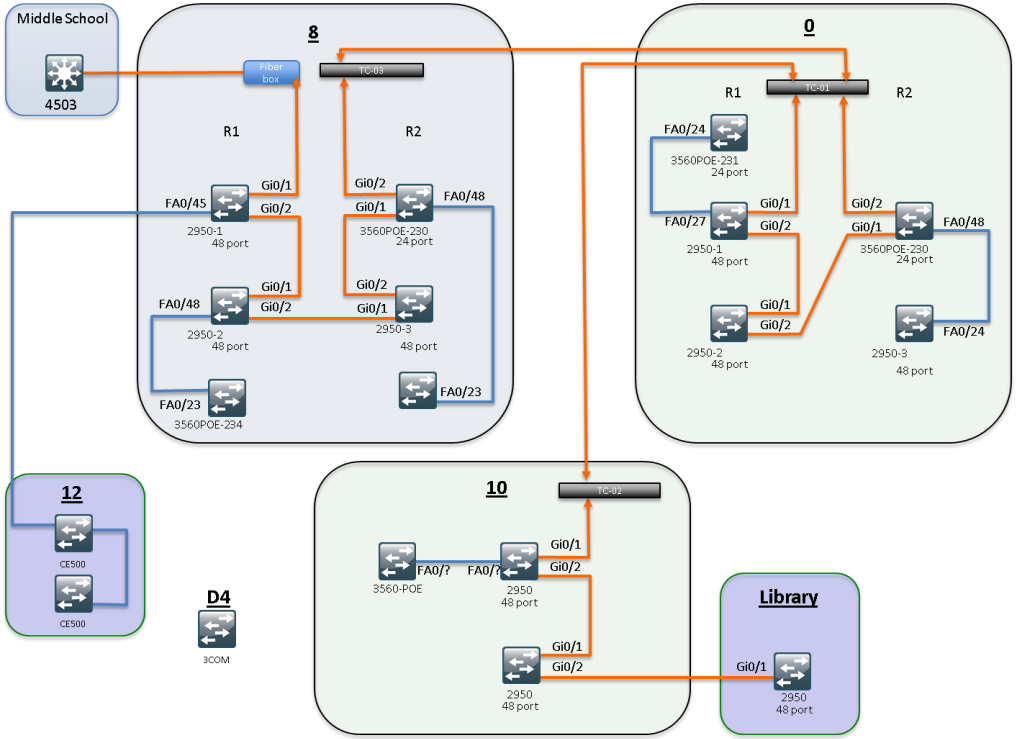

Hard to troubleshoot

Poor performance

Campus wired LAN design and high availability

Agenda

- What is high availability?
- Campus network foundations and structured design
- **Campus wired LAN design and high availability**
  - **Connecting the devices**
  - Considerations with the traditional multilayer campus design
  - Layer-3 access design
  - Layer-2 and simplified distribution design
  - Routed access design
  - New requirements driving new options for campus design
- Campus wireless LAN design and high availability
- Summary and conclusions

How do I choose what to build?

- Principles:
  Ease of deployment
  flexibility, scalability, security
- Hierarchical model:
  resiliency
  modularity
  load balancing
- Devices?
- Capabilities?
- Connectivity and resiliency?

Structured campus network design

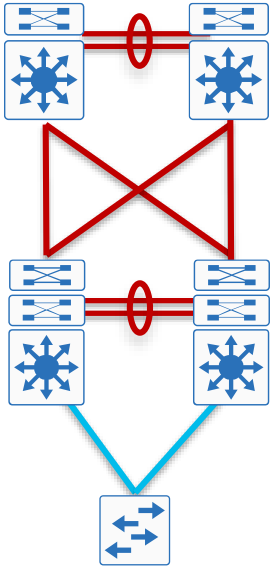

Core

Distribution

Access

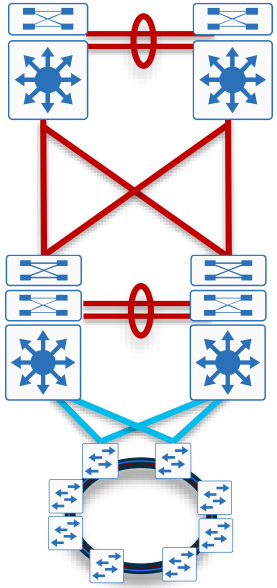

Core

Distribution

Access

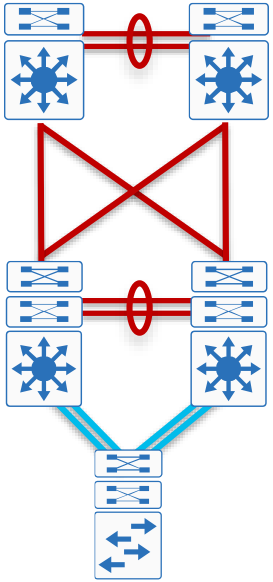

- Optimize data load-sharing, redundancy design for best application performance
  - Diversify uplink network paths with cross-stack and dual-sup access-layer switches
  - Build distributed and full-mesh network paths between Distribution and Access-layer switches

Optimizing network convergence
Failure detection and recovery

- Optimal high availability network design attempts to leverage ‘local’ switch fault detection and recovery
- Design should leverage the hardware capabilities of the switches to detect and recover traffic flows based on these ‘local’ events
- Design principle –
  Hardware failure detection and recovery is both faster and more deterministic
- Design principle –
  Software failure detection mechanisms provide a secondary, not primary, fault detection and recovery mechanism in the optimal design

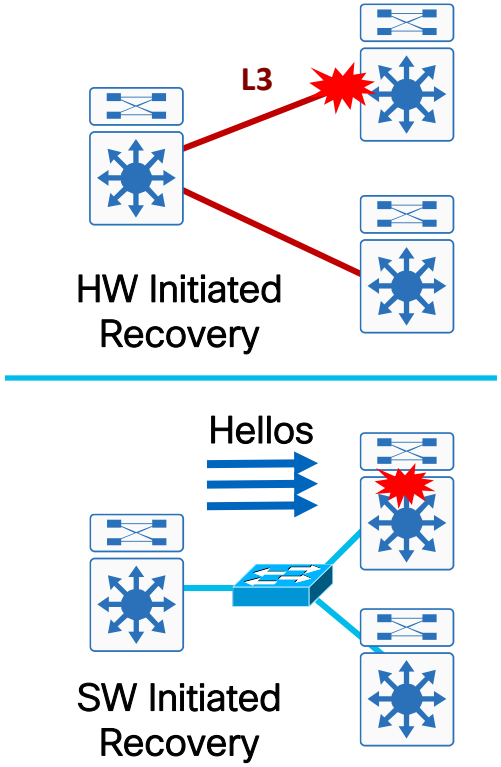

Optimizing network convergence

Layer 1 link failure fault detection

- Do not disable auto-negotiation on GigE /10GigE ports
- IEEE 802.3z and 802.3ae link negotiation define Remote Fault Indicator & Link Fault Signaling mechanisms
- IOS debounce –
  - GigE/10GigE fiber ports is 10 msec.; copper min. 300 msec.
  - NX-OS debounce – Currently 100 msec. by default
  - All 1G and 10G SFP / SFP+ based interfaces (MM, SM, CX-1) changing to a default of 10 msec.
  - RJ45 based Copper interfaces on NX-OS remains 100 msec.
- Design principle: Understand how hardware choices and tuning impact

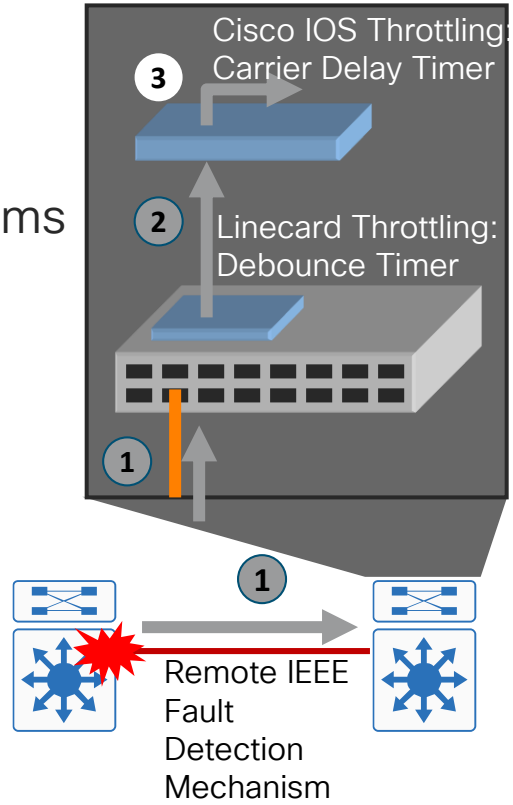

Optimizing network convergence
Layer 2 software fault detection (e.g. UDLD)

- While 802.3z and 802.3ae link negotiation provide for L1 fault detection, hardware ASIC failures can still occur
- UDLD – L2 based keep-alive mechanism confirms bi-directional L2 connectivity
- Switch ports with UDLD send UDLD protocol packets (at L2) containing:
  port's own device / port ID
  neighbor's device / port IDs seen by UDLD on that port
- If port does not see its own device / port ID echoed by incoming UDLD packets, the link is considered unidirectional and is shutdown
- Design principle –
  Redundant fault detection mechanisms required
  (SW as a backup to HW as possible)

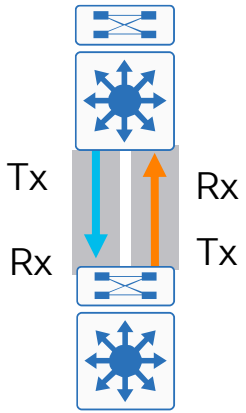

UDLD Keepalive

Optimizing network convergence Layer 2 and 3 – Why use routed interfaces?

L3 routed interfaces allow faster convergence than L2 switchport with an associated L3 SVI

~ 8 msec loss

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


```
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
```

~ 200-250 msec. loss

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


```
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
```

Agenda

- What is high availability?
- Campus network foundations and structured design
- **Campus wired LAN design and high availability**
  - Connecting the devices
  - **Considerations with the traditional multilayer campus design**
  - Layer-3 access design
  - Layer-2 and simplified distribution design
  - New requirements driving new options for campus design
- Campus wireless LAN design and high availability
- Summary and conclusions

Optimizing the Layer 2 design – spanning tree

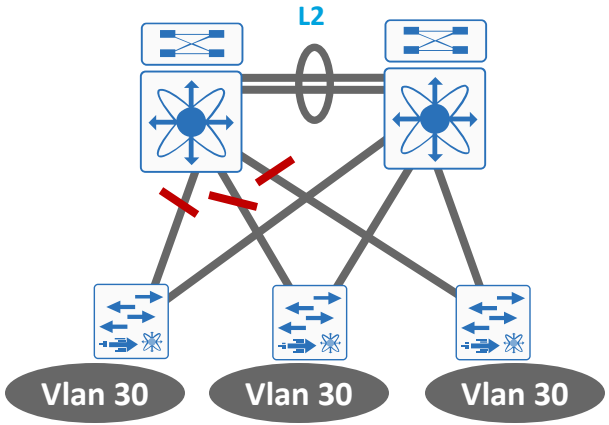

- At least some VLANs span multiple access switches
- Layer 2 loops
- Layer 2 and 3 running over link between distribution
- Blocked links
- More typical of a “classic” data center design

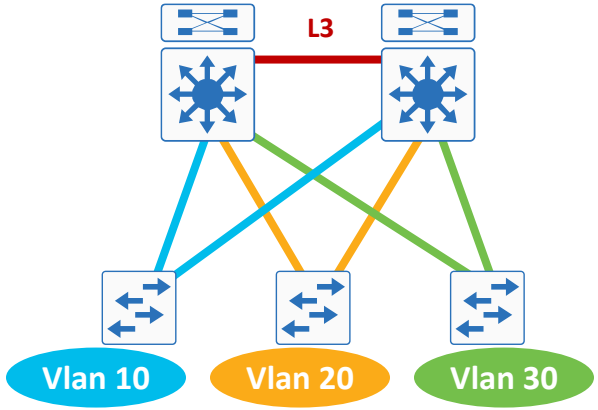

- Each access switch has unique VLANs
- No Layer 2 loops
- Layer 3 link between distribution
- No blocked links
- More typical of a campus LAN design

Optimizing the Layer 2 design
Non-STP-blocking topologies converge fastest

- When STP is not blocking uplinks, recovery of access to distribution link failures is accomplished **based on L2 CAM updates** not on the Spanning Tree protocol recovery
- Time to restore traffic flows is based on: Time to detect link failure + Time to purge the HW CAM table and begin to flood the traffic
- No dependence on external events (no need to wait for Spanning Tree convergence)
- Behavior is **deterministic**

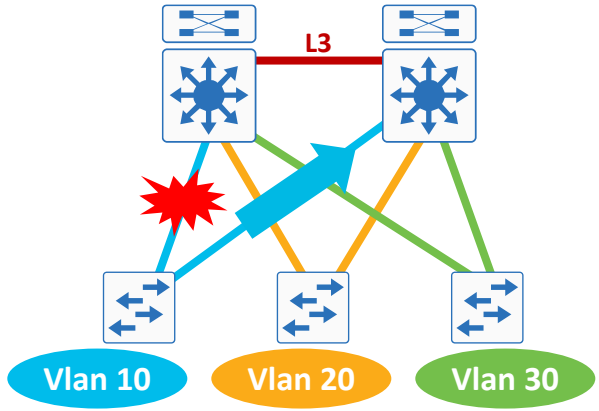

- All links rorwarding – In an environment with all Links active, traffic is restored based on **HW recovery**

Optimizing the Layer 2 design PVST+, Rapid PVST+, MST

- PVST+ (pre 802.1D-2004) - traditional spanning tree
- Rapid-PVST+ (802.1w)
  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 BackboneFast for any indirect link failures
- Rapid PVST+
  Scales to large size (up to 16,000 logical ports)
  Easy to implement, proven, scales
- MST (802.1s)
  Permits very large scale STP implementations (up to 75,000 logical ports)

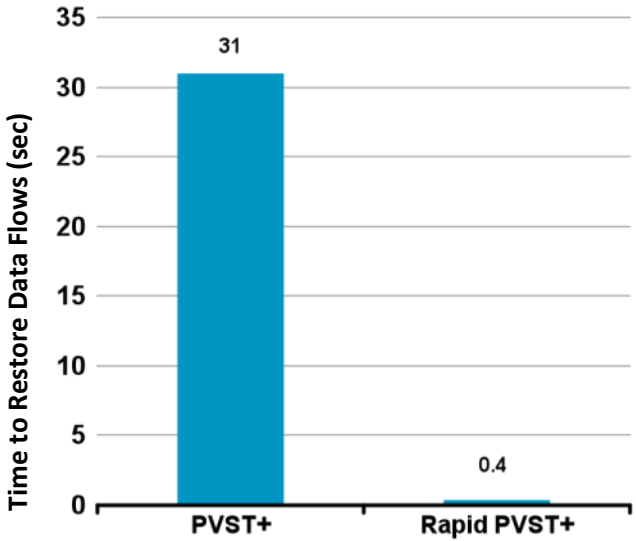

Optimizing the Layer 2 design
STP toolkit – PortFast and BPDU guard

- PortFast is configured on edge ports to allow them to quickly move to forwarding bypassing listening and learning and avoids TCN (Topology Change Notification) messages
- BPDU guard can prevent loops by moving PortFast configured interfaces that receive BPDUs to errdisable state
- BPDU guard prevents ports configured with PortFast from being incorrectly connected to another switch
- When enabled globally, BPDU guard applies to all interfaces that are in an operational PortFast state

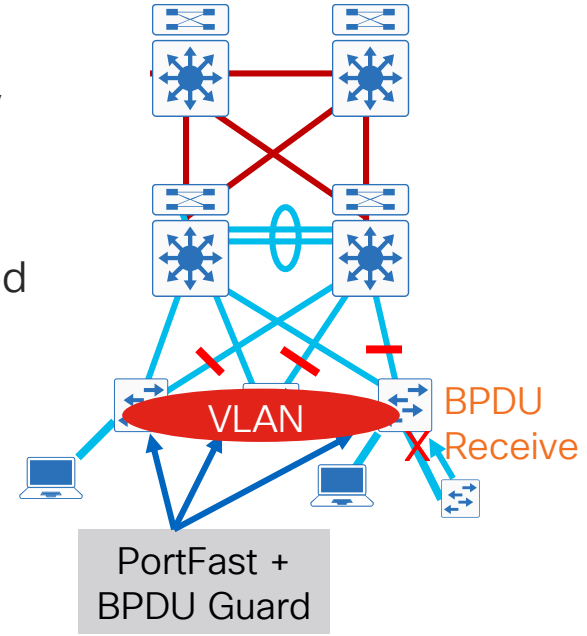


```
Switch(config-if)#spanning-tree portfast
Switch(config-if)#spanning-tree bpduguard enable
```

```
1w2d: %SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port FastEthernet3/1 with BPDU Guard enabled. Disabling port.
1w2d: %PM-4-ERR_DISABLE: bpduguard error detected on Fa3/1, putting Fa3/1 in err-disable state
```

Optimizing the Layer 2 design
STP best practices for campus

Reference

- The root bridge should stay where you put it
  - Define the STP primary (and backup) root
  - Rootguard
  - Loopguard or bridge assurance
  - UDLD
- There is a reasonable limit to broadcast and multicast traffic volumes
- Configure storm control on backup links to aggressively rate limit broadcast and multicast

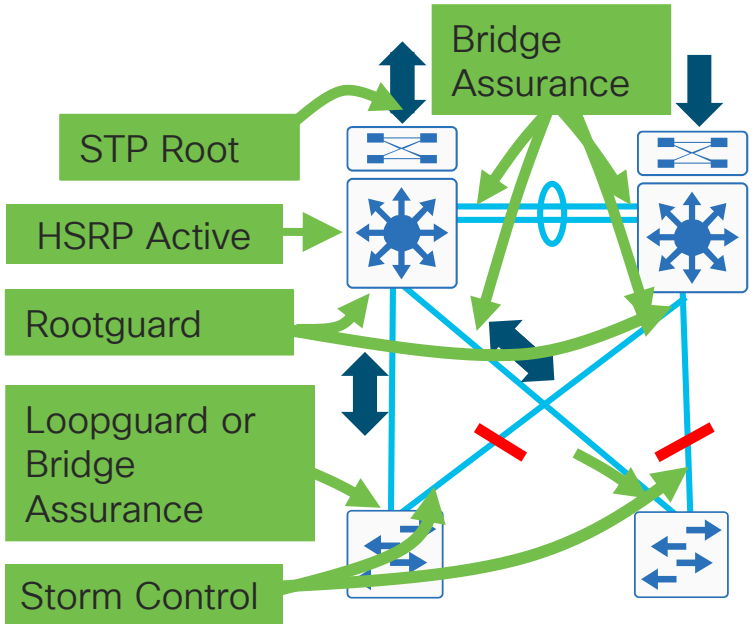

Layer 2 access with Layer 3 distribution First hop redundancy protocols (FHRP)

- HSRP, GLBP, and VRRP:
  provide a resilient default gateway / first hop address to end stations
- A group of routers act as a single logical router providing first hop router redundancy
- Protect against multiple failures
  - Distribution switch failure
  - Uplink failure
- Default recovery is ~10 Seconds

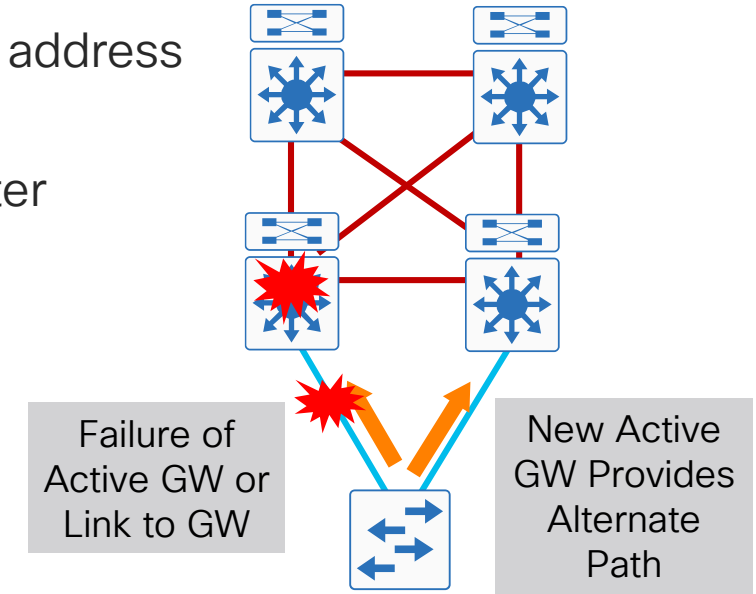

First hop redundancy

Subsecond timers improve convergence

HSRP Config

```
interface Vlan4
 ip address 10.120.4.2 255.255.255.0
 standby 1 ip 10.120.4.1
 standby 1 timers msec 250 msec 750
 standby 1 priority 150
 standby 1 preempt
 standby 1 preempt delay minimum 180
```

GLBP Config

```
interface Vlan4
 ip address 10.120.4.2 255.255.255.0
 glbp 1 ip 10.120.4.1
 glbp 1 timers msec 250 msec 750
 glbp 1 priority 150
 glbp 1 preempt
 glbp 1 preempt delay minimum 180
```

VRRP Config

```
interface Vlan4
 ip address 10.120.4.1 255.255.255.0
 vrrp 1 description Master VRRP
 vrrp 1 ip 10.120.4.1
 vrrp 1 timers advertise msec 250
 vrrp 1 preempt delay minimum 180
```

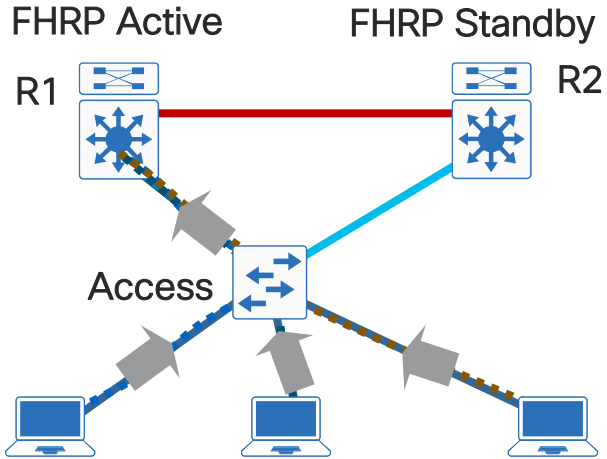

HSRP is widely used with Its rich feature set

GLBP facilitates uplink load balancing – not optimal for L2 looped topology

VRRP for multi-vendor interoperability

HSRP, GLBP and VRRP provide millisecond timers and excellent convergence performance

Critical for VoIP and video recovery in < 1 second

cisco Live!

HSRP preemption–why it is desirable

- Spanning tree root and HSRP primary are aligned
- When spanning tree root is re-introduced, traffic takes a two-hop path to HSRP active
- **HSRP preemption** allows HSRP to follow the spanning tree topology

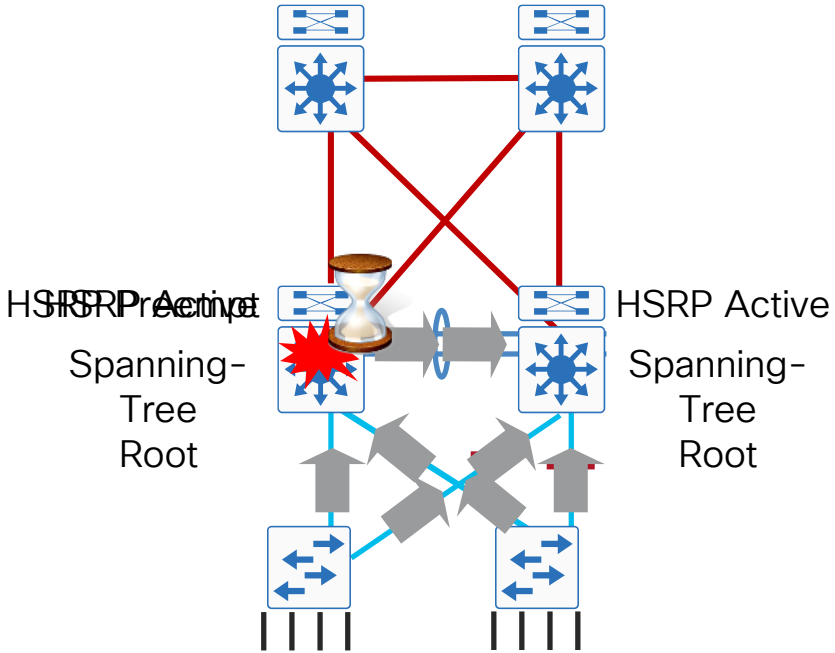

Without Preempt Delay, HSRP Can Go Active Before the Switch Is Completely Ready to Forward Traffic – L1 (Linecards), L2 (STP), L3 (IGP Convergence)

FHRP design considerations
Preempt delay needs to be longer than boot time

- HSRP is not always aware of the status of the entire switch and network
- Ensure that you provide enough time for the entire (full or partial), L1 (line cards), L2 (STP), L3 (IGP convergence)
- Tune delay and preempt delay conservatively, as the network is already forwarding data

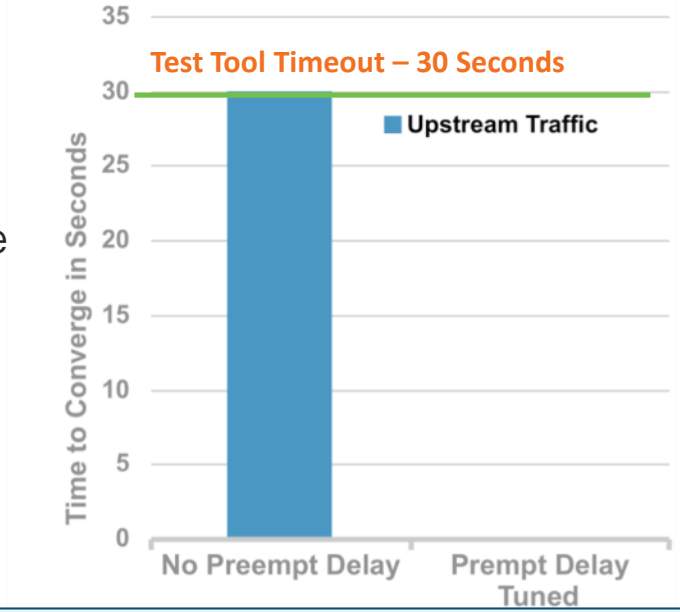


```
interface Vlan402
. . .
 standby delay minimum 60 reload 600
 standby 1 ip 10.147.102.1
 standby 1 timers msec 250 msec 750
 standby 1 priority 110
 standby 1 preempt delay minimum 60 reload 600
 standby 1 authentication ese
 standby 1 name HSRP-Voice
 hold-queue 2048 in
```

`standby delay`: Controls time interface needs to be up before HSRP starts.
`preempt delay`: Controls time to wait after HSRP establishes a neighbour relationship.
Configure both.

Sub-second timer considerations HSRP, GLBP, OSPF, PIM

- Evaluate your network before implementing any sub-second timers
- Certain events can impact the ability of the switch to process sub-second timers
  - Application of large ACL
  - OIR of line cards in Catalyst 6500/6800
- Control plane traffic volume also impacts ability to process
  - 250 / 750 msec GLBP & HSRP timers are only valid in designs with less than 150 VLAN instances (Catalyst 6x00 in the distribution)
  - Spanning Tree size

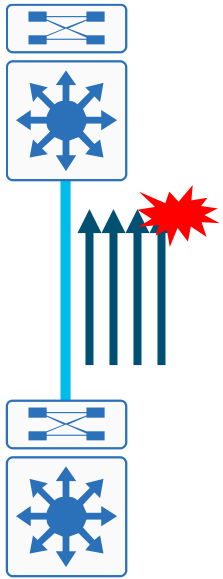

FHRP design considerations– asymmetric routing (unicast flooding)

- Alternating HSRP Active between distribution switches can be used for upstream load balancing
- This can cause a problem with unicast flooding
- ARP timer defaults to four hours and CAM timer defaults to five minutes
- ARP entry is valid, but no matching L2 CAM table exists
- In many cases when the HSRP standby needs to forward a frame, it will have to unicast flood the frame since its CAM table is empty

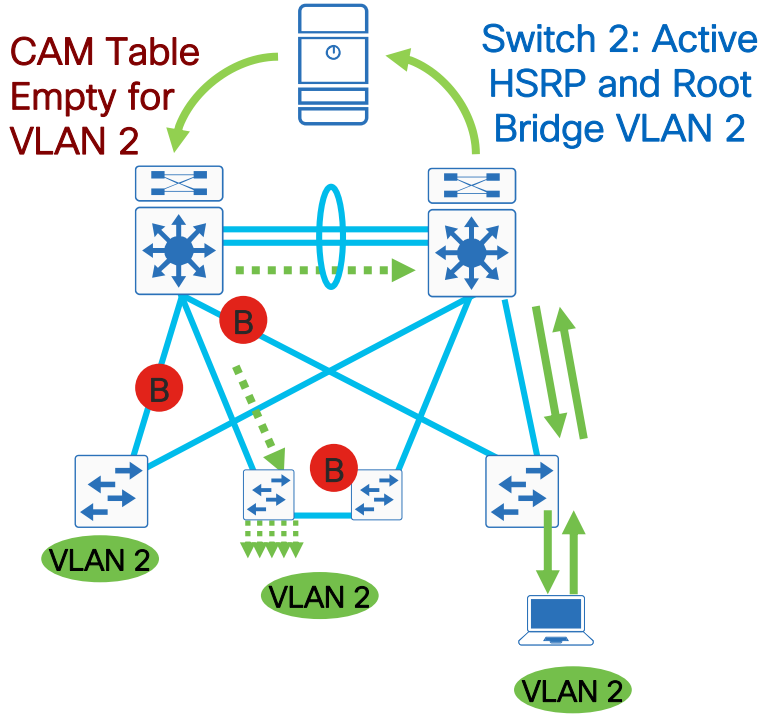

FHRP design considerations– asymmetric routing (unicast flooding) solutions

- Using ‘V’ based design with unique voice and data VLANs per access switch, this problem has no user impact
- Don’t deploy stacking switches (ie. daisy-chained switches) that depend on spanning tree for managing stack interconnects
- Tune ARP timer to 270 seconds and leave CAM timer to default, unless ARP > 10,000, change CAM timers
- Deploy MultiChassis EtherChannel with StackWise Virtual (SWV), Virtual Switching System (VSS), or Virtual Port Channel (vPC) in the distribution block

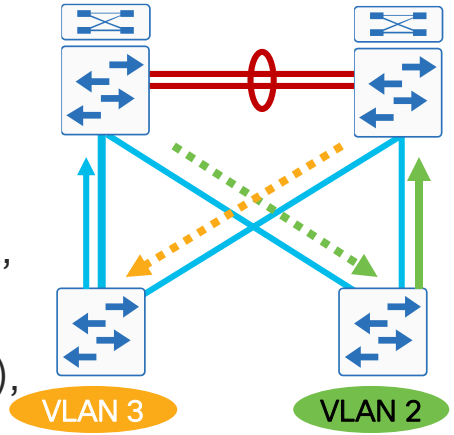

CAM timers traditionally default to 5 minutes to allow for MAC addresses (devices) to move in the network. It is safe to increase the CAM timers if the client devices will generate unicast or multicast traffic to refresh the CAM table.

Even with faster convergence from RPVST+ we still have to wait for FHRP convergence

- FHRP protocol based forwarding topologies
  - Load balancing based on Per-Port or Per-VLAN
- Protocol-based fault detection and recovery –
  - Configure per-VLAN aggressive timers to protect user experience impact within <1 second boundary
- Limited network scale for system reliability
- Sub-second protocol timers must be avoided on SSO capable networks

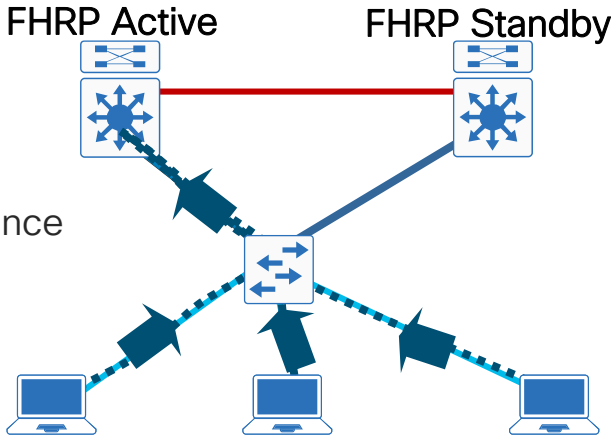

HSRP Config

```
interface Vlan2
 ip address 10.120.2.2 255.255.255.0
 standby 1 ip 10.120.2.1
 standby 1 timers msec 250 msec 750
 standby 1 priority 150
 standby 1 preempt
 standby 1 preempt delay minimum 180
```

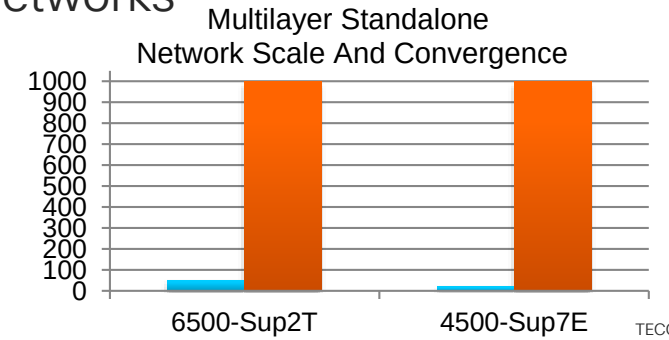

Multilayer campus network design– It is a good solid design, but…

- Utilizes multiple control protocols
  - Spanning tree (802.1w), HSRP / GLBP, EIGRP, OSPF
- Convergence is dependent on multiple factors –
  - FHRP – 900msec to 9 seconds
  - Spanning tree – Up to 50 seconds
- Load balancing –
  - Asymmetric forwarding
  - HSRP / VRRP – per subnet
  - GLBP – per host
- Unicast flooding in looped design
- STP, if it breaks badly, has no inherent mechanism to stop the loop

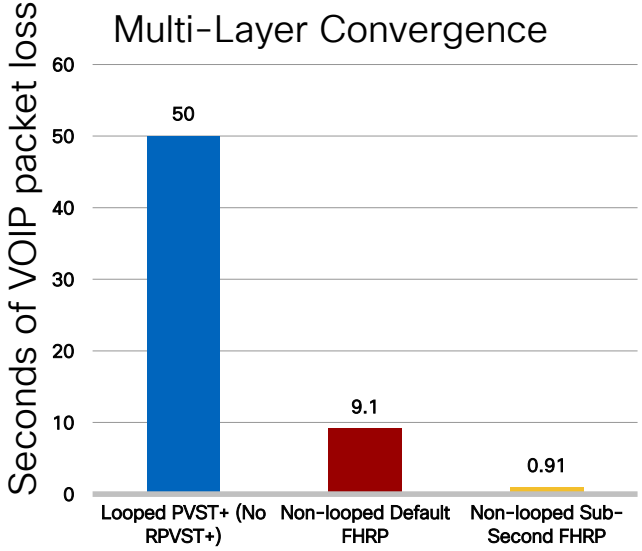

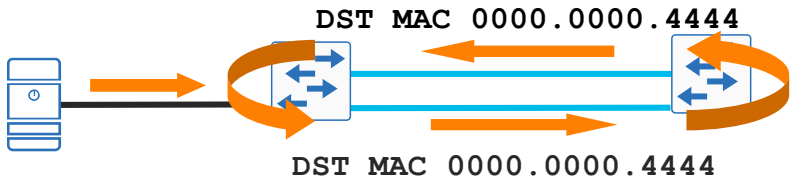

Campus wired LAN design

Option 1: Traditional multilayer campus (BRKCRS-2031)

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

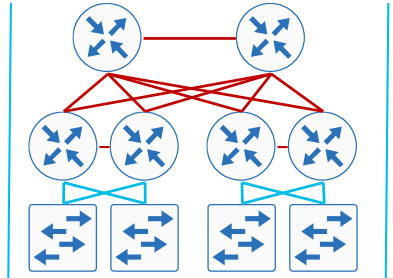

Physical topology: 2 core 2 dist./acc.

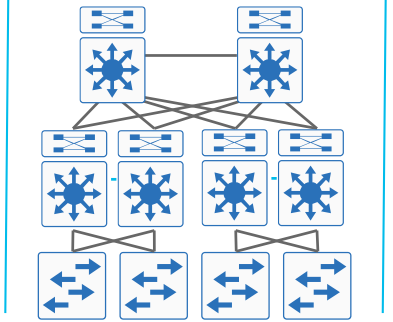

- Common design since the 1990's
- Complex configurations (prone to human error) related to spanning-tree, load balancing, unicast and multicast routing
- Requires heavy performance tuning resulting from reliance on FHRPs (HSRP, VRRP, GLBP)

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	

Agenda

- What is high availability?
- Campus network foundations and structured design
- **Campus wired LAN design and high availability**
  - Connecting the devices
  - Considerations with the traditional multilayer campus design
  - **Layer-3 access design**
  - Layer-2 and simplified distribution design
  - New requirements driving new options for campus design
- Campus wireless LAN design and high availability
- Summary and conclusions

Transforming multilayer campus

Before: Layer 3 distribution with Layer 2 access

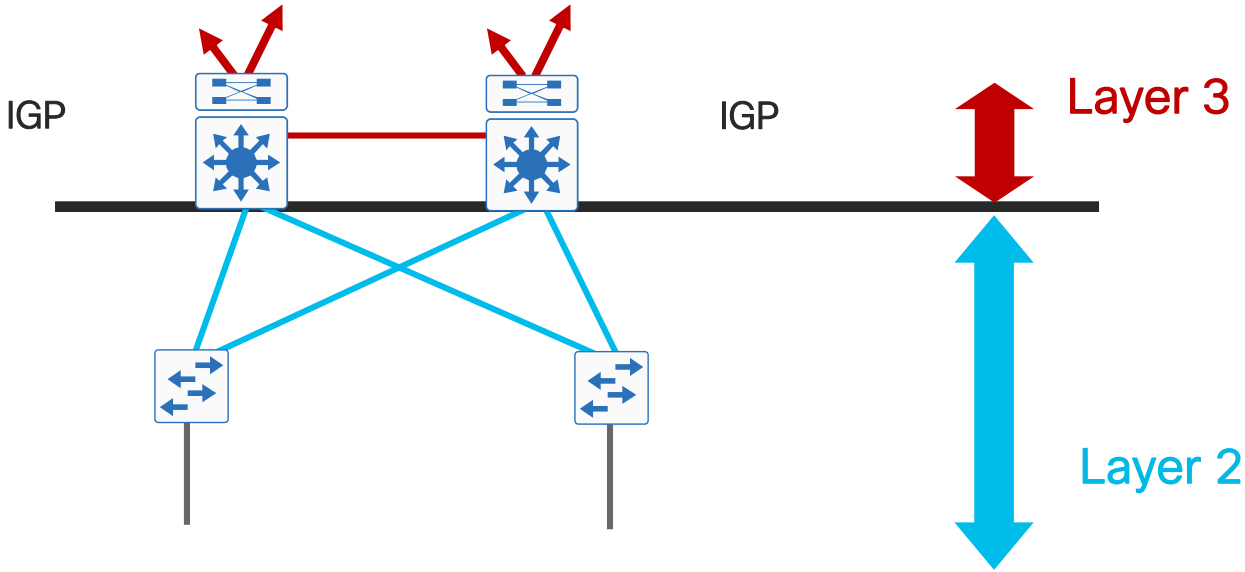

Simplification with routed access design

After: Layer 3 distribution with Layer 3 access

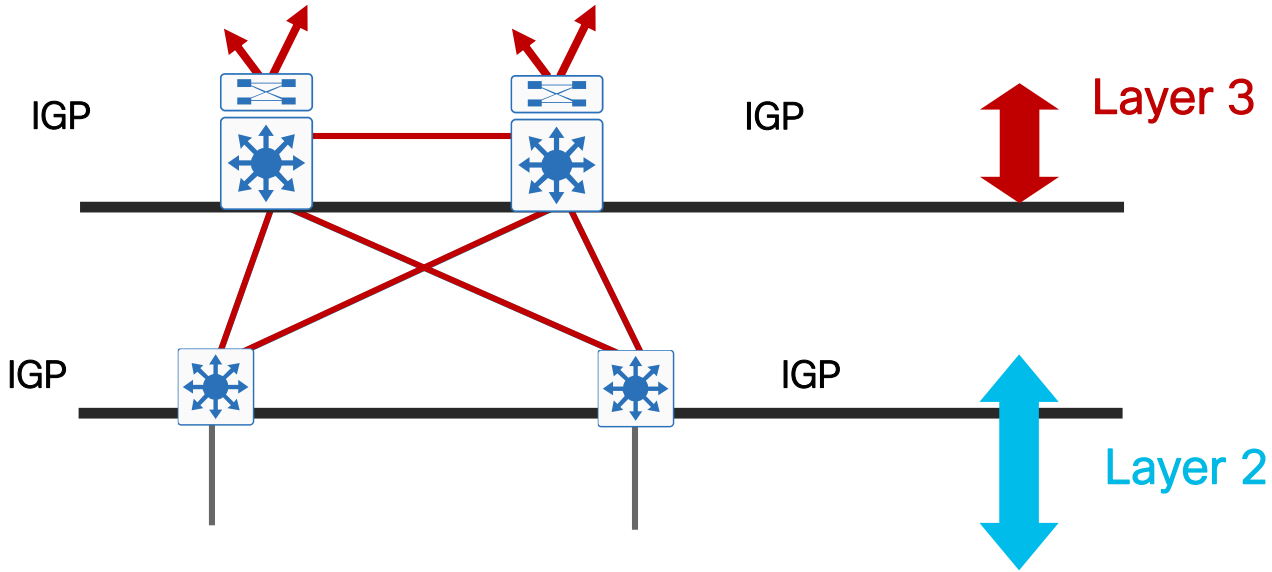

- Move the Layer 2 / 3 demarcation to the network edge
- Leverages Layer 2 only on the access ports, but builds a Layer 2 loop-free network
- **Design motivations** – Simplified control plane, ease of troubleshooting, highest availability

Routed access advantages

Simplified control plane

- Simplified Control Plane
  - **No STP** feature placement (root bridge, loopguard, ...)
  - **No default gateway** redundancy setup/tuning (HSRP, VRRP, GLBP ...)
  - **No matching of STP/HSRP priority**
  - **No asymmetric flooding**
  - **No** L2/L3 **multicast** topology **inconsistencies**
  - **No Trunking** Configuration Required

- L2 Port Edge features still apply:
  - Spanning Tree Portfast
  - Spanning Tree BPDU Guard
  - Port Security, DHCP Snooping, DAI, IPSG
  - Storm Control

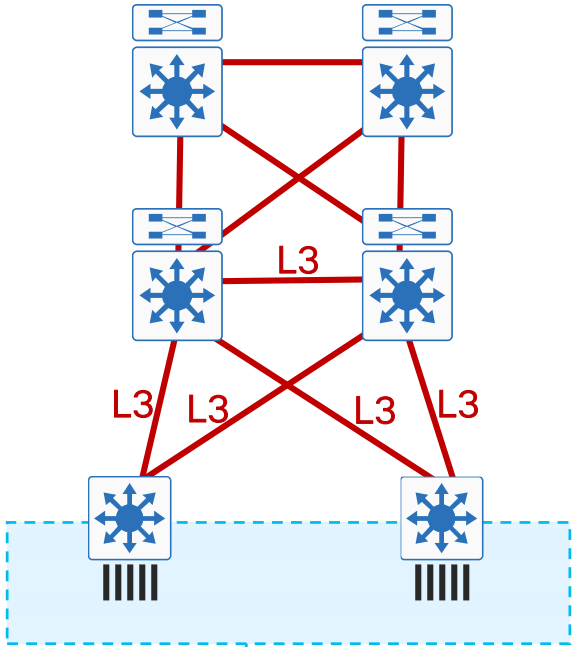

Routed access advantages

Simplified network recovery

- Routed access network recovery is dependent on L3 re-route
- Upstream traffic restoration: ECMP re-route
  - Detect link failure
  - Process SW RIB update
  - Update HW FIB
- Downstream traffic restoration: routing protocol re-route
  - Detect link failure
  - Determine new route
  - Process SW RIB update
  - Update HW FIB

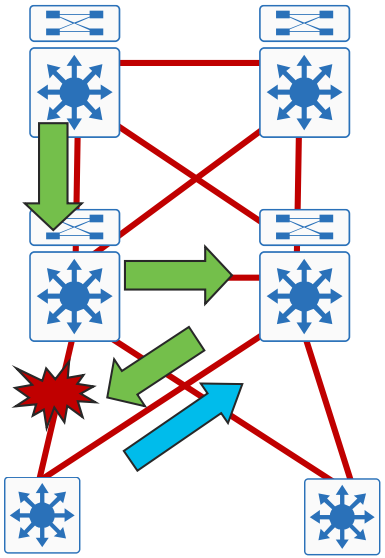

Upstream Recovery: ECMP
Downstream Recovery: Routing Protocol

Compare to…

- RPVST+ convergence times dependent on FHRP tuning
- Proper FHRP design and tuning can achieve sub-second times
- EIGRP converges <200 msec
- OSPF converges <200 msec with LSA and SPF tuning

Why isn’t routed access deployed everywhere?

Routed access design constraints

- VLANs don’t span across multiple wiring closet switches/switch stacks

  Does this impact your requirements?

- IP addressing changes: more DHCP scopes and subnets of smaller sizes increase management and operational complexity
- Deployed access platforms must be able to support routing features

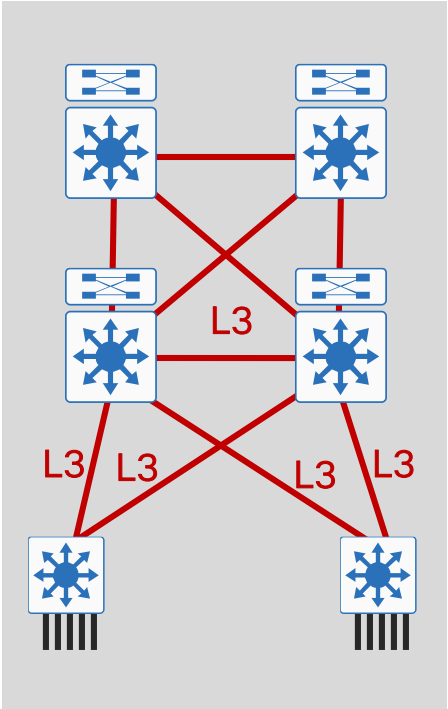

Campus wired LAN design

Option 2: Layer 3 routed access (BRKCRS-3036)

Logical topology—
L3: everywhere
L2: edge only

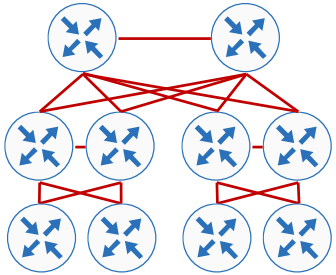

Physical topology:
2 core
2 dist./acc.

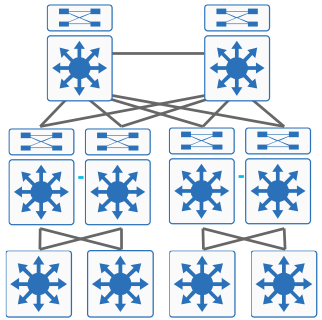

- Complexity reduced for Layer 2 (STP, trunks, etc.)
- Elimination of FHRP and associated timer tuning
- Requires more Layer 3 subnet planning; might not support Layer 2 adjacency requirements

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	

Agenda

- What is high availability?
- Campus network foundations and structured design
- **Campus wired LAN design and high availability**
  - Connecting the devices
  - Considerations with the traditional multilayer campus design
  - Layer-3 access design
  - **Layer-2 and simplified distribution design**
  - New requirements driving new options for campus design
- Campus wireless LAN design and high availability
- Summary and conclusions

Traditional multilayer campus design

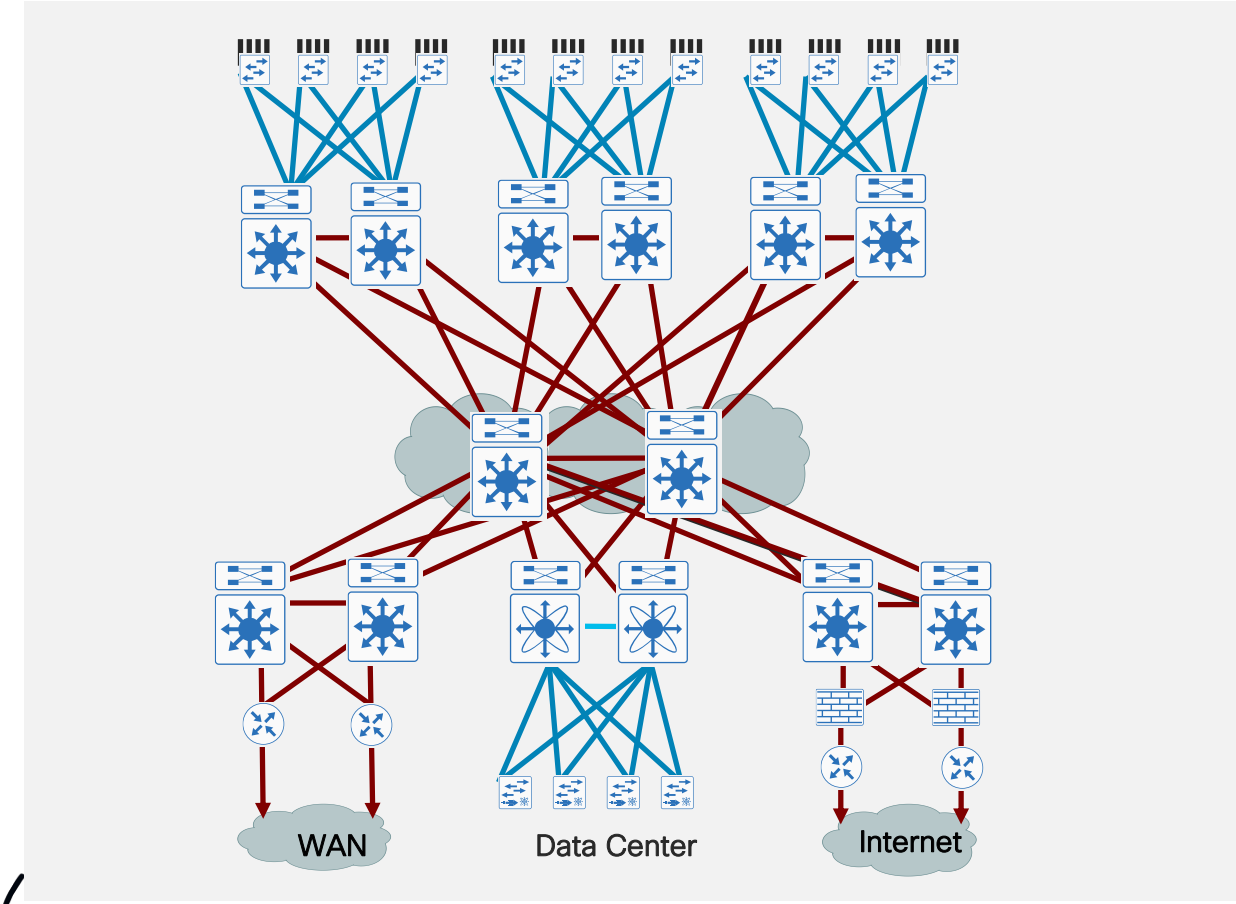

What if we could do a simplified design?

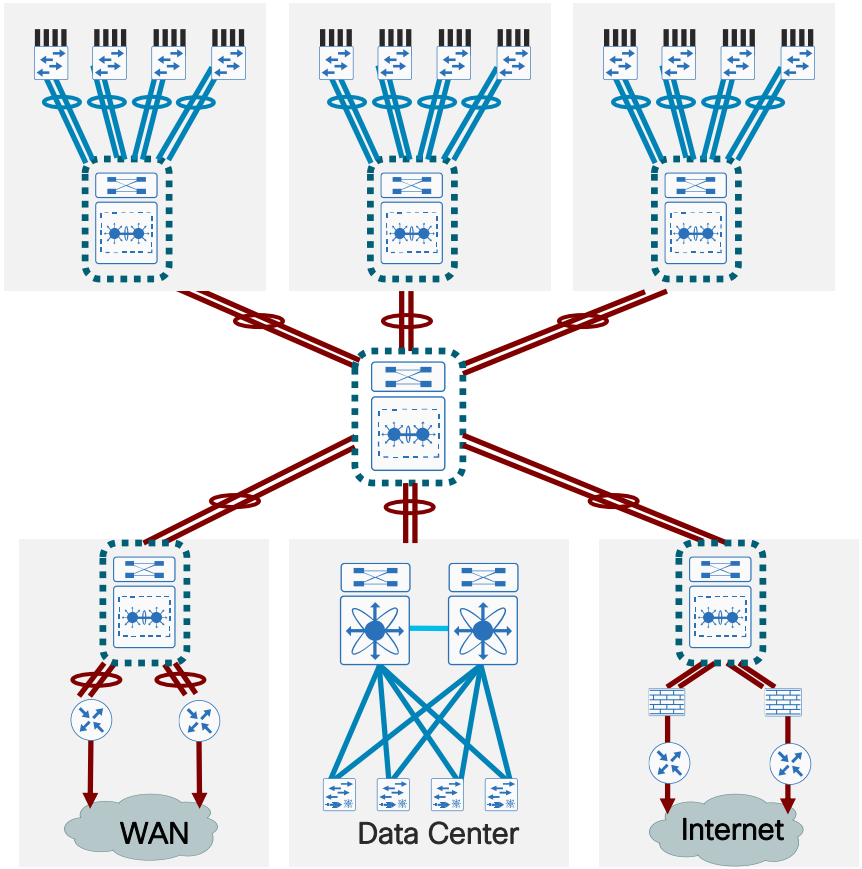

Standalone (multilayer) versus simplified

~~STP Loop~~

~~FHRP~~

~~FHRP Tunings~~

~~PIM DR Priority~~

~~PIM Tunings~~

~~Protocol Dependent Scale~~

~~Unicast Flooding~~

~~Asymmetric Forwarding~~

L2 Hardening

~~Network/System Redundancy Tradeoff~~

~~Protocol Dependent Recovery~~

~~CAM/ARP Tunings~~

~~OSPF LSA/SPF Tuning~~

~~Control/Management/Forwarding Complexity~~

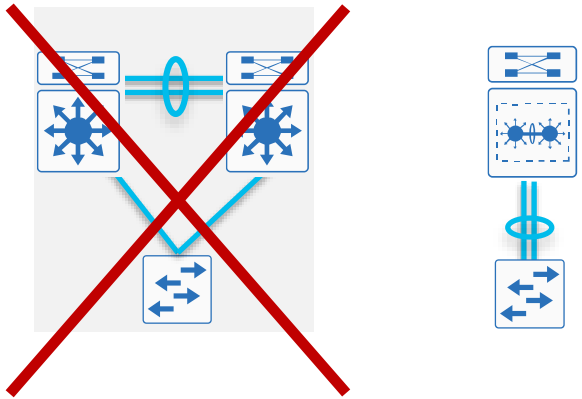

Scale-independent Recovery

Network/System Level Redundancy

Hardware Driven Recovery

Increase Unicast Capacity

Increase Multicast Capacity

Simplified Network Topologies

Control-plane Simplicity

Operational Simplicity

L2-L4 Load Sharing

Flat L2 Network

Unified system architecture

StackWise Virtual (SWV) and Virtual Switching System (VSS)

Simplified Control-Plane

- Single control-plane to manage two physical systems
- Consistent IOS software feature parity as Standalone
- Centralized programming for distributed forwarding

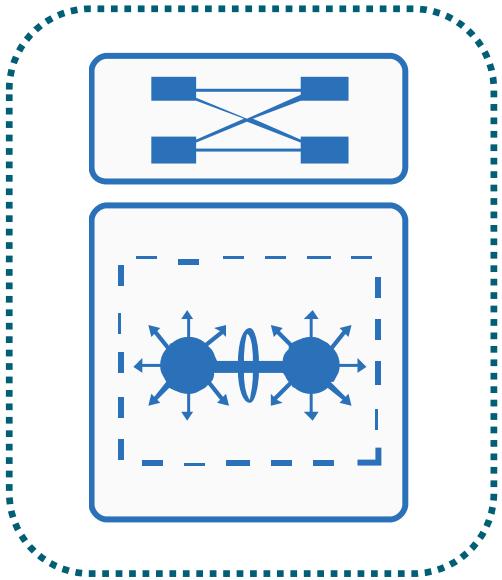

Common Management

- Single virtual system for OOB/in-band management of two physical systems
- Common SNMP MIBs, traps with advanced MIBS
- Single troubleshooting point

“How can I simplify my distribution?”

Cisco StackWise Virtual

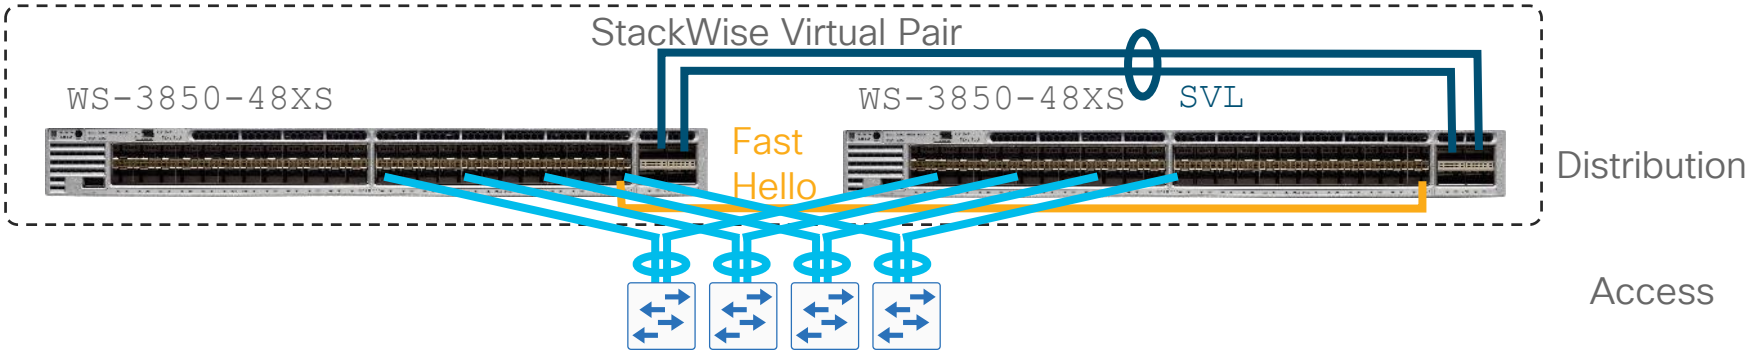

StackWise Virtual – single-homed connections

- Regardless of system modes (SWV, VSS, or standalone), single-homed connections are not recommended
- Cannot leverage distributed architecture benefits.
- Non-congruent Layer 2 or Layer 3 network design with –
  - Centralized network control-plane processing over VSL
  - Asymmetric forwarding plane. Ingress data may traverse over VSL interface and oversubscribe the ports
- Single-point of failure in various faults – Link/SFP/module failure, SSO switchover, ISSU etc..
- Cannot be trusted switch for dual active detection purposes

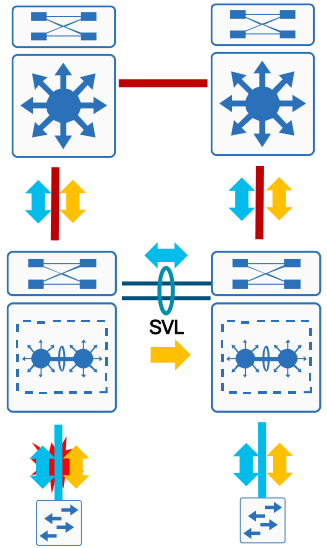

StackWise Virtual– multi-homed physical connections

- Redundant network paths per system delivers best architectural approach
  However, without MultiChassis Etherchannel on Access Layer uplinks
- Parallel Layer 2 paths between bridges builds sub-optimal topology :
  - Creates STP loop. Except for root port, all other ports are in blocking mode
  - Slow network convergence
- Parallel Layer 3 doubles control-plane processing load :
  - ACTIVE switch needs to handle control plane load of local and remote-chassis interfaces
  - Multiple unicast and multicast neighbor adjacencies
  - Redundant routing and forwarding topologies

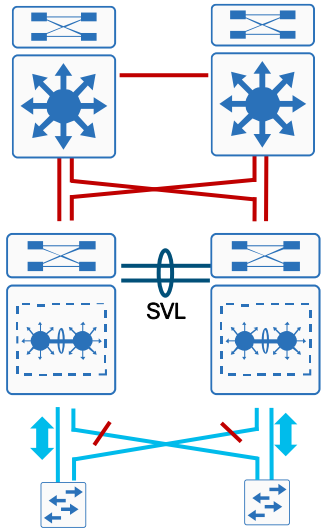

StackWise Virtual– Multichassis EtherChannel

Multichassis EtherChannel (MEC) enables
Distributed link bundling into single logical L2/L3 Interface

- MEC enables:
  - Simplified STP loop-free network topology
  - Consistent L3 control-plane and network design as traditional standalone system
  - Deterministic sub-second network recovery
- MECs can be deployed in two modes
  – Layer 2 or Layer 3

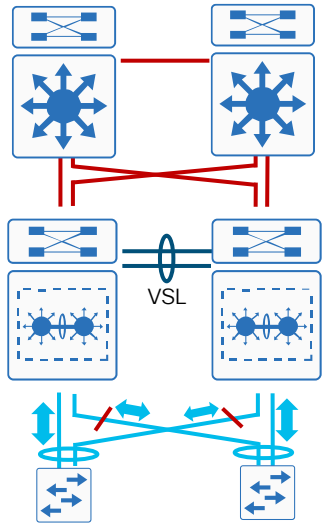

StackWise Virtual – simplified STP topology

- StackWise Virtual simplifies STP – it does not eliminate STP. Never disable STP.
- Multiple parallel Layer 2 network path builds STP loop network
- StackWise Virtual with MEC builds single loop-free network to utilize all available links.
- Distributed EtherChannel minimizes STP complexities compared to standalone distribution design
- STP toolkit should be deployed to safe-guard multilayer network

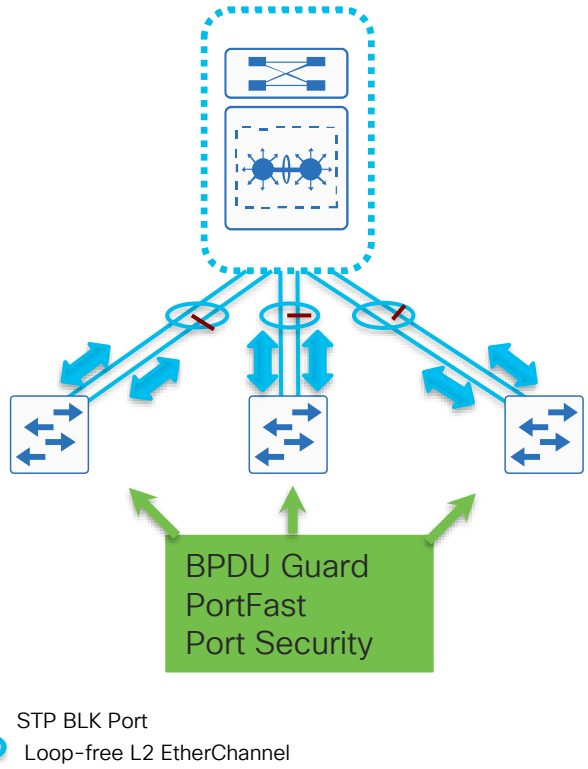

Traditional distribution design comparison

Redundant design with sub-optimal topology and complex operation

- Stabilize network topology with several L2 features:
  - STP Primary and Backup Root Bridge
  - Rootguard
  - Loopguard or Bridge Assurance
  - STP Edge Protection
- Protocol restricted forwarding topology
  - STP FWD/ALT/BLK Port
  - Single Active FHRP Gateway
  - Asymmetric forwarding
  - Unicast Flood
- Protocol dependent driven network recovery:
  - PVST/RPVST+ and FHRP Tuning

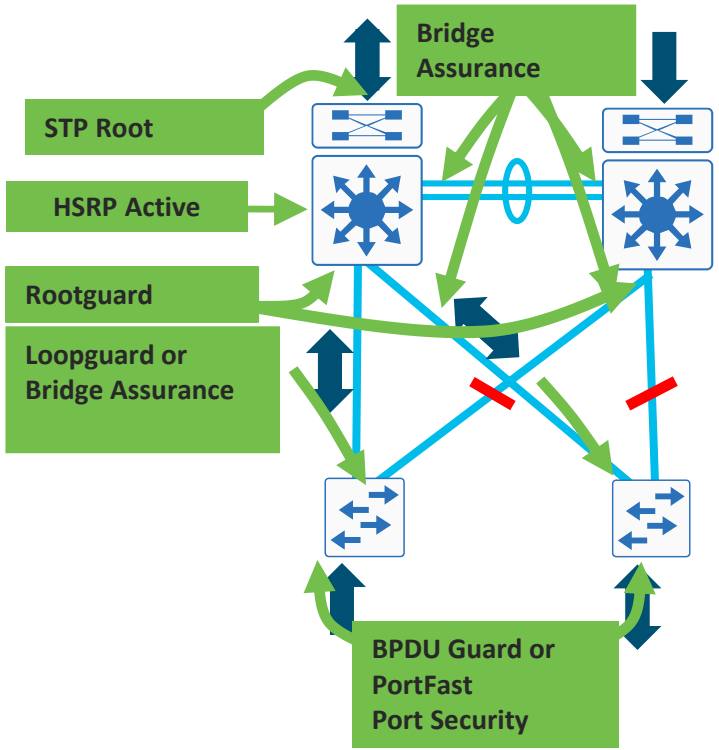

Resiliency versus performance/scale: HSRP

- Multichassis EtherChannel based forwarding topologies
  - Per-Flow Load Balancing based on Layer 2 to Layer 4 + VLANs

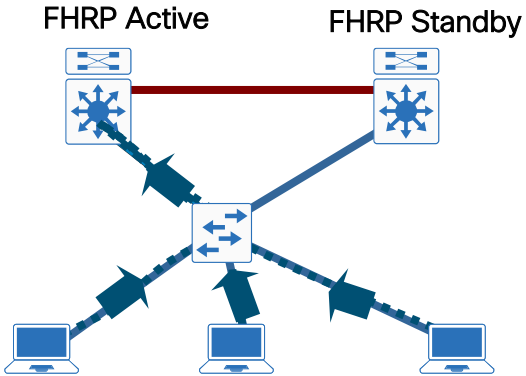

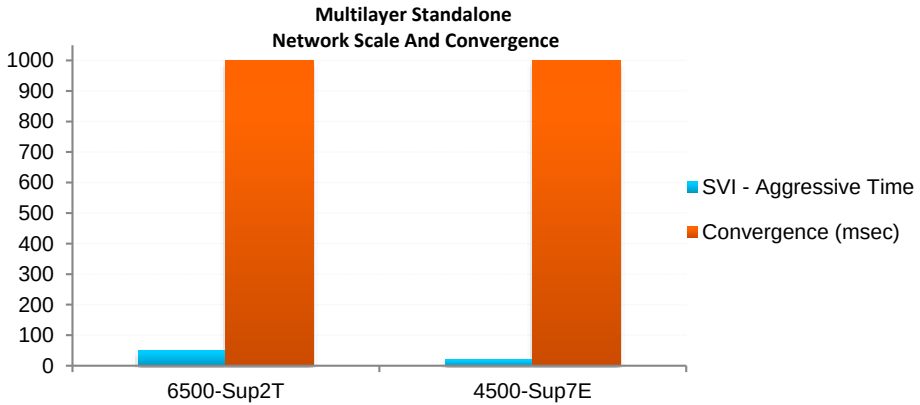

HSRP Config

```
interface Vlan2
 ip address 10.120.2.2 255.255.255.0
 standby 1 ip 10.120.2.1
 standby 1 timers msec 250 msec 750
 standby 1 priority 150
 standby 1 preempt
 standby 1 preempt delay minimum 180
```

Resiliency versus performance/scale: SW Virtual

- Multichassis EtherChannel based forwarding topologies
  - Per-Flow Load Balancing based on Layer 2 to Layer 4 + VLANs
- Hardware-Based Fault Detection and Recovery
  - Deterministic network convergence with simplistic approach
- Increases Network Scale for system reliability
- No reliability compromise to enable path and system-level Quad-Sup redundancy

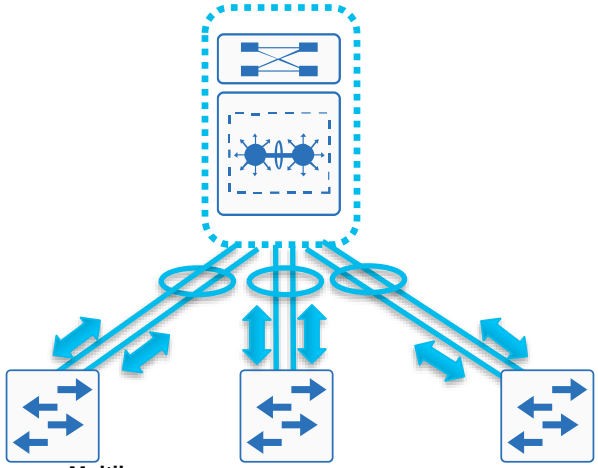

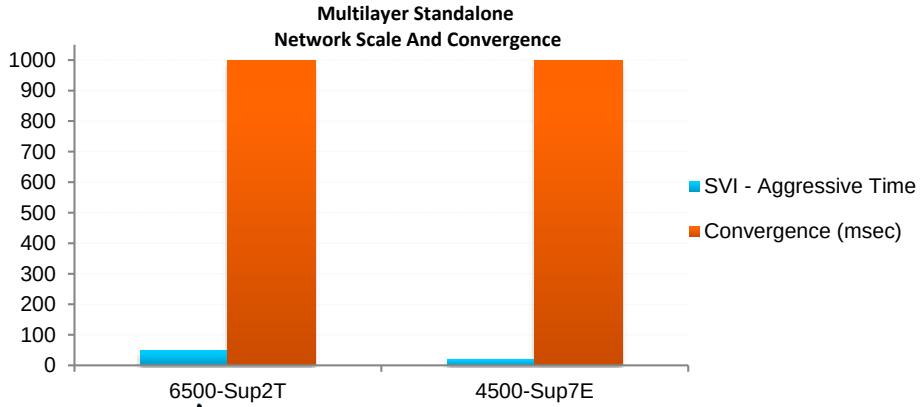

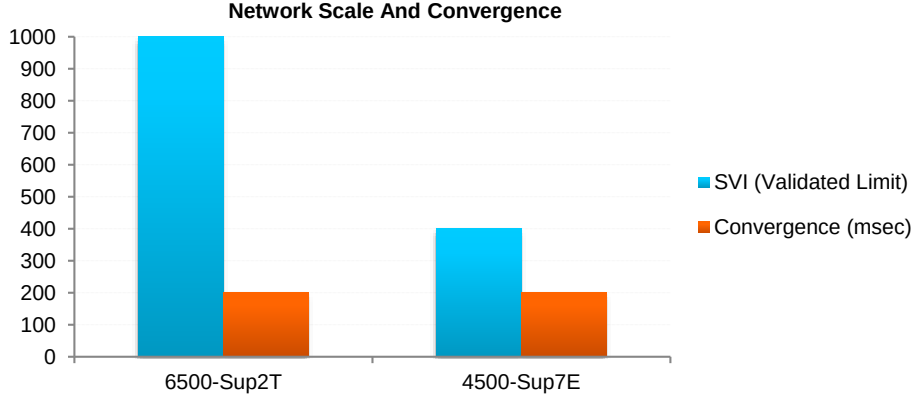

PIM timers also need to be tuned

- Multicast recovery depends on PIM DR failure detection in Layer 2 network
- PIM routers exchanges PIM expiration time in query message
  - DR Failure Detection:
    ~90 seconds (30 sec. hello * 3 multiplier)
- Tune PIM query interval to sub-sec as FHRP for faster multicast convergence
- Sub-second protocol timer must be avoided on SSO capable networks

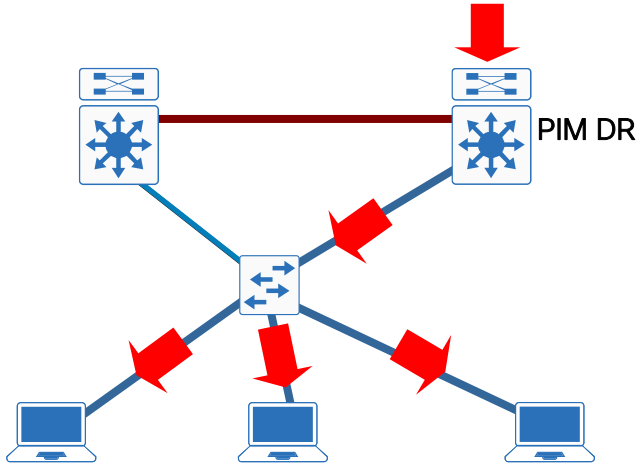


```
interface Vlan2
 ip pim sparse-mode
 ip pim query-interval 250 msec
```

Simplified, robust multicast network: SW Virtual

- Single PIM DR system in Layer 2 network to process IGMP from host receivers
- Doubles multicast forwarding performance across all Multichassis EtherChannel member links
- Optimize multicast network with PIM stub configuration
- Rapid, deterministic and simple multicast design
  - Hardware based sub-second fault detection and recovery.
  - Eliminates aggressive timer requirement and improves system performance and scalability

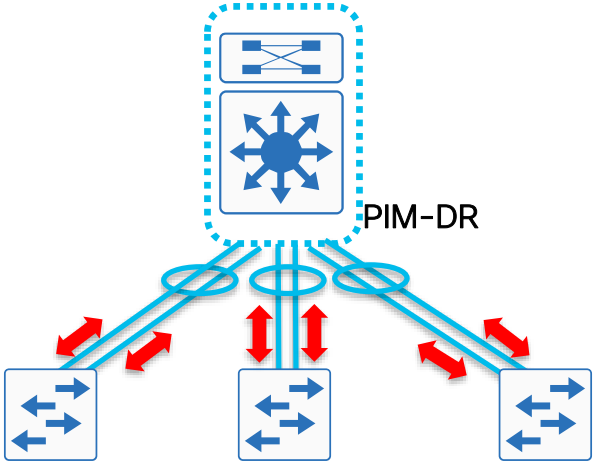


```
interface Vlan2
 ip pim passive
```

Multichassis EtherChannel performs better in any network design

- Network recovery mechanic varies in different distribution design –
  - Standalone – protocol and timer dependent
  - StackWise Virtual – hardware dependent
- StackWise Virtual logical distribution system –
  - Single P2P STP Topology
  - Single Layer 3 gateway
  - Single PIM DR system
- Distributed and synchronized forwarding table –MAC address, ARP cache, IGMP
- All links are fully utilized based on Ether-channel load balancing

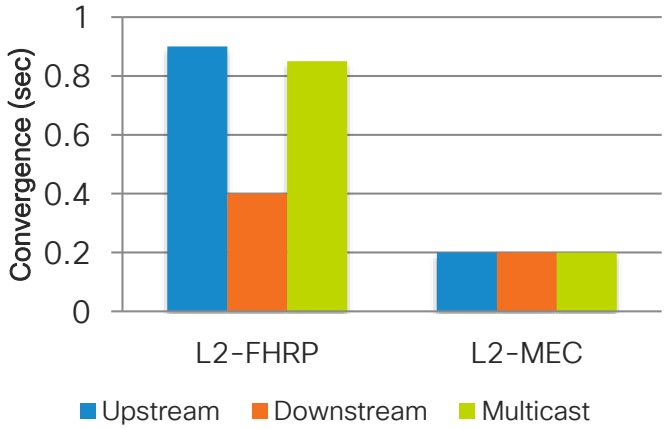

StackWise Virtual-enabled campus core design

- Extend StackWise Virtual architectural benefits to campus core layer network
- SWV-enabled core increases capacity, optimizes network topologies and simplifies system operations
- Key SWV-enabled core best practices :
  - Protect network availability and capacity with NSF/SSO
  - Simplify network topology and routing database with single MEC
  - Leverage self-engineer SWV and MEC capabilities for deterministic network fault detection and recovery

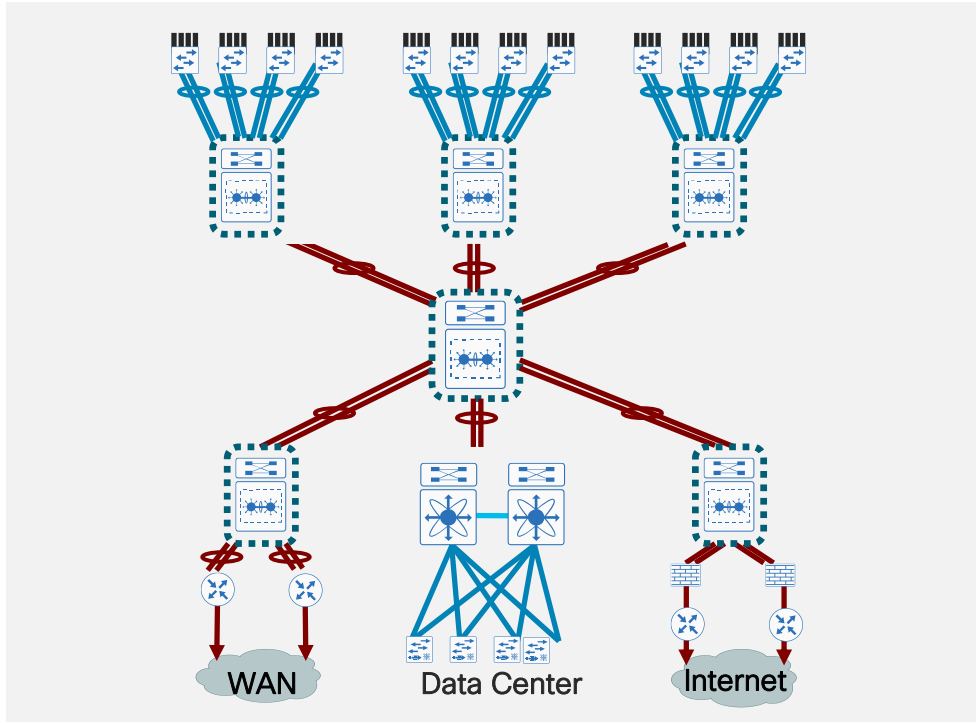

StackWise Virtual core network design options

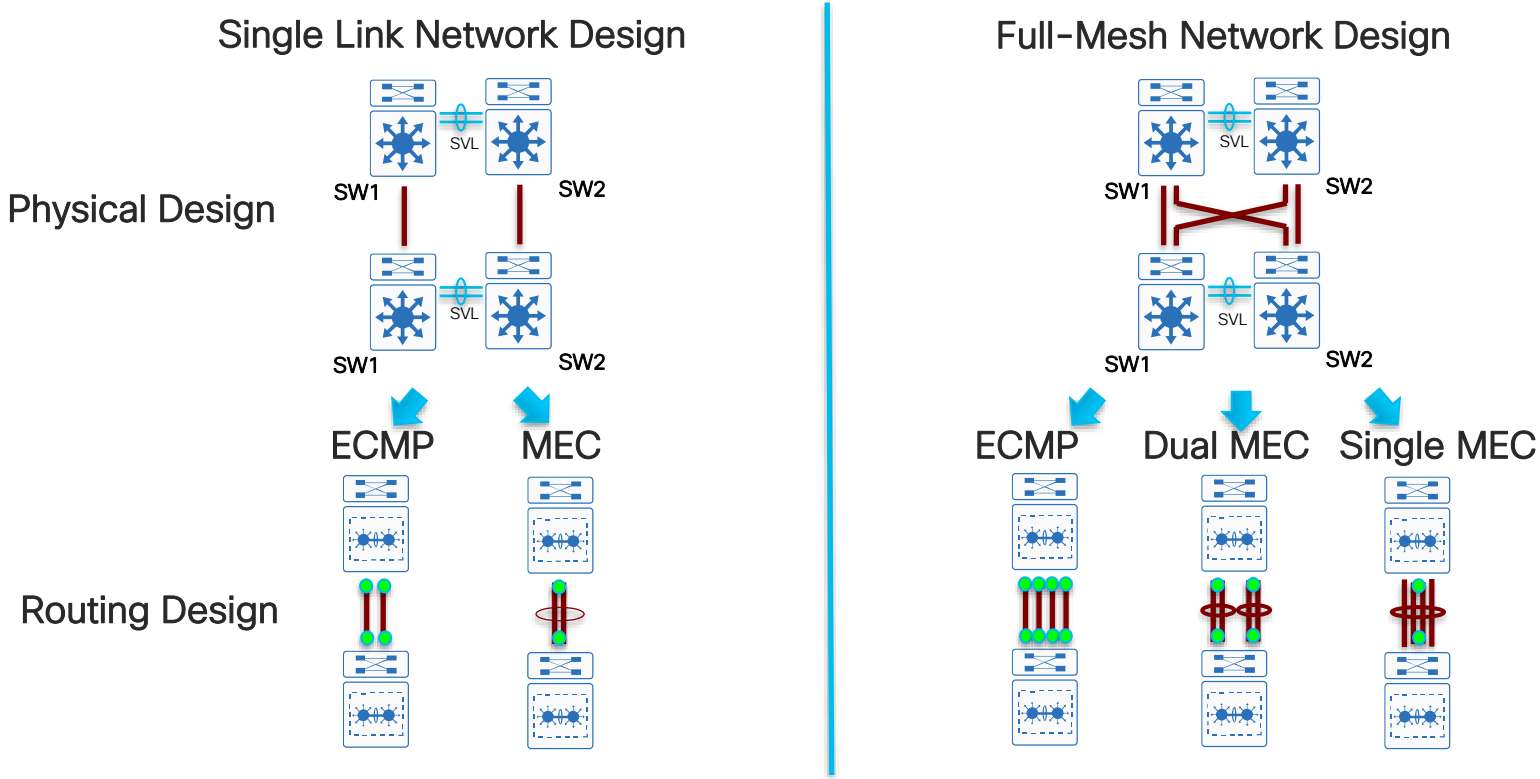

Recommended Design : Full-mesh physical network with single MEC

Summary – optimizing core performance (1/2)

HW Driven Forwarding Topology & High Availability

MEC Design

SWV-Core

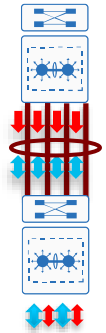

SWV-Dist

- Improved System Performance – Single MEC that reduces 50% control-plane load in Core
- Simple Topology – Abstracts hardware layer with single neighbor and single best forwarding path
- Improved Network Performance – Consistent unicast forwarding design. Increase in multicast capacity in core
- Improved App Performance – Increased unicast and multicast load sharing input variables
- Resilient: Protocol + scale-independent network recovery

ECMP Design

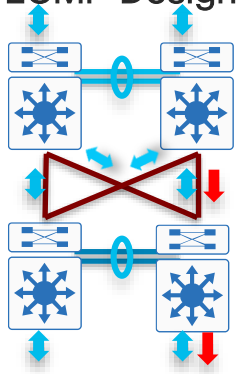

Unicast Forwarding Path
Multicast Forwarding Path

Standalone-Core

Standalone--Dist

- ECMP network doubles control-plane load and redundant topologies
- Unicast routing protocol installs ECMP. Multicast routing installs single Outgoing Interface List (OIL)
- Egress data forwarding decision is localized with 6500E/6800. Catalyst 4500E/4500X egress forwarding decision is across all ECMP links
- Protocol and scale-dependent network recovery

cisco Live!

Summary – optimizing core performance (2/2)

HW Driven Forwarding Topology & High Availability

Standalone-Core

VSS-Dist

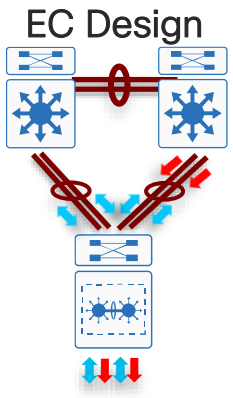

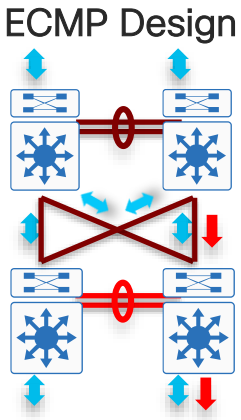

Unicast Forwarding Path
Multicast Forwarding Path

Standalone-Core

Standalone-Dist

- Dual MEC between network layer maintains original control-plane load on SWV ACTIVE system
- Dual MEC L3 unicast/multicast neighbor and ECMP best path in table
- Consistent unicast forwarding design. Increase in multicast switching capacity in core
- Increased unicast and multicast load sharing input variables
- Protocol and scale-independent network recovery

- ECMP network design doubles control-plane load and redundant topologies
- Unicast routing protocol installs ECMP best path between two chassis. Multicast routing installs single OIL
- Egress data forwarding decision is localized with 6500E. Catalyst 4500E egress forwarding decision is across all ECMP links
- Protocol and scale-dependent network recovery

cisco Live!

Simple core network design delivers deterministic network recovery

- Routing protocol independent network convergence in large scale campus core
- ECMP prefix-independent convergence (PIC) improves performance
- Cisco Express Forwarding (CEF) optimization in IOS software.
- Default behavior: no additional configuration or tuning required
- Hardware-based fault detection and recovery in MEC/EC designs

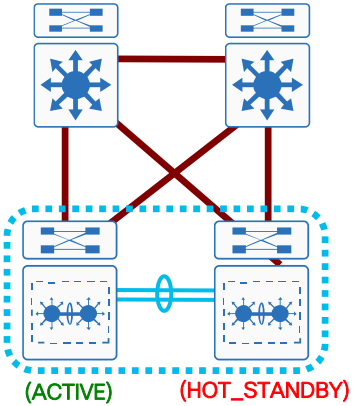

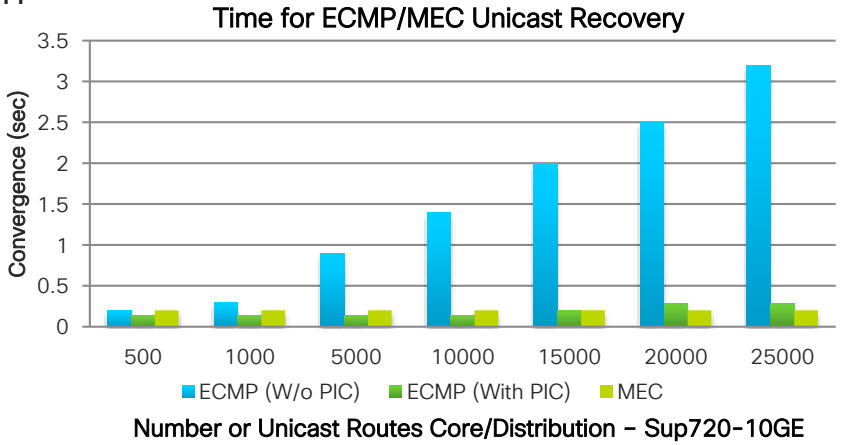

SWV core simplifies multicast operation, improves performance and redundancy (before)

- Standalone core needs anycast MSDP peering for RP redundancy
- ECMP builds single multicast forwarding path and protocol-based fault detection and recovery

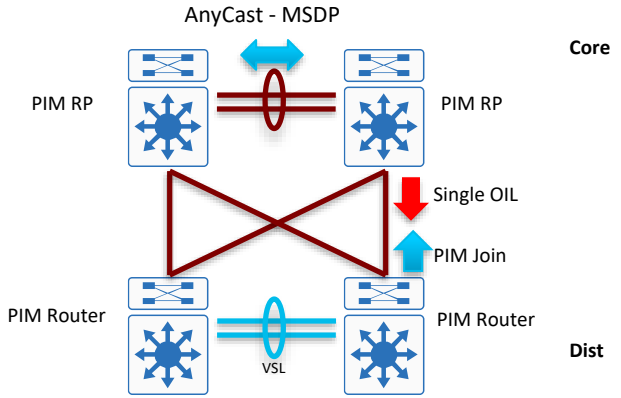

SWV core simplifies multicast operation, improves performance and redundancy (after)

- SWV based Catalyst systems enables PIM RP Redundancy with resilient technologies
- MEC increases multicast forwarding capacity by utilizing all member-links and provides hardware-based fault detection and recovery

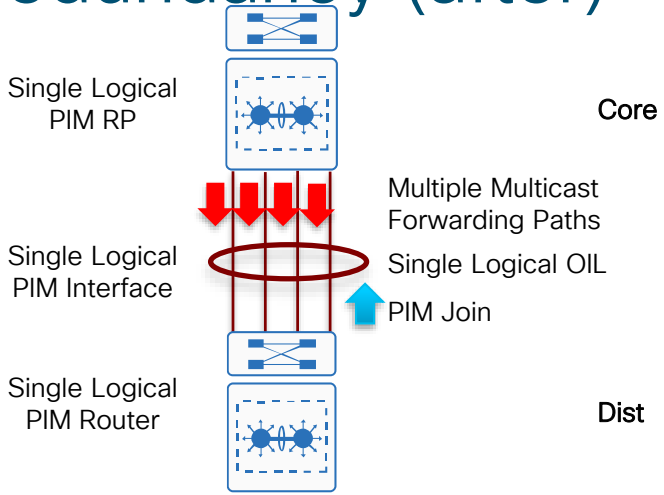


```
SWV#show ip multicast redundancy state
Multicast IPv4 Redundancy Mode:    SSO
<snip>
Stale NSF state flush timeout: 30000 ms
Current sync state: Synched

Multicast ISSU Client Status:
  PIM MIC client                ISSU compatible
  MRIB MIC client               ISSU compatible
  MFIB IPv4 MIC client          ISSU compatible
```

Simplified multicast network design delivers deterministic network recovery

- ECMP multicast recovery is mroute scale dependent could range in seconds.
- MEC/EC multicast recovery is hardware-based and recovery is scale-independent in sub-seconds

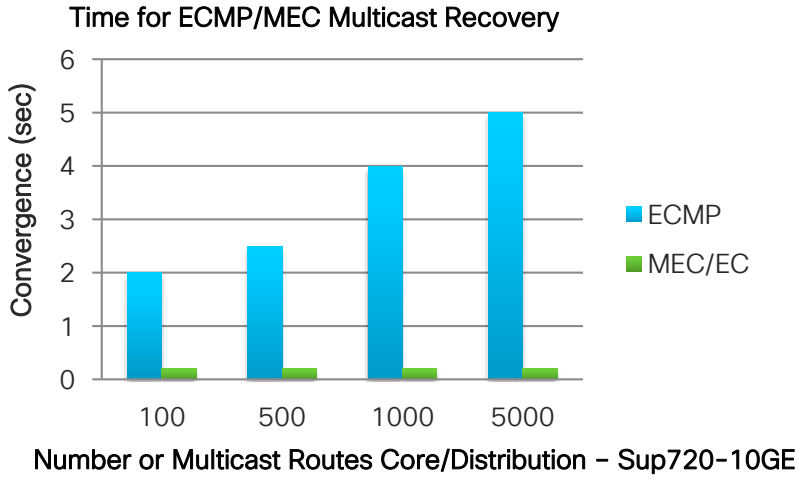

Implementing non-stop forwarding

- SWV software design is built on NSF/SSO architecture.
- Switches deployed in SWV mode must enable NSF. No configuration required on NSF helper system
- NSF capability must be manually enabled for all Layer 3 routing protocols :
  - EIGRP, OSPF, ISIS, BGP, MPLS etc..
- In VRF environment the NSF must be manually enabled on per-VRF IGP instance
- Multicast NSF capability is default ON

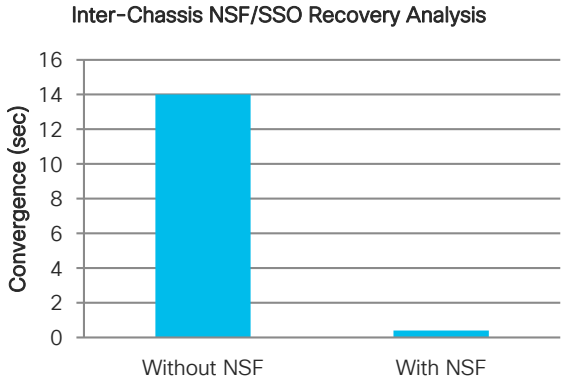

Sub-second protocol timers and NSF/SSO

- NSF is intended to provide availability through route convergence avoidance
- Fast IGP timers are intended to provide availability through fast route convergence
- In an NSF environment dead timer must be greater than:
  SSO recovery +
  Routing Protocol restart +
  time to send first hello
- Recommendation:
  Do not configure aggressive timer Layer 2 protocols, i.e. Fast UDLD
  Do not configure aggressive timer Layer 3 protocols,
  i.e. OSPF Fast Hello, BFD etc. Keep all protocol timers at default settings

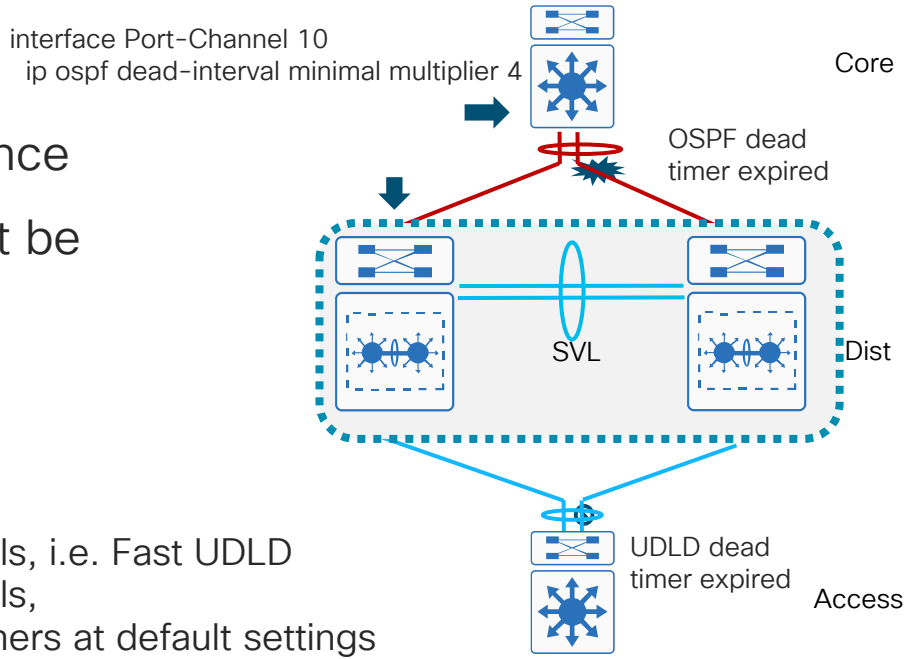

Campus wired LAN design

Option 3: Layer 2 access with “simplified” distribution (BRKCRS-1500)

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

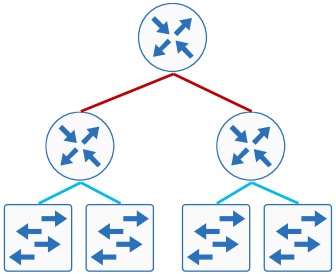

Physical topology:
2 core
2 dist./acc.

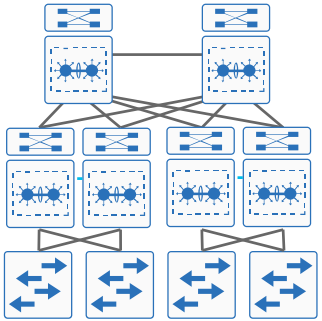

- Leading campus design for easy configuration and operation when using stacking or similar technology (StackWise Virtual, VSS)
- Flexibility to support Layer 2 services within distribution blocks, without FHRPs.
- Easy to scale and manage

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	

Agenda

- What is high availability?
- Campus network foundations and structured design
- **Campus wired LAN design and high availability**
  - Connecting the devices
  - Considerations with the traditional multilayer campus design
  - Layer-3 access design
  - Layer-2 and simplified distribution design
  - **New requirements driving new options for campus design**
- Campus wireless LAN design and high availability
- Summary and conclusions

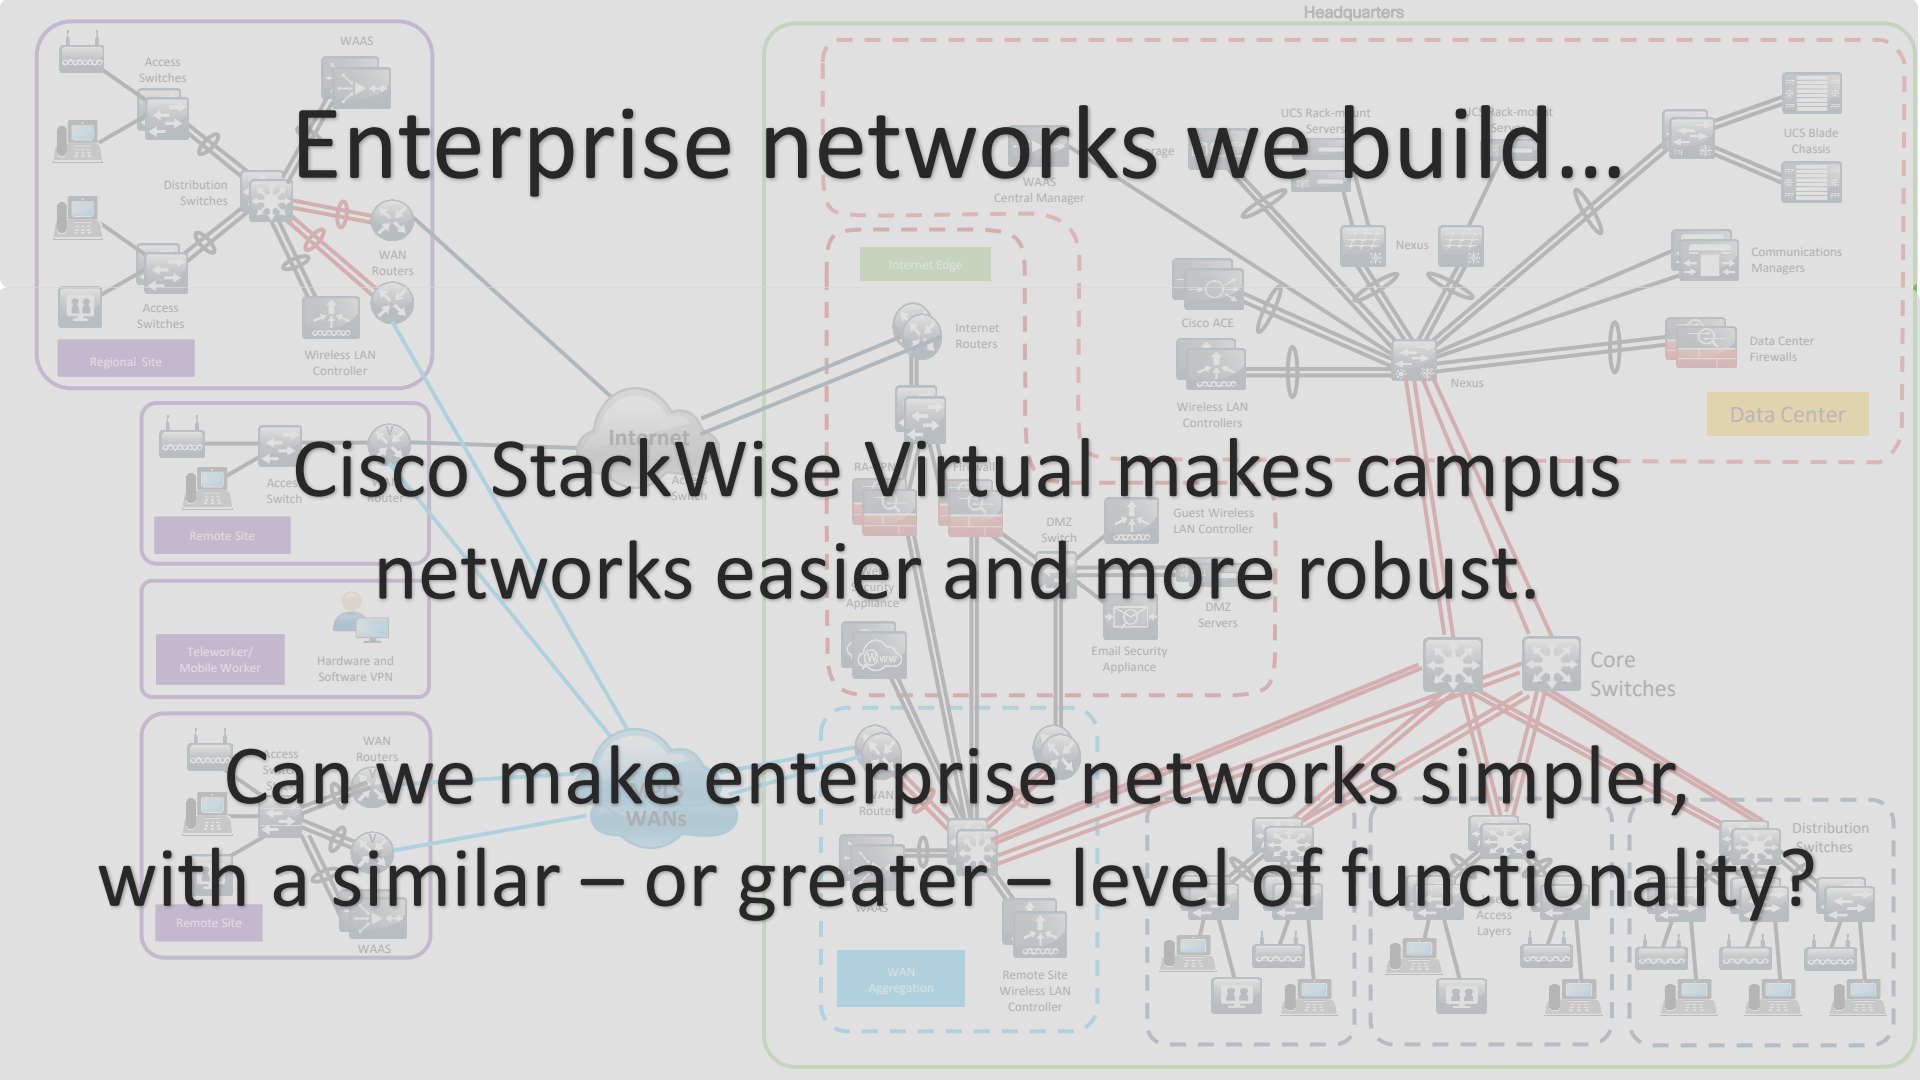

Enterprise networks we build…

Cisco StackWise Virtual makes campus networks easier and more robust.

Can we make enterprise networks simpler, with a similar – or greater – level of functionality?

What's different in your network today versus a decade ago? How does it affect availability?

Bring Your Own Device in the workspace

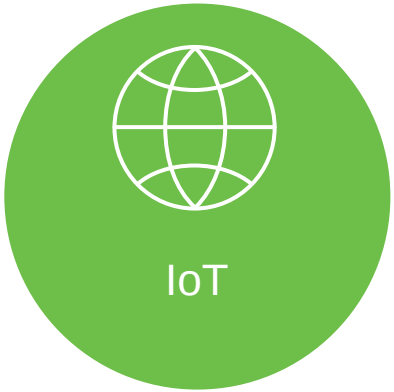

Auto-detect non-user devices devices everywhere

Networking and security advanced threats

Key challenges for traditional networks

Difficult to segment

Ever increasing number of users and endpoint types

Ever increasing number of VLANs and IP Subnets

Complex to manage

Multiple steps, user credentials, complex interactions

Multiple touch-points

Slower issue resolution

Separate user policies for wired and wireless networks

Unable to find users when troubleshooting

Traditional networks cannot keep up!

What if you could do this?

Cisco Software-Defined Access

- Enables:
  - Host mobility
  - Network segmentation
  - Role-based access control
- It is an overlay network to the network underlay
  - Control plane based on LISP
  - Data plane based on VXLAN
  - Policy plane based on TrustSec

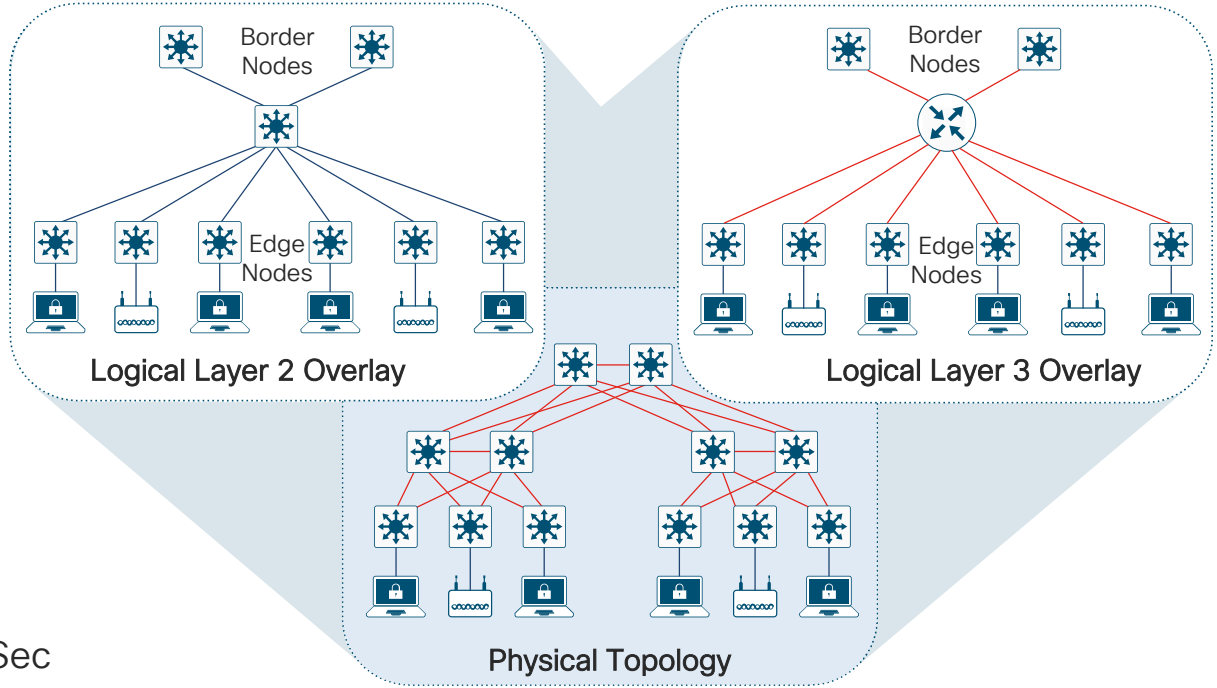

Software-Defined Access Soluton Design Guide

https://cs.co/sda-sdg

SD-Access

Why overlays?

Separate the “forwarding plane” from the “services plane”

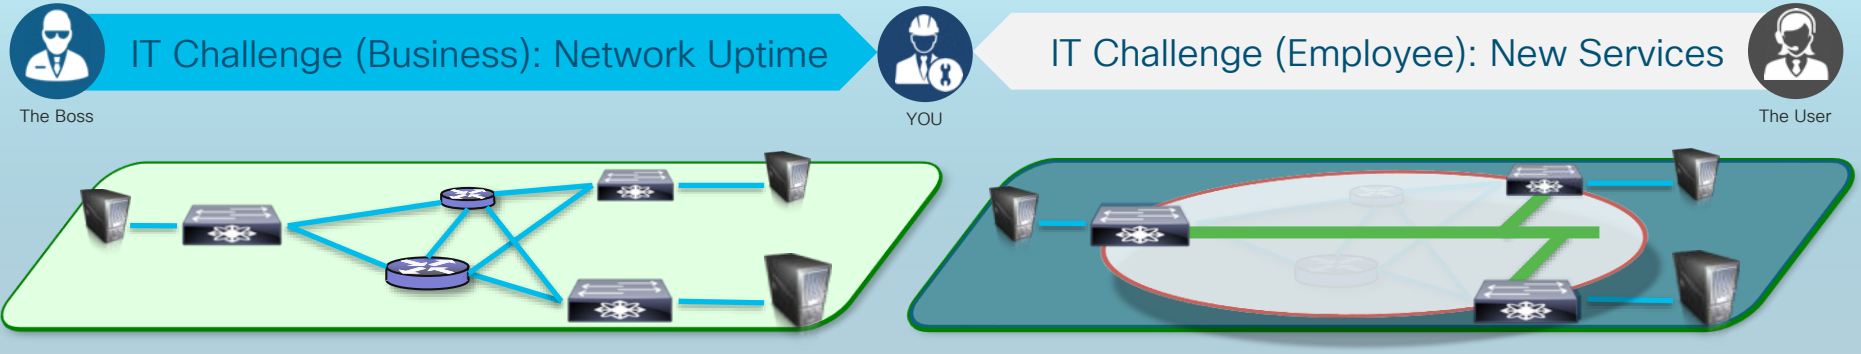

Simple Transport Forwarding

- Redundant Devices and Paths
- Keep It Simple and Manageable
- Optimize Packet Handling
- Maximize Network Reliability (HA)

Flexible Virtual Services

- Mobility – Map Endpoints to Edges
- Services – Deliver using Overlay
- Scalability – Reduce Protocol State
- Flexible and Programmable

SD-Access
Types of overlays

Hybrid L2 + L3 Overlays offer the Best of Both Worlds

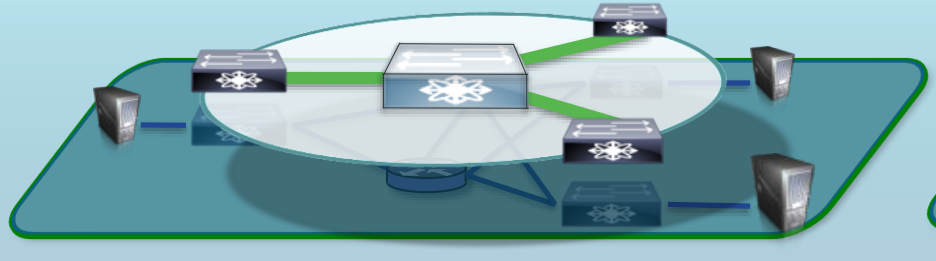

Layer 2 Overlays

- Emulates a LAN segment
- Transport Ethernet Frames (IP & Non-IP)
- Single subnet mobility (L2 domain)
- Exposure to Layer 2 flooding
- Useful in emulating physical topologies

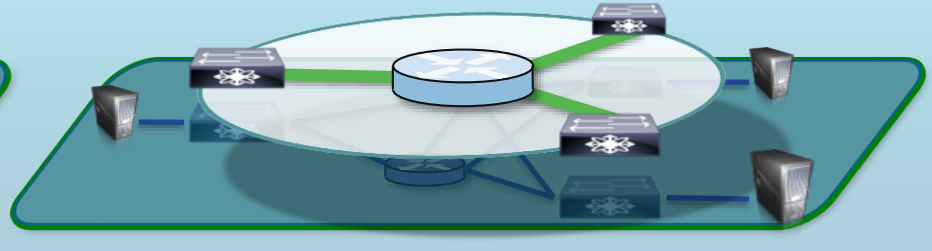

Layer 3 Overlays

- Abstract IP connectivity
- Transport IP Packets (IPv4 & IPv6)
- Full mobility regardless of Gateway
- Contain network related failures (floods)
- Useful to abstract connectivity and policy

Campus wired LAN design

Option 4: Cisco Software-Defined Access (BRKCRS-1501, many others)

Logical topology—
L2/L3: flexible overlays

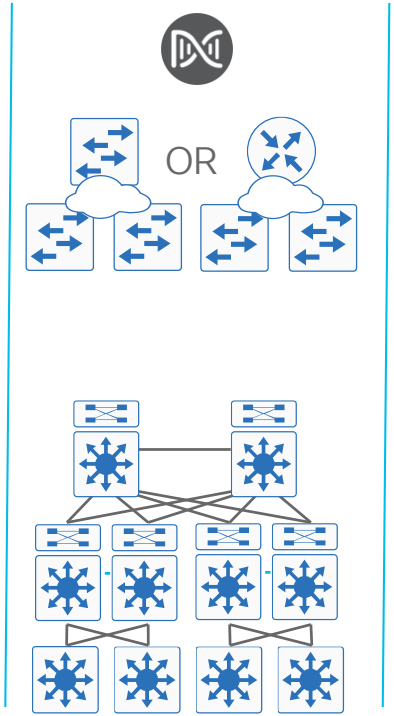

Physical topology:
2 core
2 dist./acc.

- 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	✔

Campus wired LAN design options—summary

	Traditional Multilayer Campus BRKCRS-2031	Layer 3 Routed Access BRKCRS-3036	L2 Access / Simplified Distribution BRKCRS-1500	SD-Access / Fabric for Campus BRKCRS-1501 (and many others)
Logical topology	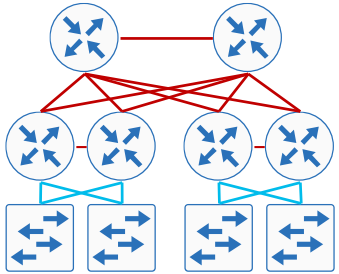	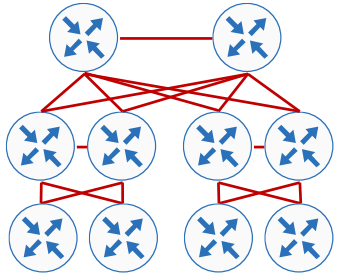	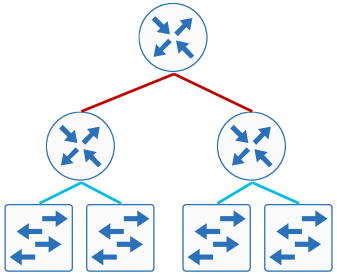	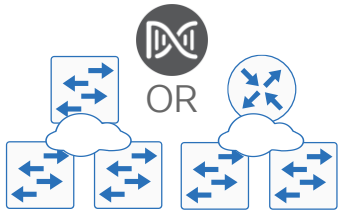
Design notes	Protocols / Tuning	L3 Planning Limited L2	Flexible, Easy, Scalable	Flexible, Tools to Simplify
Physical topology: 2 core 2 dist./acc.	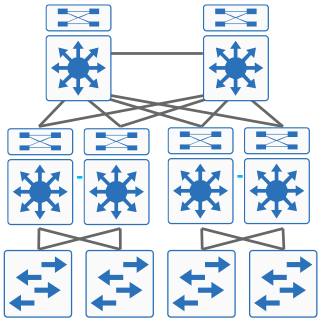	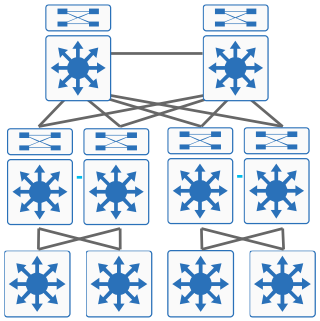	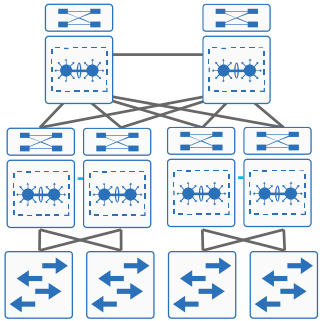	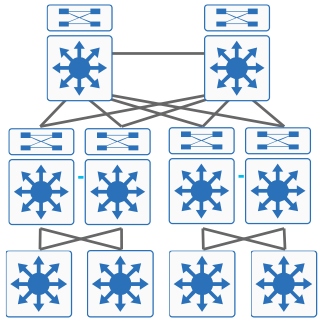

How do I get there?

Successful deployments...

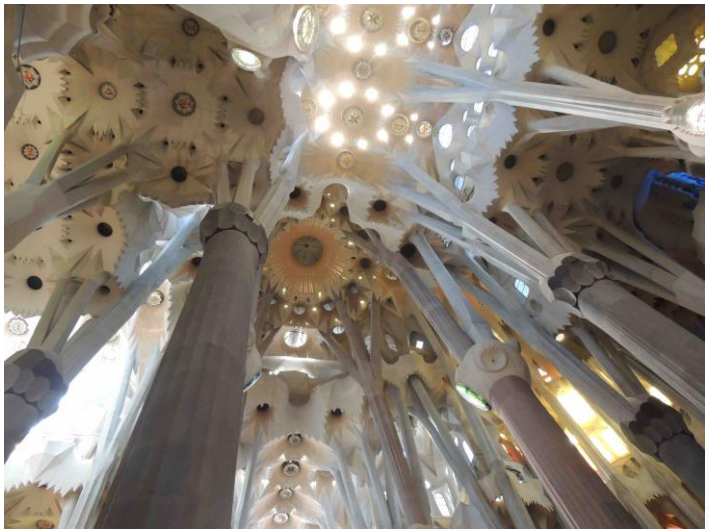

...start with a plan.

Photos showing Basílica i Temple Expiatori de la Sagrada Família

Summary – Design decisions affecting high availability in the wired campus design

- Hierarchy
- Device capabilities (or lack)
- Device interconnections (direct/indirect, media, config)
- Layering / choices for fault detection (HW and SW)
- Layer-2 application needs
- Number/complexit of protocols required for a given design

- Use of ECMP and/or MCEC
- Subsecond timers vs. SSO
- Overall design choices (multilayer vs. routed access vs. simplified distribution vs. SD-Access) and supporting protocols
- Simplifying the network and improving network availability improves other services overlaid on that network

Agenda

- What is high availability?
- Campus network foundations and structured design
- Campus wired LAN design and high availability
- **Campus wireless LAN design and high availability**
- Summary and conclusions

Campus wireless LAN design and high availability

Who connected to a wired network today?

… a typical day of a connected life…

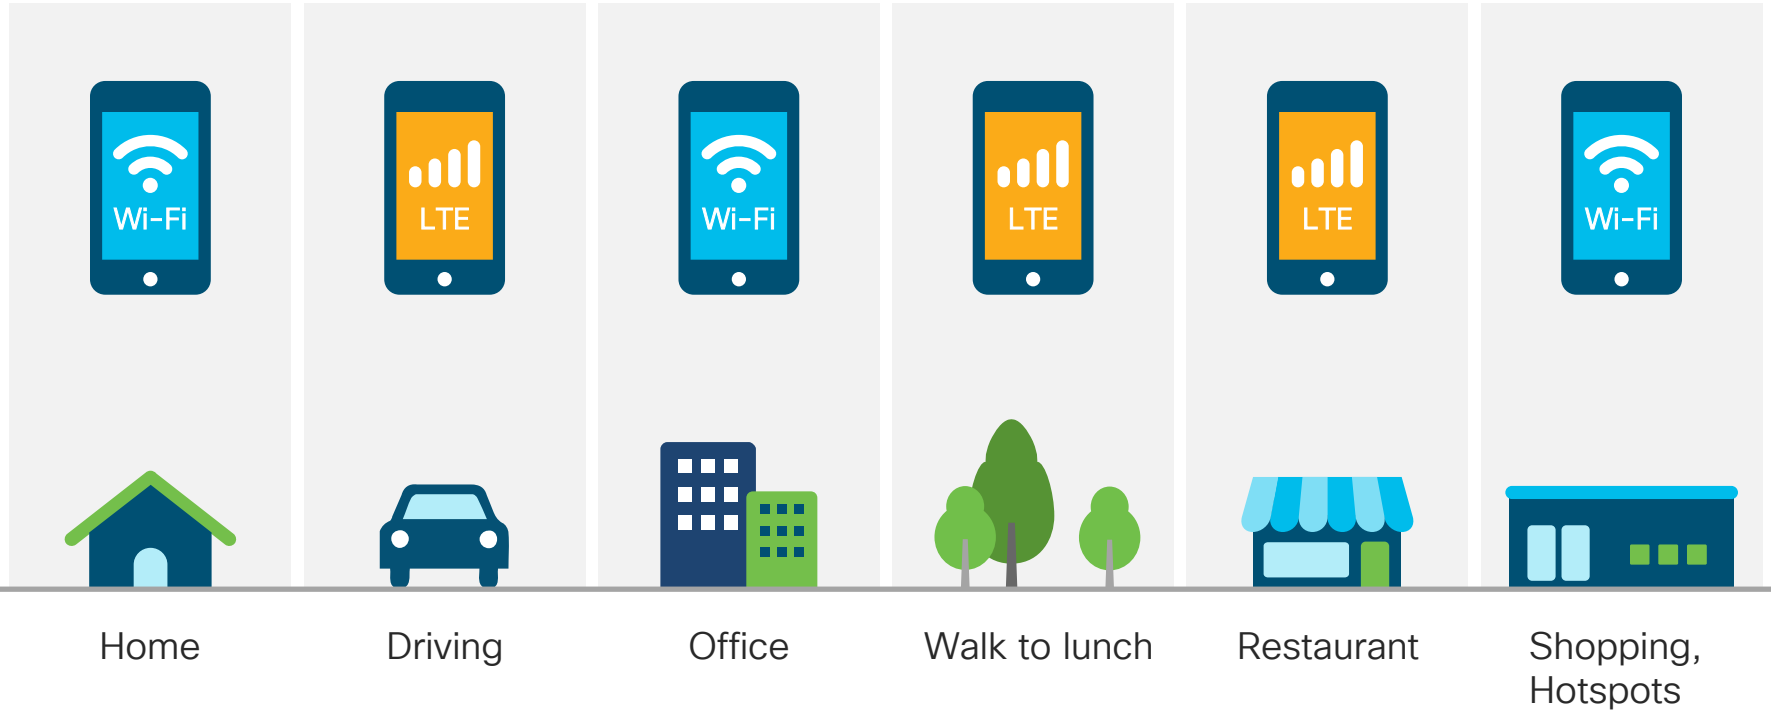


“No Wireless == No Network Access”

Section Objective

The goal of this section is to show you how to design and deploy a Highly Available wireless network **to reduce the network downtime**

Wireless High Availability concepts

- Good news: all the High Availability concepts and best practices we have seen for wired are applicable to wireless access as well
- Bad news: wireless is not wired

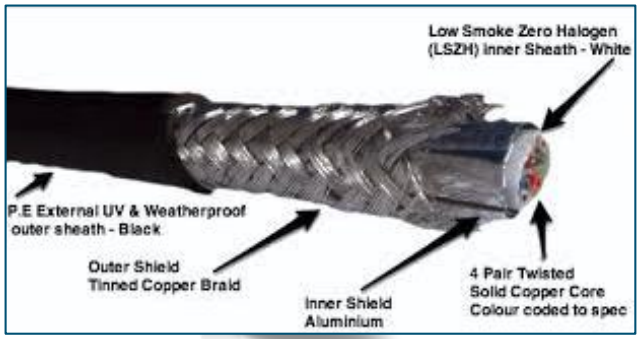

Shielded, isolated access

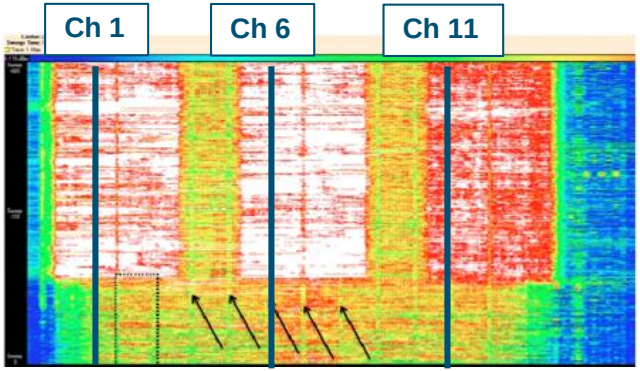

No electromagnetic protection

We use the air to transmit packets, it's a shared media, it's unlicensed….enough?

Agenda

- What to do at the Radio Frequency layer?
- HA Design and Deployment Practices
  - Central/Large Site Deployments
  - Remote/Small Site Deployments
- Wireless Controller Features for Planned Outages
- Key takeaways

RF HA – how to build redundancy at the RF layer?

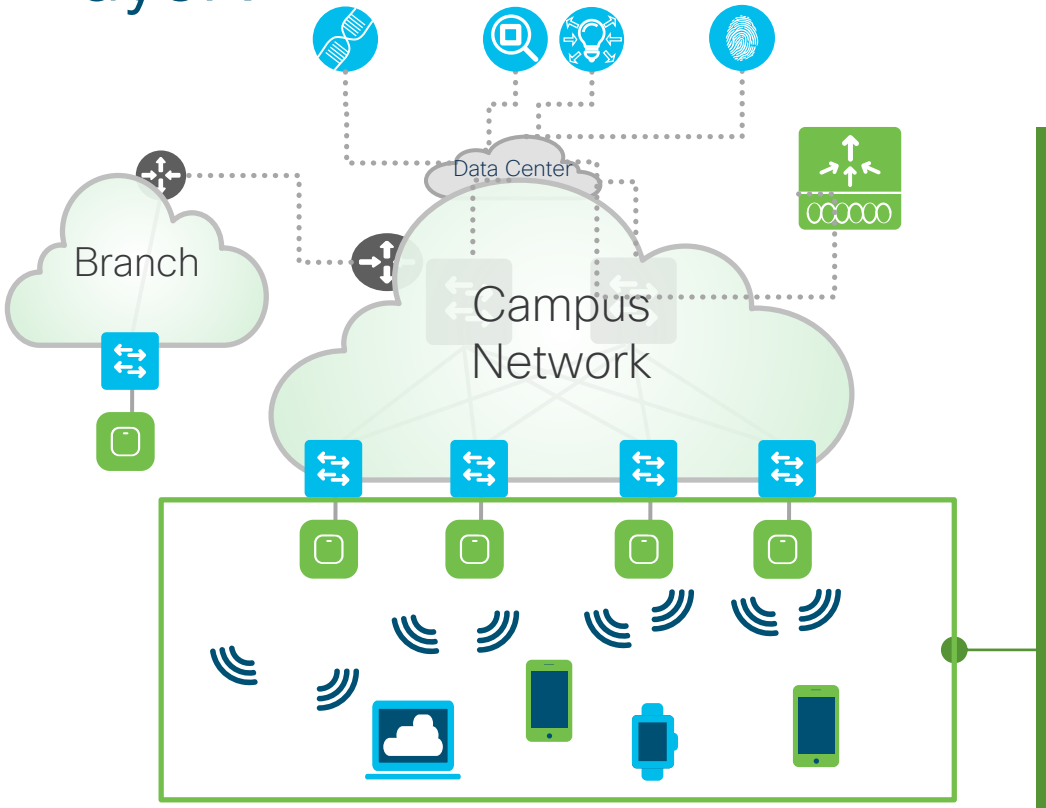

- Creating a stable, predictable RF environment (**Proper Design, Site Survey**)
- Dealing with RF that is continuously changing (**RRM and RF Management**)
- Coping with coverage holes from an AP going down (**RRM and RF Management**)

Radio Frequency (RF) High Availability: Site Survey

- Site Survey, site survey….and site survey
  - Use “Active” survey
  - Coverage vs. Capacity
  - Consider Client type (ex. Smartphone vs. Laptop)

My power is half of my brother MacBook

I try to connect to 5GHz and stay connected until the signal is REALLY bad

Radio Frequency (RF) High Availability: Site Survey

- Site Survey, site survey….and site survey
  - Use “Active” survey
  - Coverage vs. Capacity
  - Consider Client type (ex. Smartphone vs. Laptop)

My antenna gain is 4 times smaller

and then move to another BSSID if it is REALLY better

Radio Frequency (RF) High Availability

- Site Survey, site survey….and site survey
  - Use “Active” survey
  - Coverage vs. Capacity
  - Consider Client type (Smartphone vs. Laptop)
- AP positioning and antenna choice is Key
  - Use common sense
  - Light source analogy
  - Internal antennas are designed to be mounted on ceiling
  - External antennas: use same antennas on all connectors
- Tools
  - What you use is less important than how you use it
  - Use the same tool to compare results

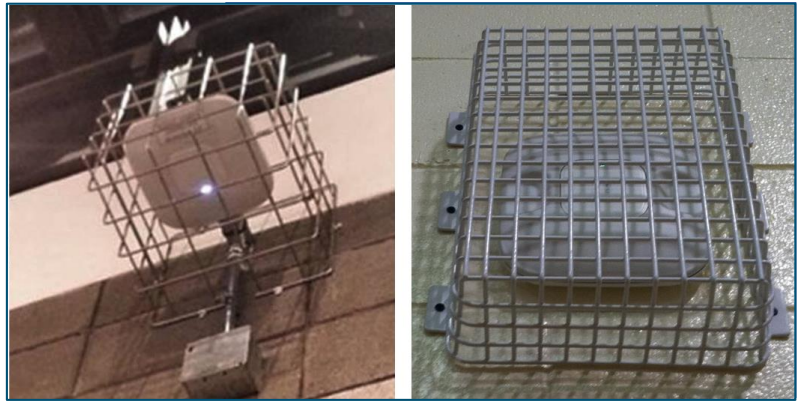

RF High Availability: Cisco RRM

- What are Radio Resource Manager (RRM)'s objectives?
  - Provide a system wide RF view of the network at the Controller (only Cisco!!)
  - Dynamically balance the network and mitigate changes
  - Manage Spectrum Efficiency so as to provide the optimal throughput under changing conditions
- What's RRM
  - DCA–Dynamic Channel Assignment
  - TPC–Transmit Power Control
  - CHDM–Coverage Hole Detection and Mitigation
- RRM best practices
  - RRM settings to auto for most deployments (High Density is a special case)
  - Design for most radios set at mid power level (lever 3 for example)
  - Use RF Profiles to customize RRM settings per Areas/Groups of APs

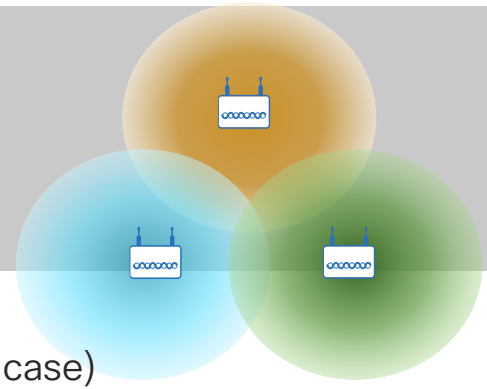

RF High Availability: Cisco RRM

RRM DCA in action

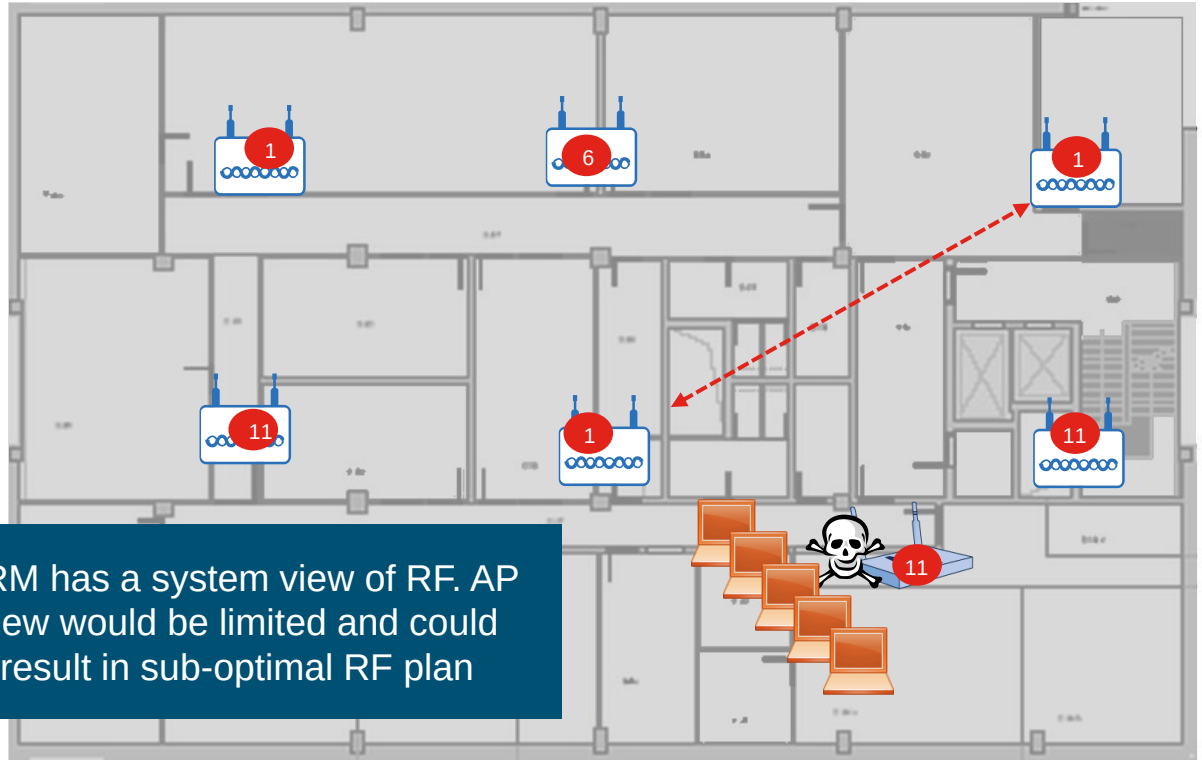

RRM has a system view of RF. AP view would be limited and could result in sub-optimal RF plan

- RRM will determine the optimal channel plan based on AP layout
- A rogue AP is detected on channel 11
- RRM will assess the RF and take a decision in less than 10min
- Channel change is triggered to improve the RF
- Note how the 3 non overlapping channels are still maintained!
- With a limited AP-based view of the RF, each AP will avoid channel 11 reducing overall network capacity

RF High Availability: Cisco RRM

RRM Channel Hole Detection Mitigation (CHDM) in action

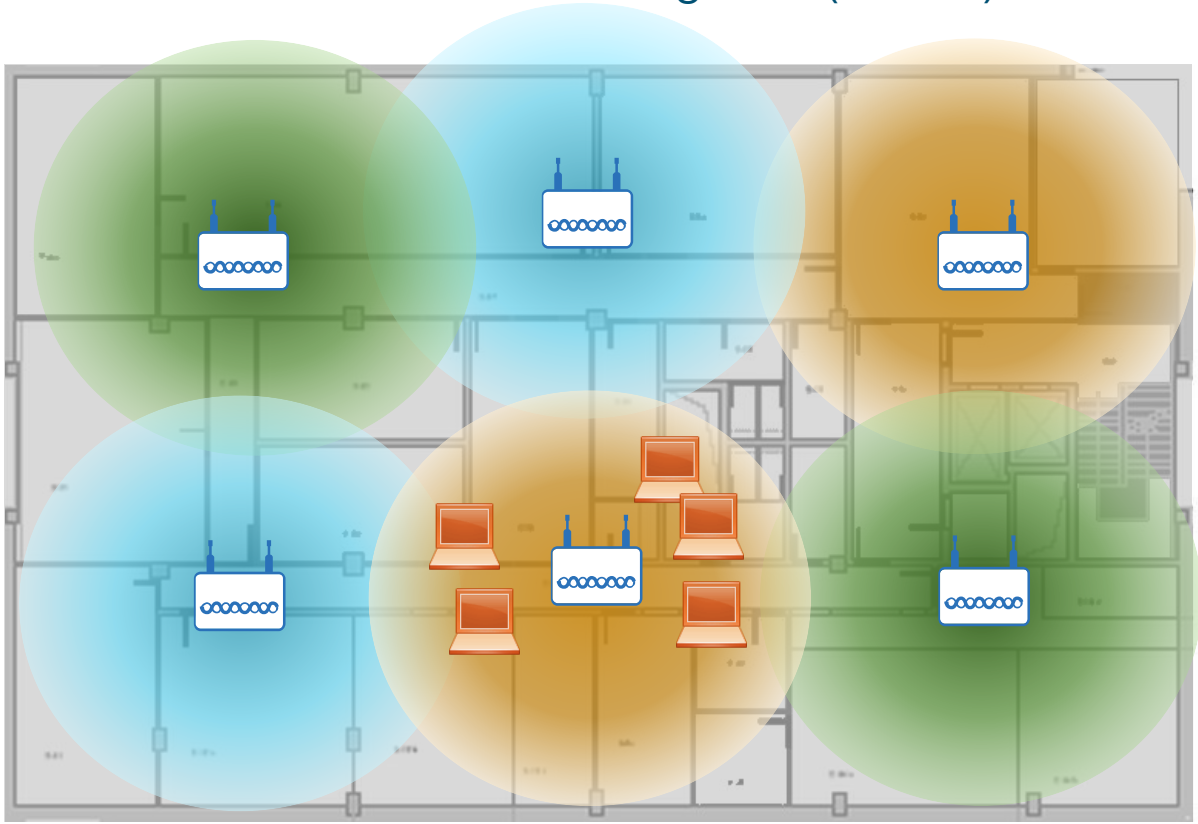

- RRM will determine the optimal Power plan based on AP layout
- Each client RSSI is tracked by AP and reported to WLC
- If an AP fails…

RF High Availability: Cisco RRM

RRM CHDM in action

RRM Details and more: Improve WLAN Spectrum Quality with Cisco's advanced RF (BRKEWN-3010)

- RRM will determine the optimal Power plan based on AP layout
- Each client RSSI is tracked by AP and reported to WLC
- If an AP fails…
- CHDM algorithms kicks in and increases power of neighboring cells within 90 secs
- Clients roam to new APs
- This happens if the CHDM conditions are met:
  - Clients are below the RSSI threshold
  - Min Failed client per AP (#3 default)
  - Coverage Exception Level per AP (25% by default)
  - Failed packets (number and %)
- These checks are needed to avoid false positives

Summary

Cisco provides well engineered Access Points, Antennas, and Radio Resource Management features in the controllers

However, you need to understand the general concepts of radio – otherwise, it is very easy to end up implementing a network in a sub-optimal way:

"RF Matters"

Agenda

- What to do at the Radio Frequency layer?
- **HA Design and Deployment Practices**
  - Central/Large Site Deployments
  - Remote/Small Site Deployments
- Wireless Controller Features for Planned Outages
- Key takeaways

Connecting Access Points and Controllers

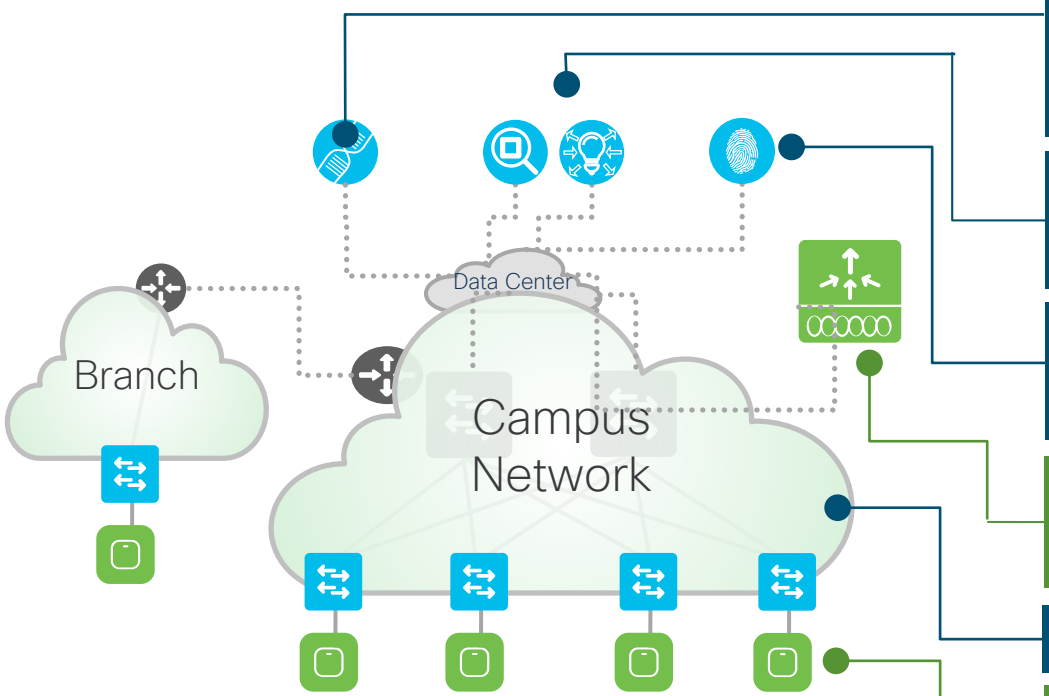

Cisco DNA Spaces
- See how people and things behave on site
- Act on insights with digitization toolkits
- Extend capabilities to drive business outcomes

Network Visibility, Automation, & Analytics
- Network Device Management
- Lifecycle Management
- Wireless Heat Maps

Identity Services
- Authentication Services
- Device Profiling
- Portal Services

Wireless Control Plane
- AP Image, Configuration
- WLAN Mobility Services (RF Analytics, Rogue, Interference)
- Appliance, Public or Private Cloud, or on AP

Wired Access and Campus Network

Wireless Access Points

HA Best Practices: Connecting an AP to the wired network

Recommendations:

- Create redundancy throughout the access layer by connecting APs to different switches/stack members/linecards
- If the AP is in Local mode, configure the port as access with SPT PortFast, BPDU guard, etc..
- If the AP is in FlexConnect mode and Local Switching, configure the port as trunk and allow only the VLANs you need

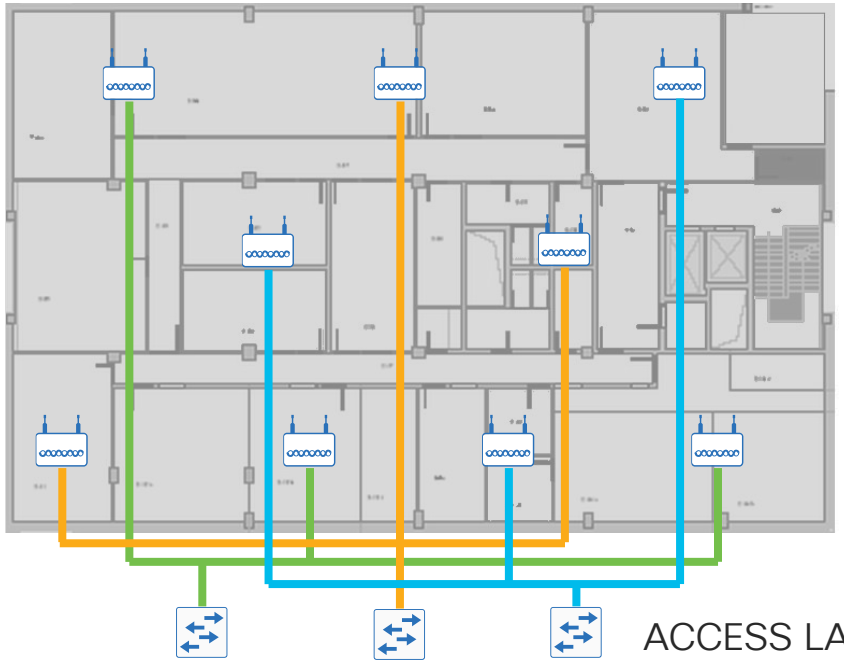

HA Best Practices: Connecting a Single Controller to the wired network

1) To a single Modular Switch or Stack

- Single L2 port-channel*
- Trunk only the required VLANs to the Controller
- Spread ports across Line Cards/Stack members

2) To Redundant Distribution Switches in a StackWise Virtual/VSS pair

- Same as Option 1
- Spread ports across VSS members

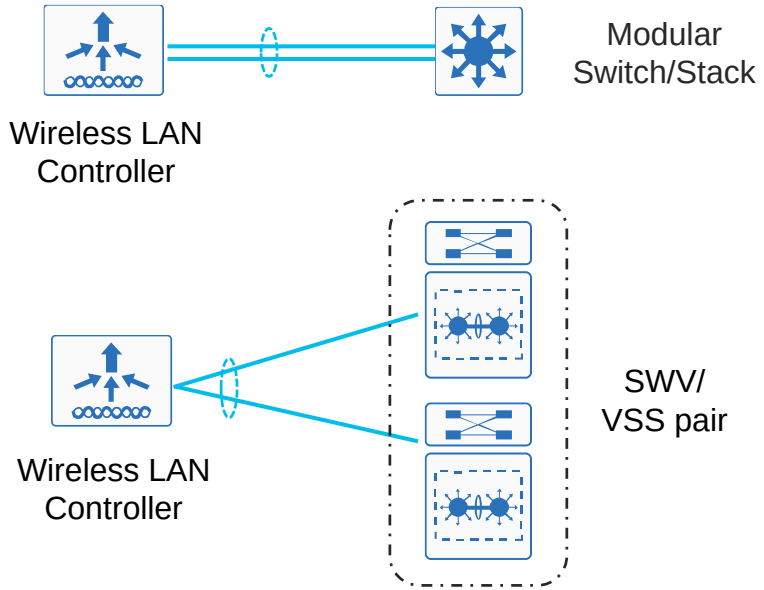

* 9800 Series: PAgP and LACP supported

HA Best Practices: Connecting HA pair to the wired network

Option 1: to single Modular Switch or Stack

- The HA pair of WLCs should be considered as separated WLCs with the same exact configuration
- Ports on both WLCs are UP but only the ones on the Active WLC are forwarding data traffic
- On WLC side: use same physical ports are connected to the network, for ex.: port 1-4 on WLC1 and port 1-4 on WLC2

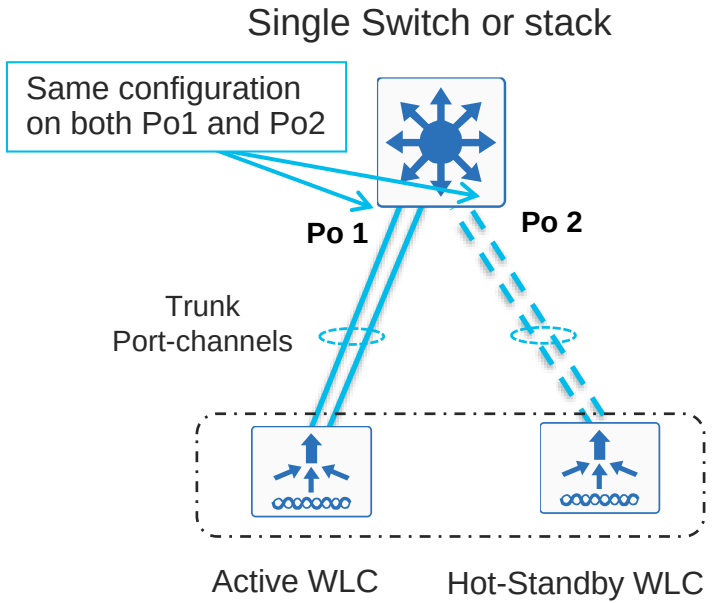

HA Best Practices: Connecting a Client SSO Controller Cluster to the wired network (SWV/ VSS)

Option 2: to StackWise Virtual/ VSS pair

- Use EtherChannel from each Wireless Controller to Distribution StackWise Virtual/ VSS
- Spread the links in each EtherChannel among the two physical switches: **this will prevent a Wireless Controller switchover upon a failure of one of the StackWise Virtual/ VSS switch**
- Keep in mind: Switch scale for ARP and MAC table
- Same applies if switch is a stack/VSL pair/modular switch

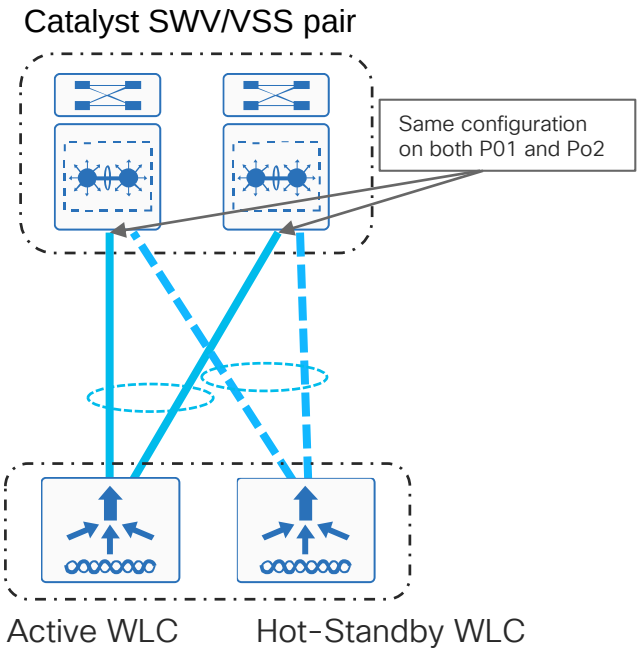

HA Best Practices: Connecting a Client SSO Controller Cluster to the wired network (HSRP)

Option 3: to HSRP pair

- Controller devices are connected to 2 HSRP routers (Active and Standby).
- Failover of HSRP Active to Standby induces a switchover of Wireless Controller HA pair.
- The AP/Clients are up after an SSO. It is a seamless transition and there are no drops on the client.

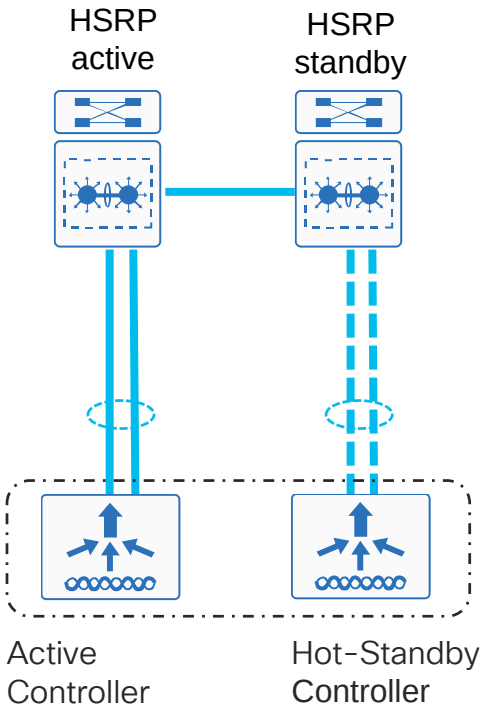

Summary: **Supported** SSO Topologies**

**9800 Series: from IOS XE 17.1.x

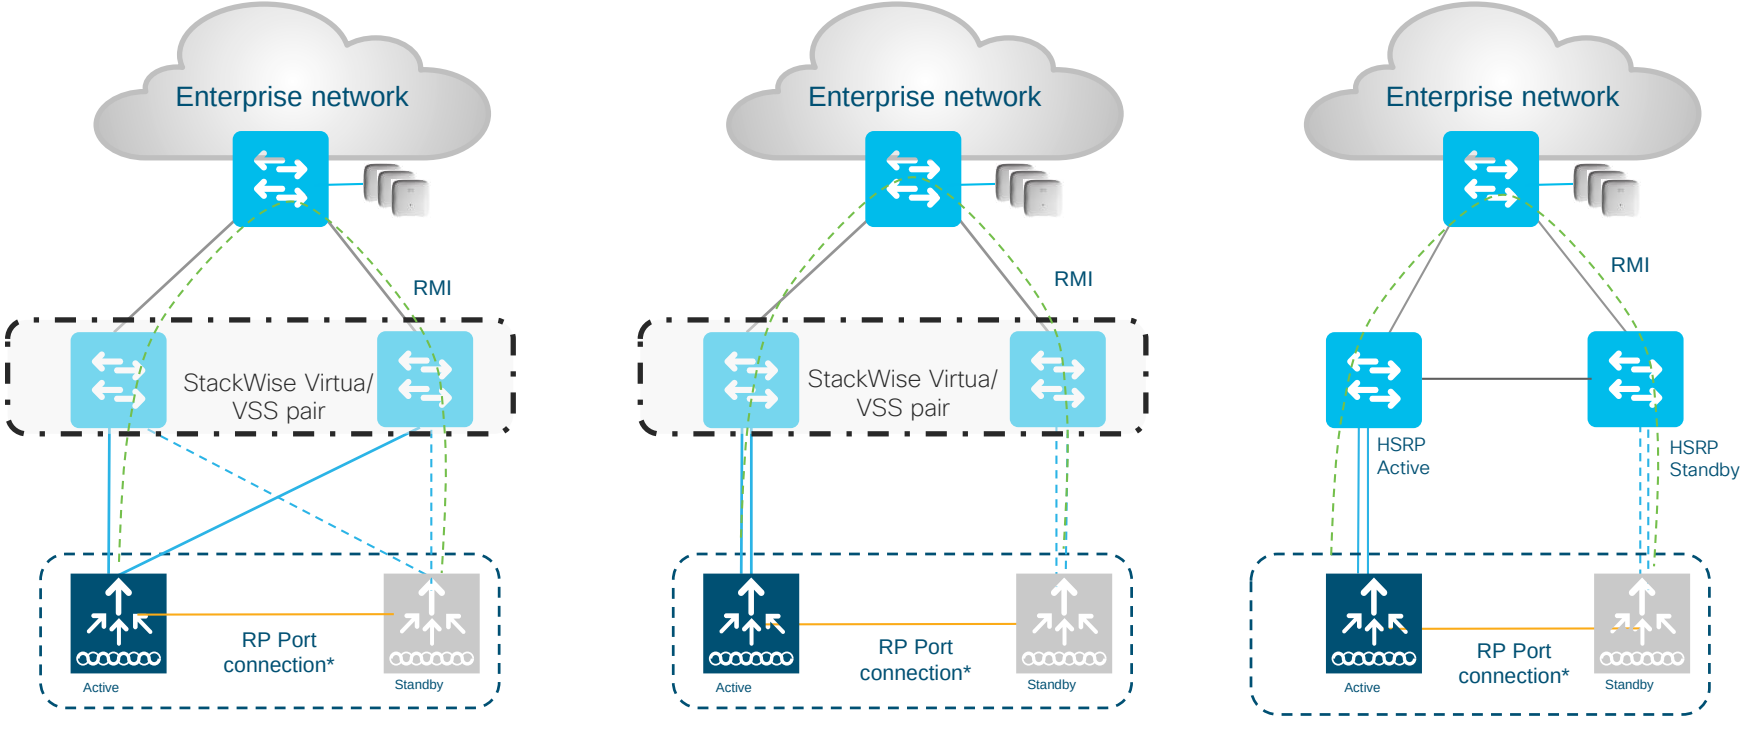

VSS Pair with Split links | VSS Pair – no Split links | HSRP

*Note: RP can be connected back-to-back or via L2 switches

Cisco Wireless Controller Options

EWC
50-100 APs

9800-L
250/500 APs

Catalyst 9800-40
2000 APs

Catalyst 9800-80
6000 APs

Catalyst 9800 Controller Series

C9800 on Switch
(SD-Access only)

Catalyst 9800-Cloud
(private and public)

Catalyst 9800-Cloud (private)
3000-6000 APs

200 APs | 1000 APs | 2000 APs | 3000 APs | 6000 APs

Mobility Express
50-100 APs

WLC 3504
150 APs

WLC 5520
1500 APs

WLC 8540
6000 APs

AireOS WLCs

cisco Live!

Cisco Catalyst 9800 Series – Wireless benefits

Powered by IOS XE
Open and Programmable
Trustworthy Solutions
Modular operating system

Resilient

- Zero downtime with software updates and upgrades
- In Service Software upgrade (ISSU)
- RF/RRM based Rolling AP upgrades

Secure

- Automated macro and micro segmentation with SD-Access
- Detect encrypted threats with Encrypted Traffic Analytics (ETA)

Intelligent

- Deploy in infrastructure of choice and cloud of choice
- Programmable
- Enhanced analytics with Cisco DNA Center

9800 Series Supported SSO Topologies with IOS XE 16.12.x and earlier (No Gateway Check or RMI)

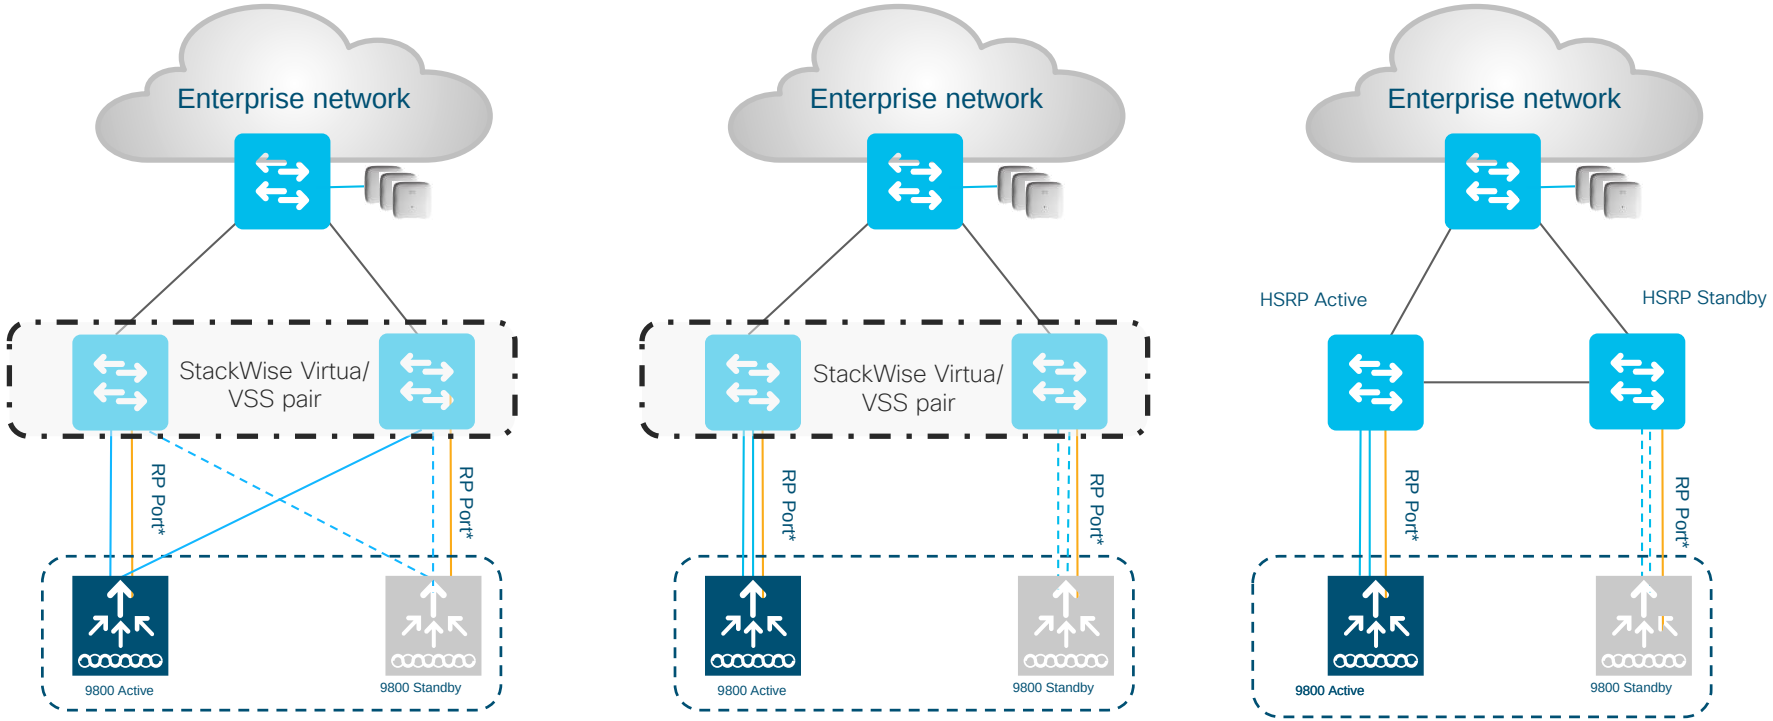

VSS Pair with Split links | VSS Pair – no Split links | HSRP

*Note: RP Redundancy Port (RP) connected to the respective uplink switches/routers

9800 Series SSO Behavior from 17.1

Reference

- Redundancy Management Interface (RMI) introduced
- Gateway Check using RMI introduced
- Dual Active detection
- Direct RP connection (back-to-back or via dedicated switches) supported in case of VSS with split links and HSRP

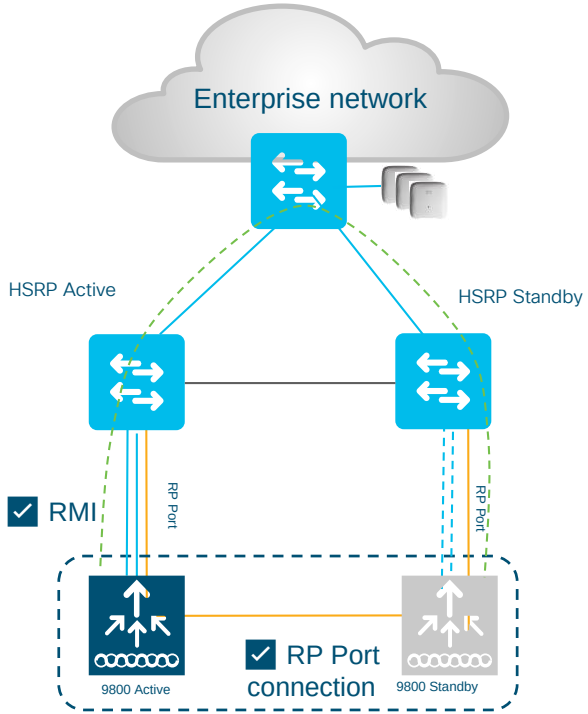

Wireless Controller modes fitting different requirements

Centralized

Ease of Deployment and management for large campuses. Cloud and non-Cloud options.

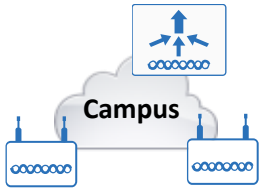

SDA-Wireless

Policy Segmentation and consistent wired-wireless management

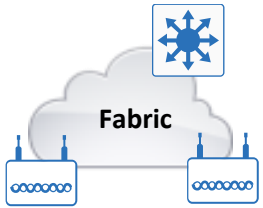

Flex Connect

Eliminate the need for a Controller at every Site for a distributed deployment. Cloud and non-Cloud options.

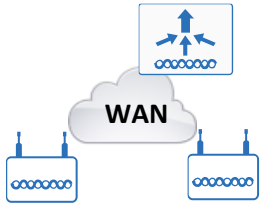

Mobility Express and EWC

Simplified Controller-less deployment for distributed deployments and small sites

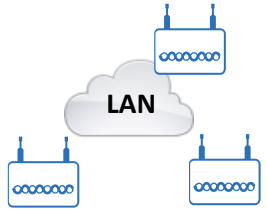

Agenda

- What to do at the Radio Frequency layer?
- HA Design and Deployment Practices
  - **Central/Large Site Deployments**
  - Remote/Small Site Deployments
- Wireless Controller Features for Planned Outages
- Key takeaways

Centralized Mode High Availability: SSO and N+1

	Requirements	Benefits
Client SSO	• Catalyst 9800 Series • 5520, 8540, 3504 WLC • **L2 connection** • Same HW+SW Version • 1:1 box redundancy	Active Client State is synched AP state is synched No Application downtime
N+1 Redundancy (Deterministic/Stateless HA, a.k.a.: primary/secondary/tertiary)	Each Controller has to be configured separately	Available on all controllers **Crosses L3 boundaries** Flexible: 1:1, N:1, N:N

N+1 Redundancy

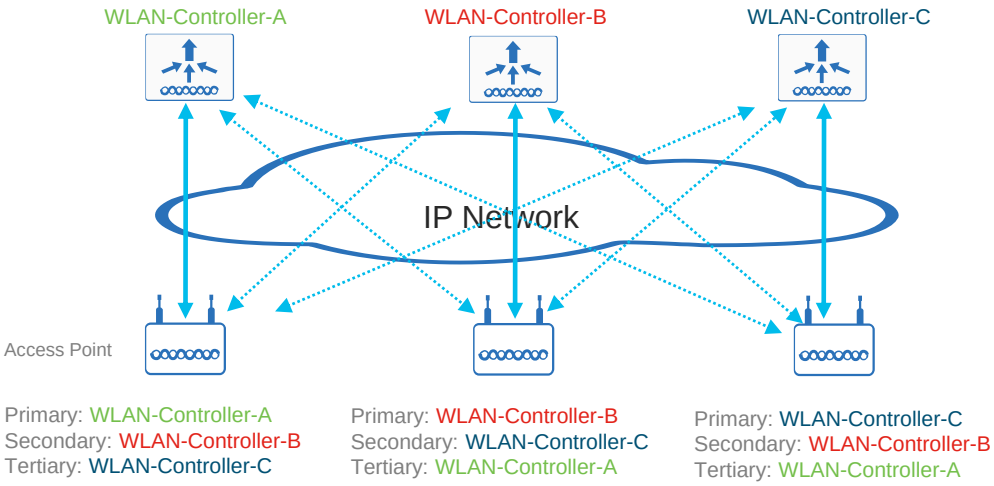

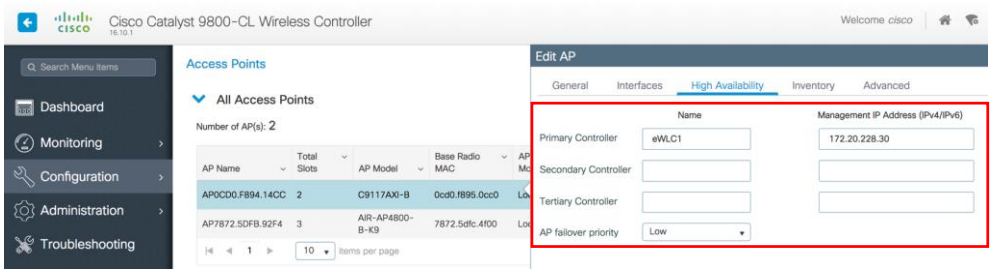

- Administrator statically assigns APs a primary, secondary, and/or tertiary controller
  - Assigned from controller interface (per AP) or Prime Infrastructure (template-based)
  - You need to specify Name and IP if WLCs are not in the same Mobility Group
- **Pros:**
  - Support for L3 network between WLCs
  - Flexible redundancy design options:1:1, N:1, N:N:1
  - WLCs can be of different HW and SW (*)
  - “Fallback” option in the case of failover
  - Can overload APs on controllers (using AP priority)
- **Cons:**
  - Stateless redundancy. There is a network downtime when the WLC fails
  - More upfront planning and configuration

(*) AP will need to upgrade/downgrade code upon joining

N+1 Redundancy

Global backup Controllers

Catalyst 9800 Controller Series

Configuration > AP Join >…

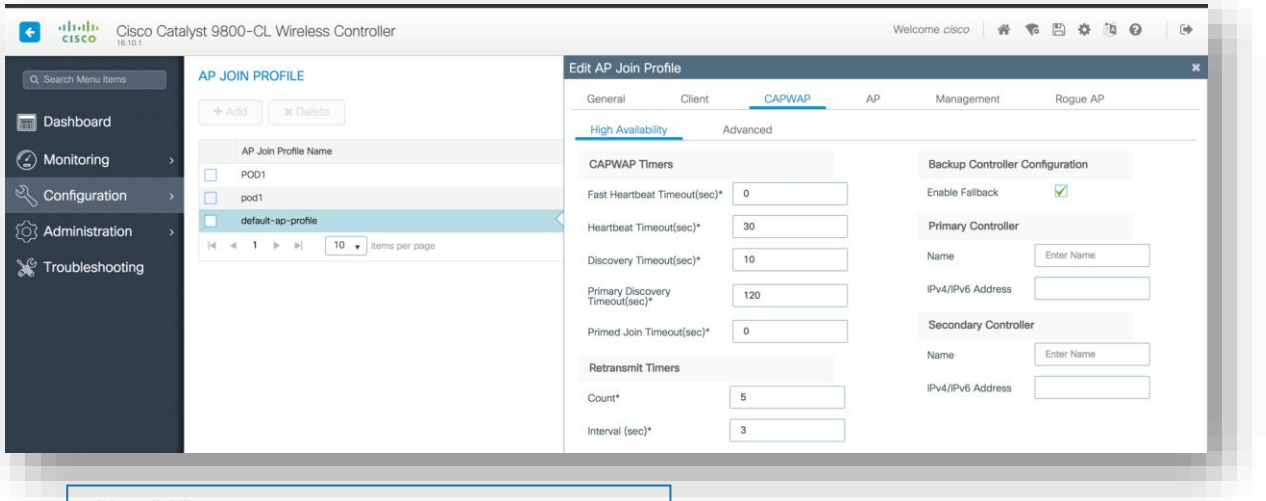

AireOS

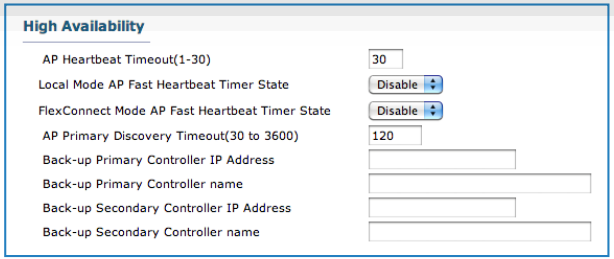

Wireless > High Availability

- Used if there are no primary/secondary/tertiary WLCs configured on the AP
- The backup controllers are added to the primary discovery response message to the AP

N+1 Redundancy

AP Failover mechanism

< 30-45 sec (*)

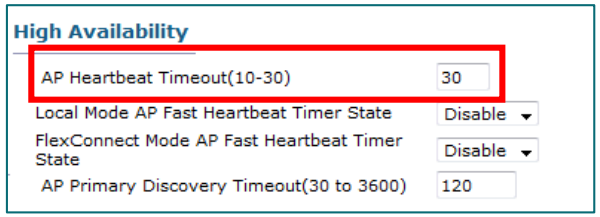

Reference

When configured with Primary and backup Controllers:

- AP uses heartbeats to validate current WLC connectivity
- Upon loosing a heartbeat to the Primary, AP sends 5 consecutives heartbeats every 3 second (default)
  - Configurable to minimum of 3 keepalive every 2 sec
- If no reply, AP declares the WLC dead and starts the join process to the first backup WLC candidate:
  - Backup is the first alive WLC in this order: primary, secondary, tertiary, global primary, global secondary.
- With N+1 Failover, AP goes back to discovery state just to make sure the backup WLC is UP and then immediately starts the JOIN process
- With N+1, AP periodically checks for Primary to come back online and falls back to it (AP fallback can be disabled)

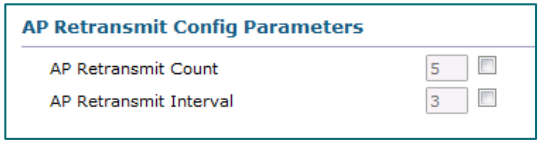

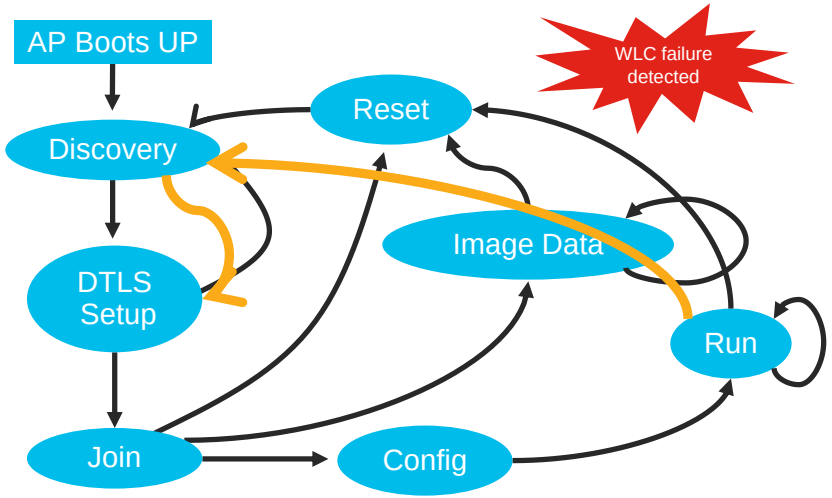

(*) With Fast Heartbeat and minimum values for keepalive

N+1 Redundancy

AP Failover Priority

- Assign priorities to APs: Critical, High, Medium, Low
- Critical priority APs get precedence over other APs when joining controller
- If backup controller doesn't have enough licenses/capacity existing lower priority APs will be dropped to accommodate higher priority APs.

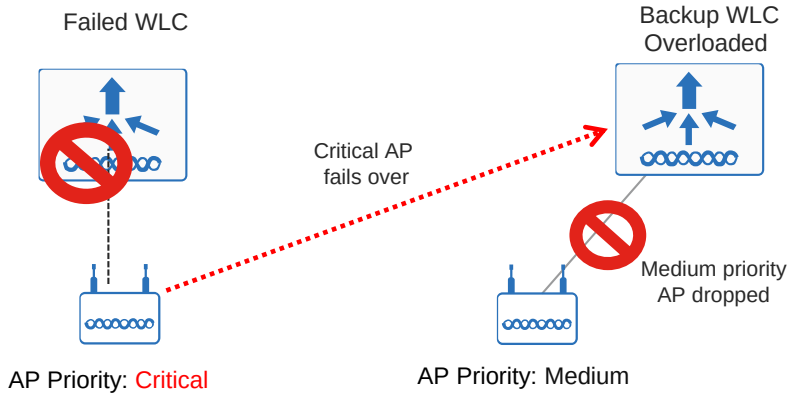

Reference

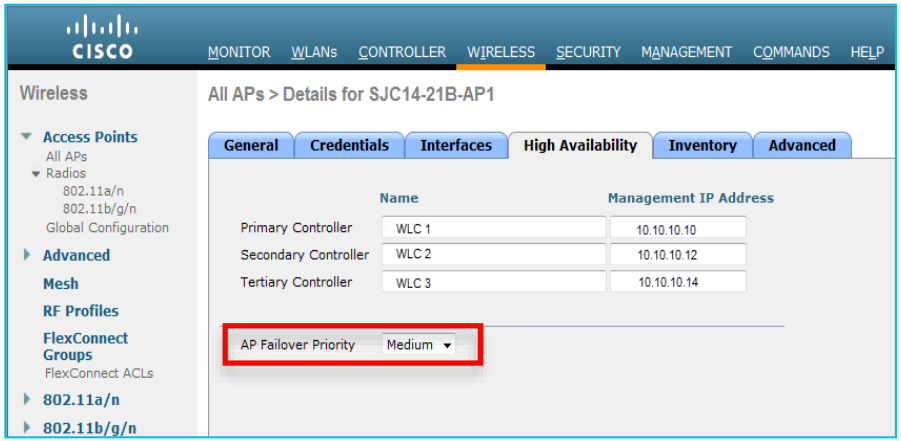

Summary N+1 Redundancy

- Most common Design is N+1 with redundant WLC in a geographically separate location across L3 Campus
- Can provide 30-45 sec of downtime when use faster heartbeat to detect failure
- Use AP priority in case of oversubscription of redundant WLC

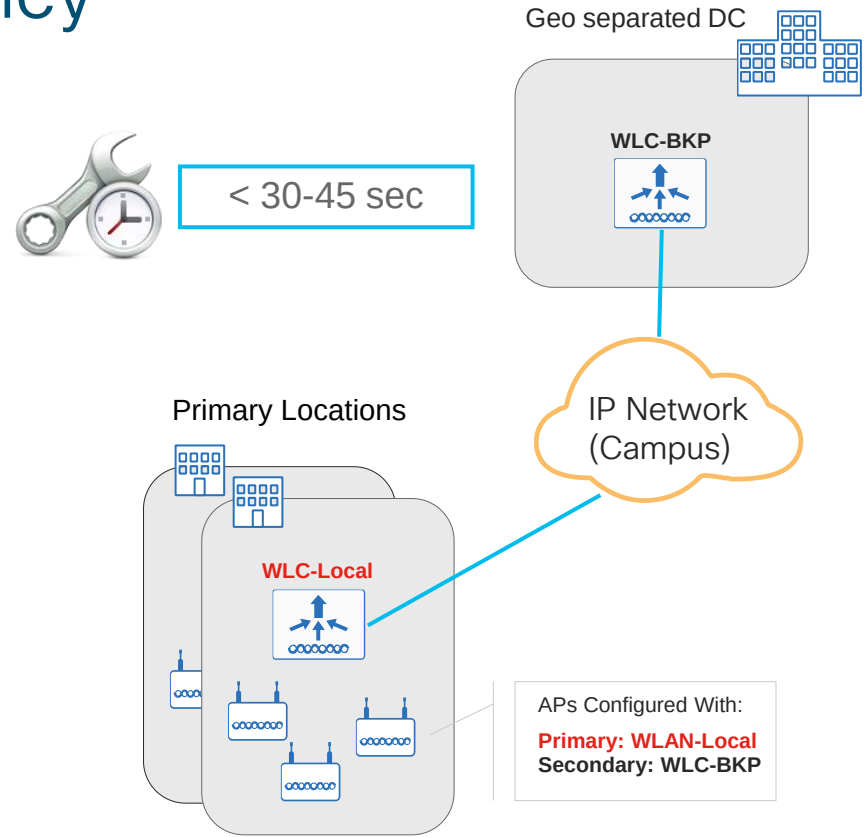

For more info:
http://www.cisco.com/en/US/docs/wireless/technology/hi_avail/N1_HA_Overview.html

Centralized Mode High Availability: SSO and N+1

	Requirements	Benefits
Client SSO	• Catalyst 9800 Series • 5520, 8540, 3504 WLC • L2 connection • Same HW+SW Version • 1:1 box redundancy	Active Client State is synched AP state is synched No Application downtime
N+1 Redundancy (Deterministic/Stateless HA, a.k.a.: primary/secondary/tertiary)	Each Controller has to be configured separately	Available on all controllers Crosses L3 boundaries Flexible: 1:1, N:1, N:N

Stateful Switchover (Client SSO)

< 1 sec

Sub-second failover and zero SSID outage

- HA Pairing is possible only between the **same type of hardware** and **software** versions
- True **Box to Box High** Availability i.e. 1:1
  - One WLC in **Active state** and second WLC in **Hot Standby state**
  - Secondary continuously monitors the health of Active WLC via L2 connection (Redundancy Port).
- Configuration on Active is synched to Standby WLC
  - This happens at startup and incrementally at each configuration change on the Active
- What else is synched between Active and Standby?
  - Licenses, AP CAPWAP state, Clients in “RUN” state
- There is no preemption in Controller SSO: when the failed Active WLC comes back online it will joining as Hot Standby

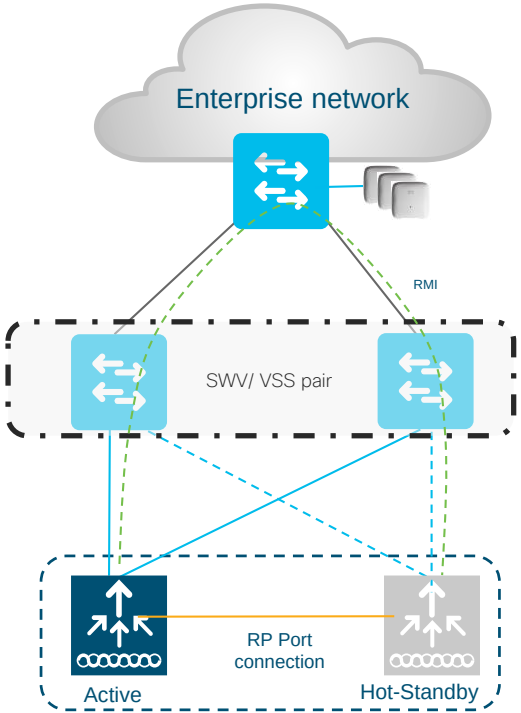

C9800 **Private Cloud** Deployment: Client SSO High Availability

ESXi KVM

C9800-CL-K9

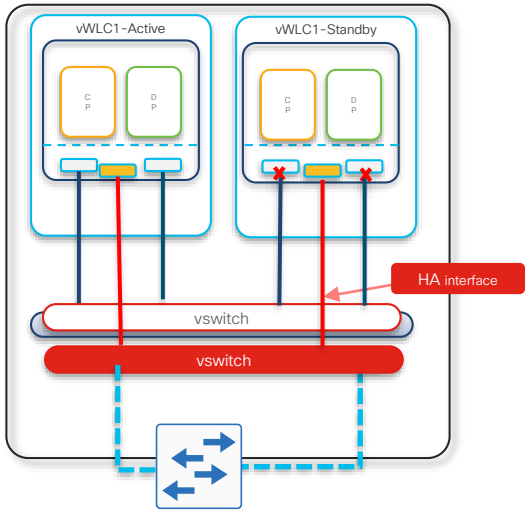

Redundancy Port Connectivity

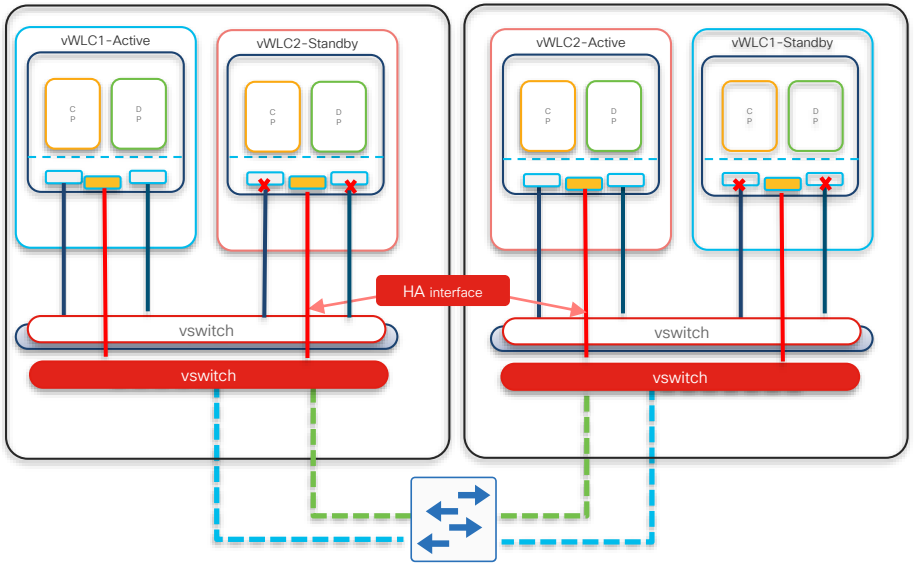

Redundancy Port Connectivity
RP via L2

Redundancy on Catalyst 9800 Wireless Controller

Configuration

- Both C9800-40-K9 and C9800-80-K9 Wireless controllers have two RP Ports:
  - RJ-45 Ethernet Redundancy port
  - SFP Gigabit Ethernet Port
- If both the Redundancy Ports are connected, SFP Gigabit Ethernet port takes precedence:
  - HA between RJ-45 and SFP Gigabit RP ports is not supported.
  - Use only Cisco supported SFPs
  - When HA link is up through RJ-45, SFPs on HA port should not be inserted even if there is no link between them.

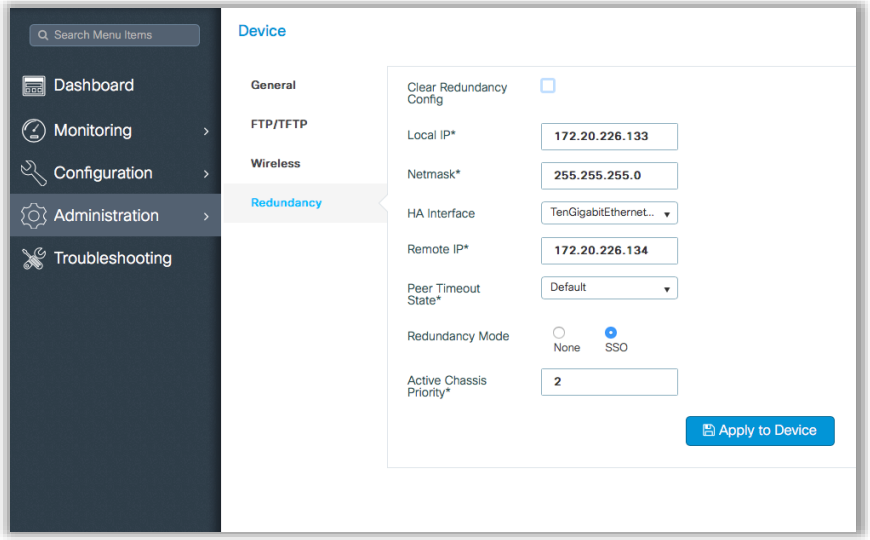

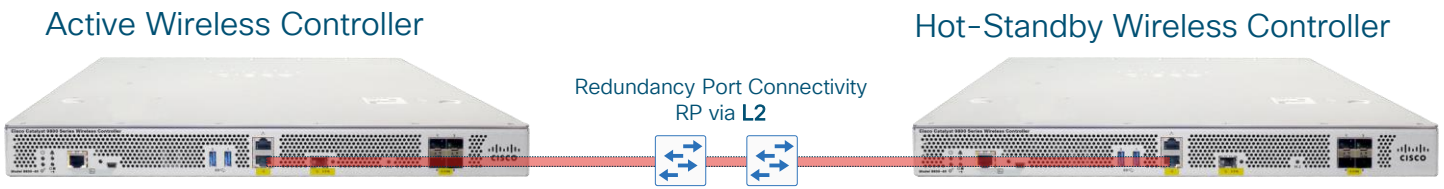

Stateful Switchover (SSO)

Failover sequence

1. Redundancy role negotiation and config sync
2. APs associates with Active controller
3. Client associates with Active through AP
4. Active failure: notify peer / or missing keep alive
5. Standby WLC sends out GARP
6. Standby becomes Active:
   - AP DB and Client DB are already synced to standby controller
   - AP CAPWAP tunnel session intact
   - Client session intact, client does not re-associate*

< 1 sec

Effective downtime for the client is:
Detection time + Switchover time

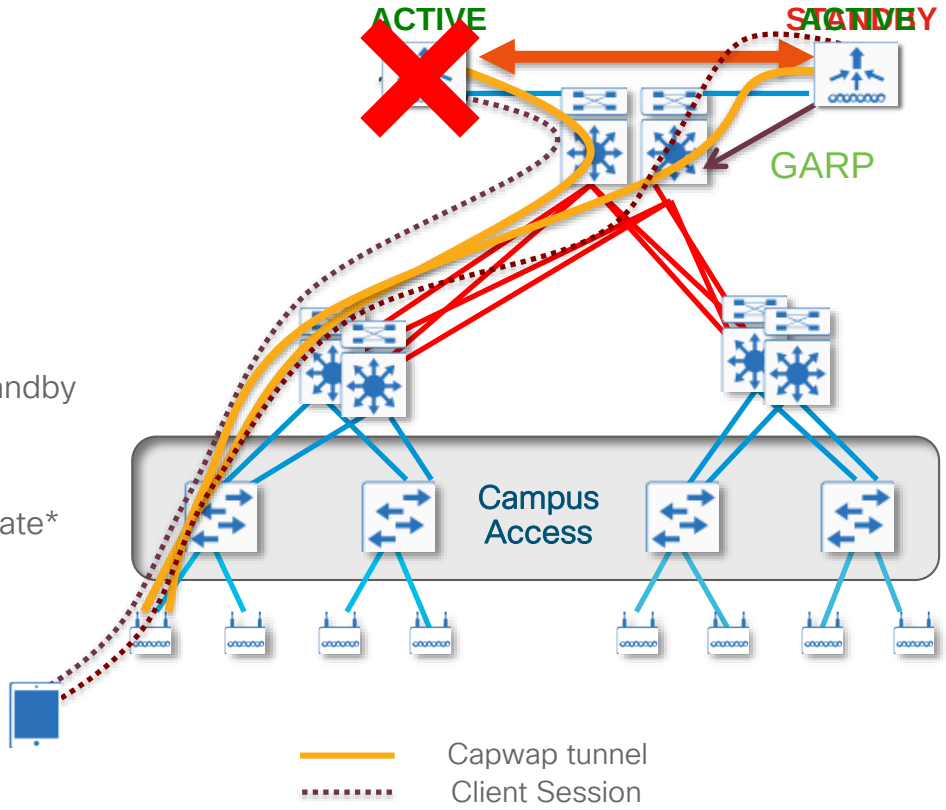

HA Deployment Best Practices for Campus

N+1	SSO	SSO + N+1	SSO + SSO
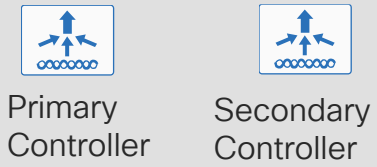	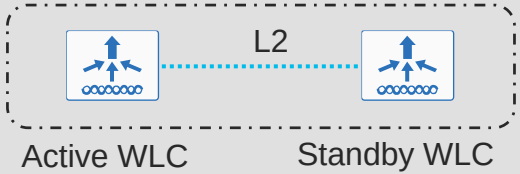	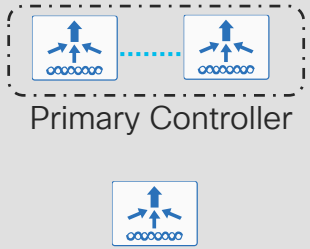	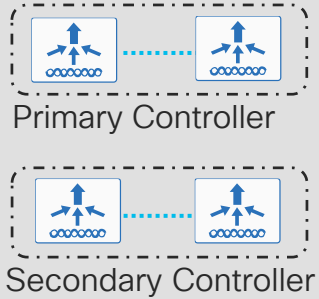
• Approx. 30 Sec failover time (AP+Client affected) • No Config Synch (risk: Config mismatch) • AP loadbalancing • L2 or L3	• Sub-Second Failover (Client+AP not affected) • Config Synch • One active, one standby (no AP loadbalancing) • L2 connection needed	adds redundancy and simplifies operation during maintenance (e.g. SW Updates)	adds redundancy and simplifies operation during maintenance (e.g. SW Updates)

Multi-site Campus: Combine SSO with N+1

- SSO pair can act as the Primary Controller and be deployed with single Secondary and Tertiary WLC
- Network downtime:
  - No network downtime for single controller failure in the Primary DC
  - On failure of both Active and Hot-standby WLC, APs will fall back to secondary/ tertiary controller
- Recommendations:
  - Make sure that AP Fallback is enabled
  - Use AP Failover priority in case of oversubscription of the backup WLC
  - Useful to reduce downtime for SSO pair software upgrade

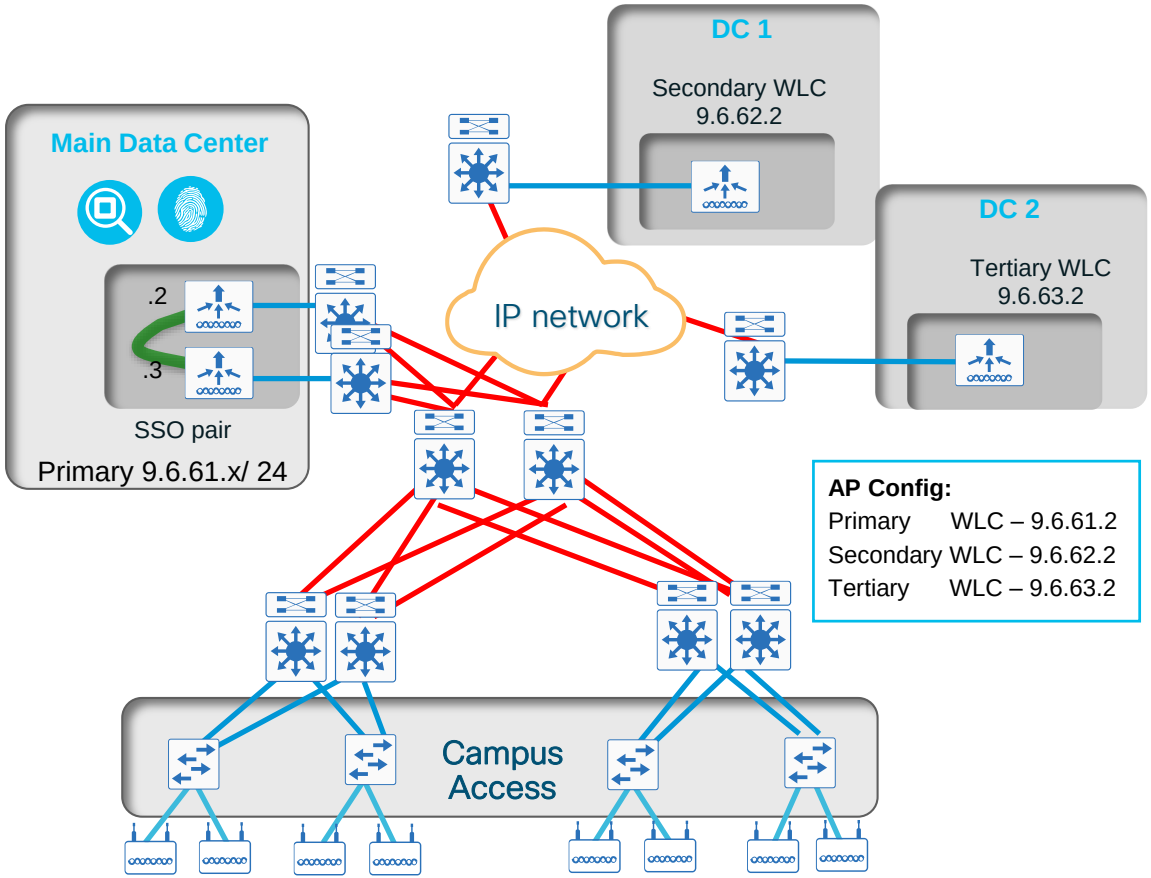

Multi-Site Campus: SSO everywhere!

- Each site can be its own separated SSO architecture
- Full site redundancy by assigning primary, secondary, tertiary to the APs.
- Max level of High Availability: no network downtime upon controller failure within any site.

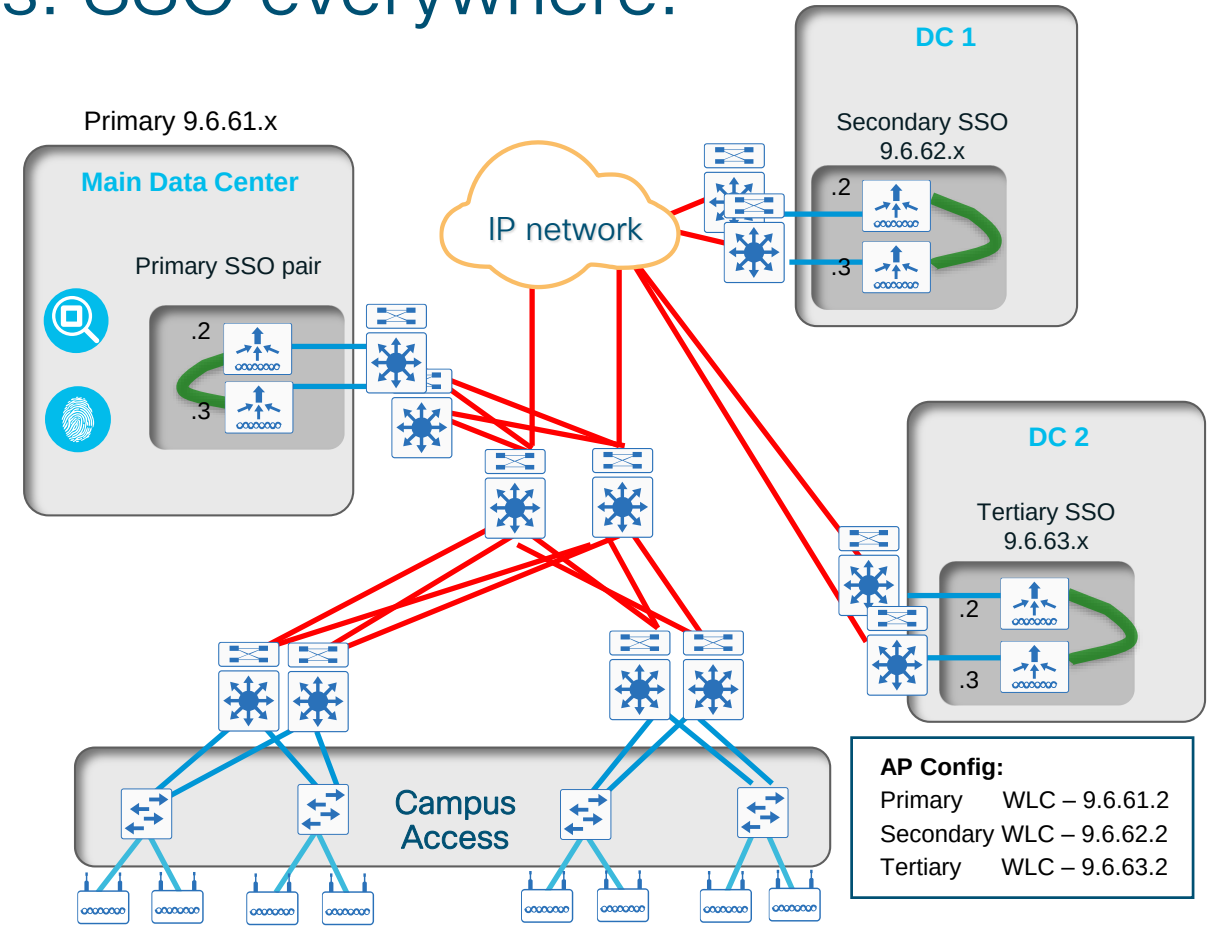

Key Considerations: Campus HA Deployment Best Practices

- What is the acceptable downtime for your business applications?
  - No downtime? Go with Stateful Switchover (Client SSO).
  - Are 30 sec to few minutes ok? Go with N+1 to have more deployment flexibility
- What is the downtime to upgrade a HA pair and how to minimize it?
  - Catalyst 9800 Wireless Controller: use built-in Rolling SW Upgrade
  - AireOS Controllers (details for reference only):
    - Plan for additional backup controller
    - Use Prime Infrastructure Rolling SW Updates Feature

Agenda

- What to do at the Radio Frequency layer?
- HA Design and Deployment Practices
  - Central/Large Site Deployments
    - SDA
  - Remote/Small Site Deployments
- Wireless Controller Features for Planned Outages
- Key takeaways

Software Defined Access: Bringing Intent Based Networking to Life

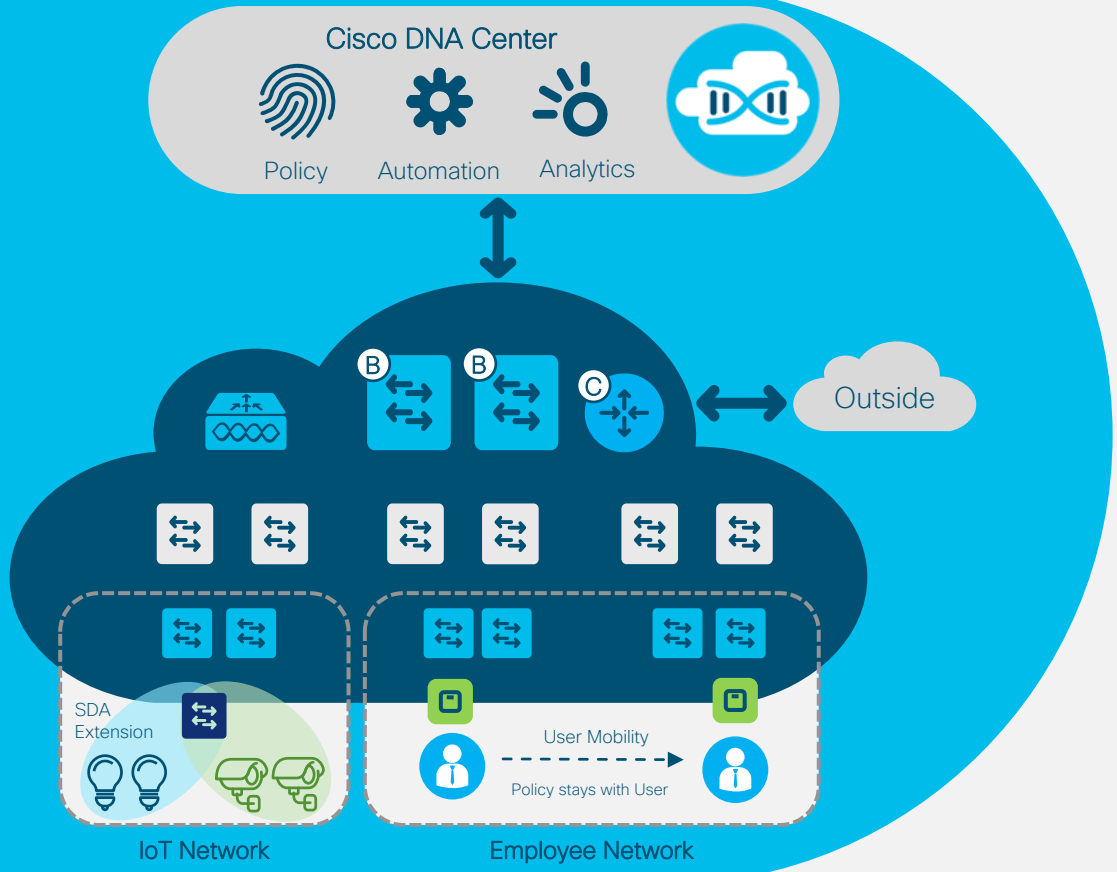

Automated Network Fabric

Single Fabric for Wired & Wireless with simple Automation

Identity-Based Policy & Segmentation

Decouples Security & QoS from VLAN and IP Address

Insights & Telemetry

Analytics and Insights into User and Application behavior

SD-Access Wireless: Redundancy Considerations (Controller outside Fabric)

To a shared Service block

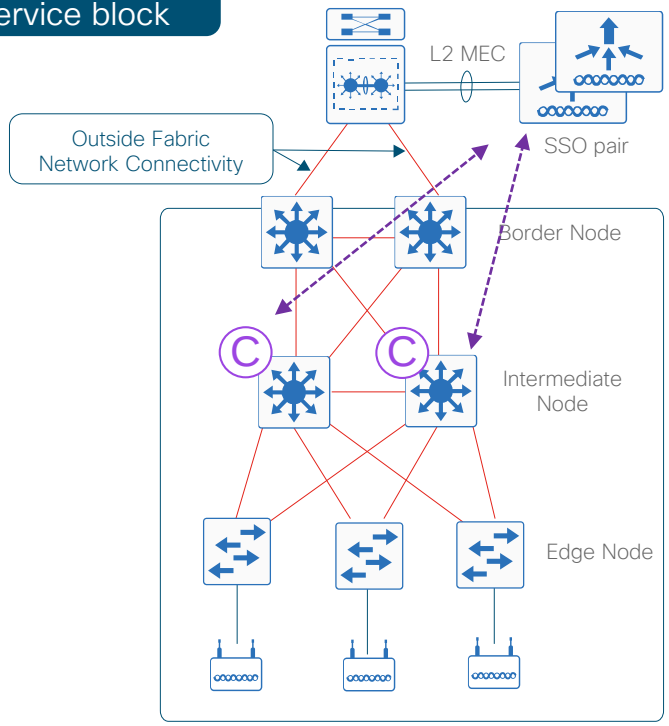

SD-Access Fabric

- WLC registers wireless clients in Host Tracking DB
- Control Plane (CP) redundancy is supported in Active / Active configuration
- WLC is configured with two CP nodes with information sync across both
- Stateful redundancy with WLC SSO pair. Active WLC updates Control nodes

Directly to the pair of FBs

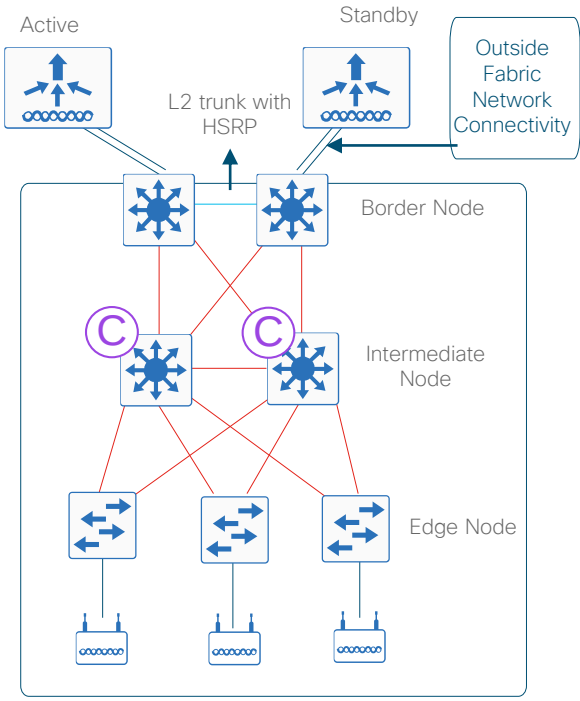

SD-Access Fabric

HA with SD-Access Embedded Wireless

9300, 9400, 9500
DNA 1.3.x

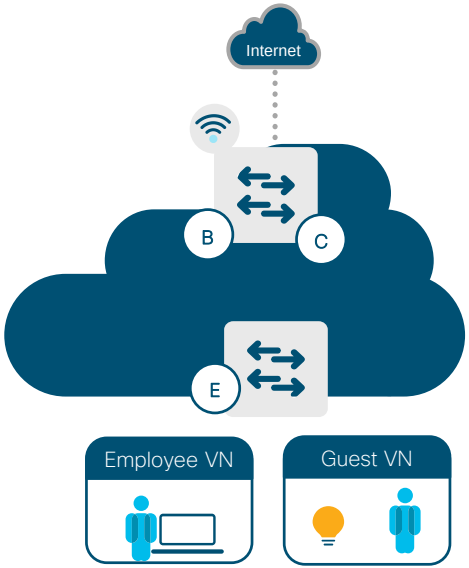

Co-Located Border + CP with
Cat9800 Embedded Wireless

9300, 9400, 9500
DNA 1.3.x

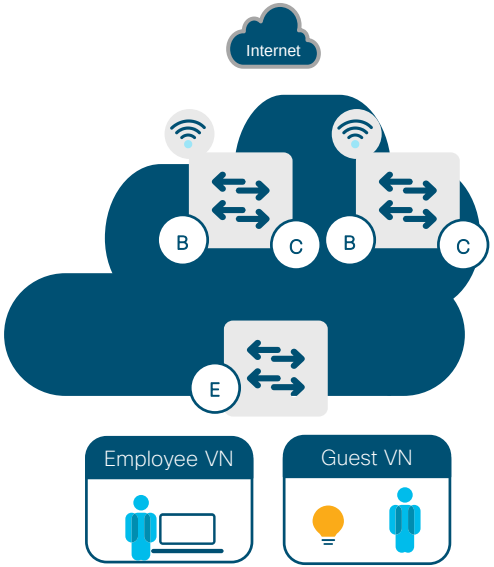

Multiple Co-Located Border + CP with
Cat9800 Embedded Wireless

- You can have up to #2 for scale to 400 APs
- The 9800 WLCs will be configured in the same Mobility Group for roaming
- SSO HA is supported within the stack but NOT across stacks
- Automated N+1 support → on roadmap

cisco Live!

SD-Access Embedded Wireless Fabric in a Box

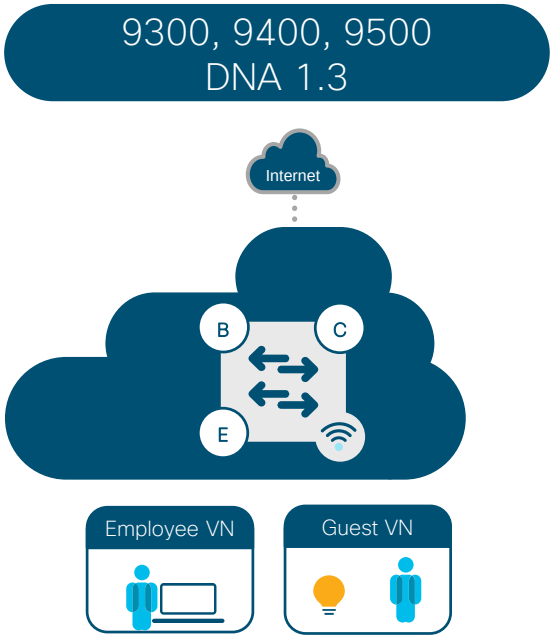

Fabric in a Box

with Cat9800 Embedded Wireless

- Only one FiaB per Fabric site
- SSO supported within the stack

AP Scale	Client Scale
100 (200 in 16.11)	4000

Platforms supporting SD-Access Wireless

Optimized for Distributed Braches → Small and Medium Campus → Medium and Large Campus

On Switch

- Cisco IOS® XE Software
- Cat 9300, Cat 9400, Cat 9500
  - 200 AP, 4k Clients
- SD-Access wireless with Embedded Cat9800 Software Package

On Private Cloud

- Cisco IOS® XE Software
- C9800-CL
  - 1k AP, 10k Clients
  - 3k AP, 32k Clients
  - 6k AP, 64k Clients^
- Scale on demand

On Appliance

- Cisco IOS® XE Software
  - C9800-40-K9
  - C9800-80-K9
  - C9800L
- Cisco AireOS Software:
  - WLC 3504 (SW8.8)
  - WLC 5520 (SW8.8)
  - WLC 8540 (SW8.8)

Agenda

- What to do at the Radio Frequency layer?
- HA Design and Deployment Practices
  - Central/Large Site Deployments
  - Remote/Small Site Deployments
- Wireless Controller Features for Planned Outages
- Key takeaways

HA Deployment Best Practices: Remote Site/Small Site Key Design Questions

Local Controller

Controller (Appliance/virtual)

- Specific per branch configuration
- Independency from WAN quality
- Reduced configuration on switches
- Full feature support
- L3 roaming supported

Mobility Express and EWC

- Specific per branch configuration
- Independency from WAN quality
- low hardware footprint (Controller running on Access Point)

FlexConnect

- Single pane of Mgmt. & Troubleshooting
- Reduced branch footprint
- Built-in resiliency
- Perfect fit for centralized IT Team

- HA questions:
  - Is the remote site independent from the Central site from an operation prospective?
  - What is the traffic flow of your application? Are the APP servers centrally located?
  - Is there a local Internet breakout? How do you authenticate new users if WAN/Controller is down? Where is the AAA server located?

Local Controller Summary

“Do your clients need full Enterprise feature set (even if WAN is down)?”

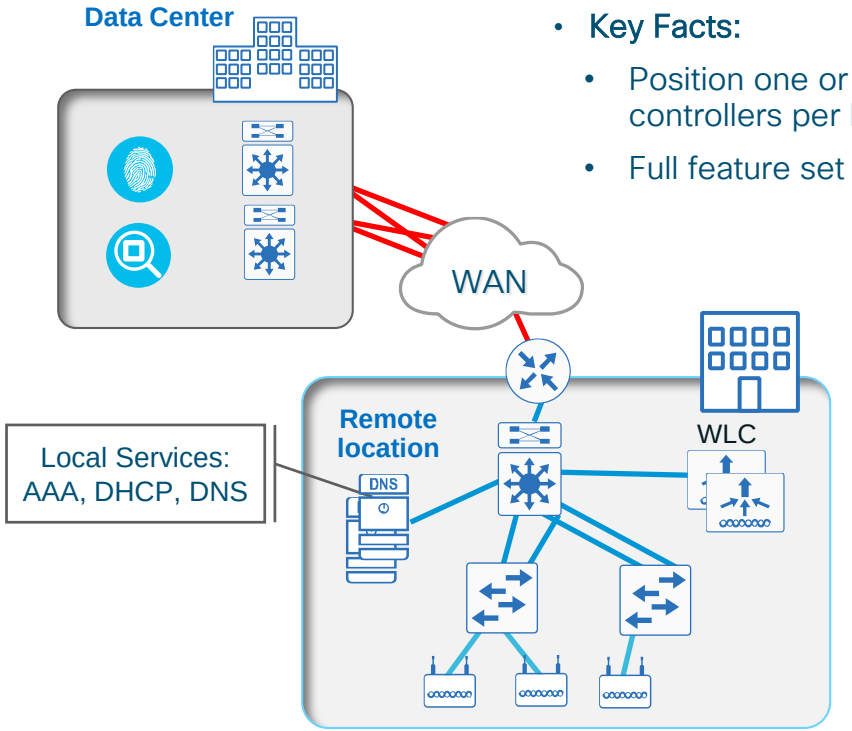

- Key Facts:
  - Position one or two controllers per branch
  - Full feature set available

When to use:

- WAN Bandwidth and latency is a concern
- Simple configuration on the switch port connected to the Access Point desired
- Branch/local IT staff requires configuration outside of corporate standard
- L3 Roaming is needed

High Availability:

- Full features available if WAN is down
- use N+1 or SSO for site controller redundancy
- Local Authentication, DHCP, DNS required for full WAN Independency

Keep in Mind:

- Need to manage each site individually
- Prime Infrastructure should be considered for central manageability

cisco Live!

HA Deployment Best Practices: Remote Site/Small Site Key Design Questions

Local Controller

Controller (Appliance/virtual)

- Specific per branch configuration
- Independency from WAN quality
- Reduced configuration on switches
- Full feature support
- L3 roaming supported

Mobility Express and EWC

- Specific per branch configuration
- Independency from WAN quality
- low hardware footprint (Controller running on Access Point)

FlexConnect

- Single pane of Mgmt. & Troubleshooting
- Reduced branch footprint
- Built-in resiliency
- Perfect fit for centralized IT Team

EWC Controller Function embedded on Cisco Catalyst access point

or “Mobility Express” for Catalyst APs ☺

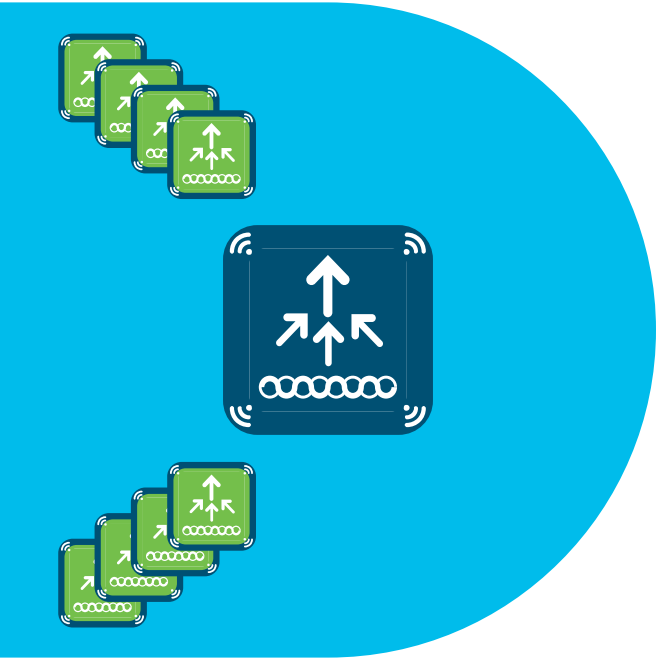

Runs 9800 Series Cisco IOS® XE wireless controller on Cisco Catalyst access points	Modern OS, scalable, open and programmable, supports telemetry
Supports advanced enterprise feature set	HA, SMU, adaptive wireless IPS (aWIPS), Cisco Umbrella™, NetFlow, ICAP
Flexible management options	Use mobile app, WebUI, and Cisco DNA Center to deploy, manage, and monitor
Investment protection	Migrate access points to controller for more than 100 access points

EWC on Catalyst AP vs. Mobility Express

EWC on 9100 Series

"9800 Controller running on Catalyst Access Point"

- ✓ Full enterprise Feature set
- ✓ Same deployment architecture as Mobility Express
- ✓ Same IOS XE look and feel across all Catalyst 9800 Series Controllers (GUI and CLI)
- ✓ Support Wave 2 APs (x800 Series) as subordinate
- ✓ Enhanced HA (SMU, AP Service Pack/Device Pack)
- ✓ Scale: 50-100 Access Points

Mobility Express on W2 APs

"AireOS Controller running on W2 Access Point"

- ✓ Reduced feature set/new GUI
- ✓ ME only runs on Wave 2 APs (x800 Series), other APs including Catalyst 9100 can operate as subordinate
- ✓ Scale: 50-100 Access Points

EWC on Cisco Catalyst 9100 access points*

Reference

*requires IOS XE 16.2.2

Ideal for single or multisite small to medium-sized enterprise deployments

C9115AX-EWC

- 50 APs, 1000 clients
- 4x4 + 4x4
- MU-MIMO, OFDMA
- Spectrum Intelligence
- Bluetooth 5
- 1x 2.5 Multigigabit
- USB
- Integrated or external antenna

C9117AX-EWC

- 50 APs, 1000 clients
- 8x8 + 4x4
- MU-MIMO, OFDMA (only DL)
- Spectrum Intelligence
- Bluetooth 5
- 1x 5 Multigigabit
- USB
- Integrated antenna only

Mission critical

Best suited for high-density enterprise branch deployments

Powered by Cisco RF ASIC

C9120AX-EWC

- 100 APs, 2000 clients
- 4x4 + 4x4
- MU-MIMO, OFDMA
- Cisco RF ASIC
- Dual 5 GHz, HDX
- RF signature capture
- 1x 2.5 Multigigabit
- Integrated or external antenna

Best in class

Powered by Cisco RF ASIC

C9130AX-EWC

- 100 APs, 2000 clients
- 8x8 + 4x4 or 4x4 + 4x4 + 4x4
- Tri-radio (dual 5 GHz + 2.4 GHz), HDX
- Cisco RF ASIC
- RF signature capture
- Decrypted data packet ICAP
- 1x 5 Multigigabit
- 8-port smart antennas

Software feature parity across APs | Supports up to 100 APs, 2000 clients | Supports Wave 2 APs as client serving | Cisco DNA Assurance with ICAP

What about 802.11ac Wave 2 access points?
→ Supports client serving mode

Reference

Ideal for small to medium-sized deployments

Mission critical

Indoor


1815w


1815i, 1815m


1832


1842


1852


2802


3802


4800

Outdoor


1540


1560

All 802.11ac Wave 2 access points can connect to the embedded wireless controller

Deploying the Cisco Embedded Wireless Controller (and Mobility Express)

- EWC-capable access points can be connected to an access port or a trunk port on the switch, depending on the deployment method
- Management traffic is always untagged

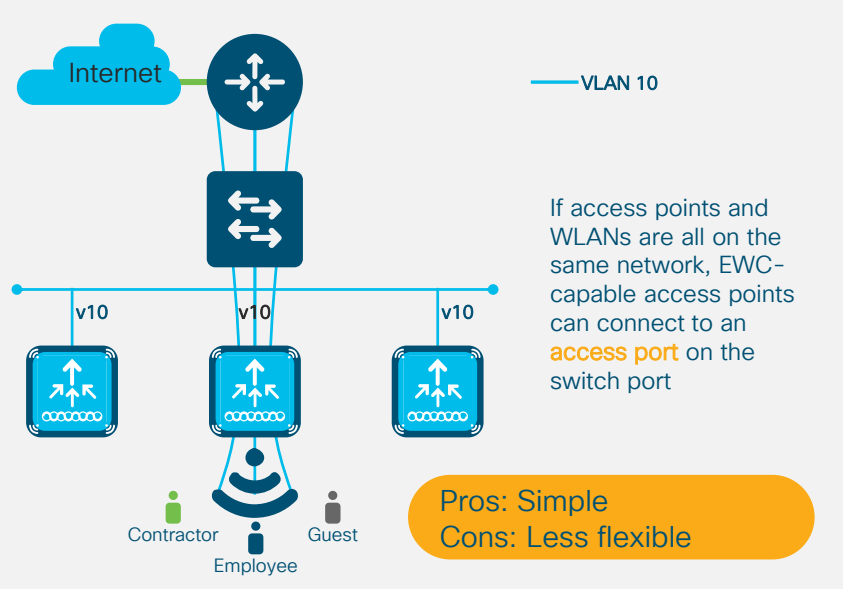

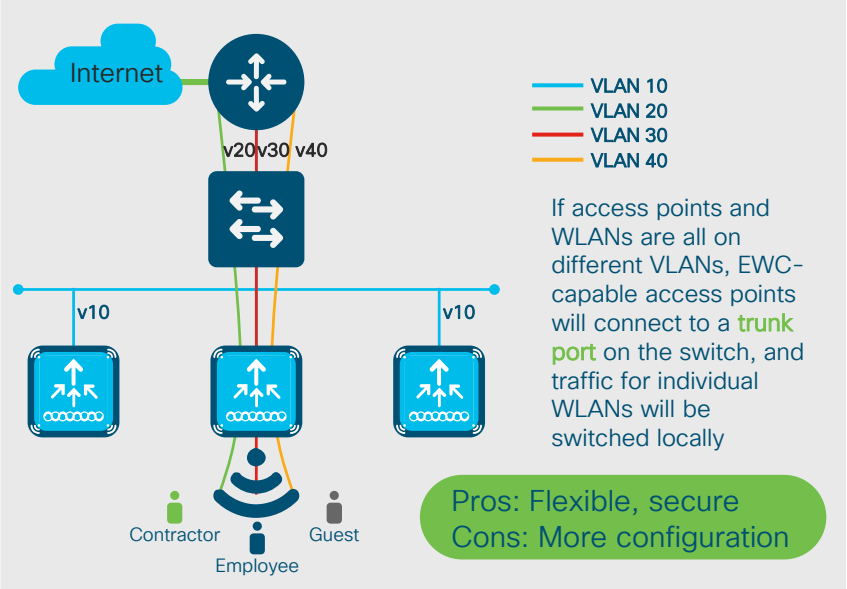

EWC on Catalyst access points: Resiliency

Always-on **network**
- APs continue to switch data traffic

Always-on **clients**
- Users and endpoints continue to stay connected

Always-on **services**
- Less than 10 seconds downtime of services

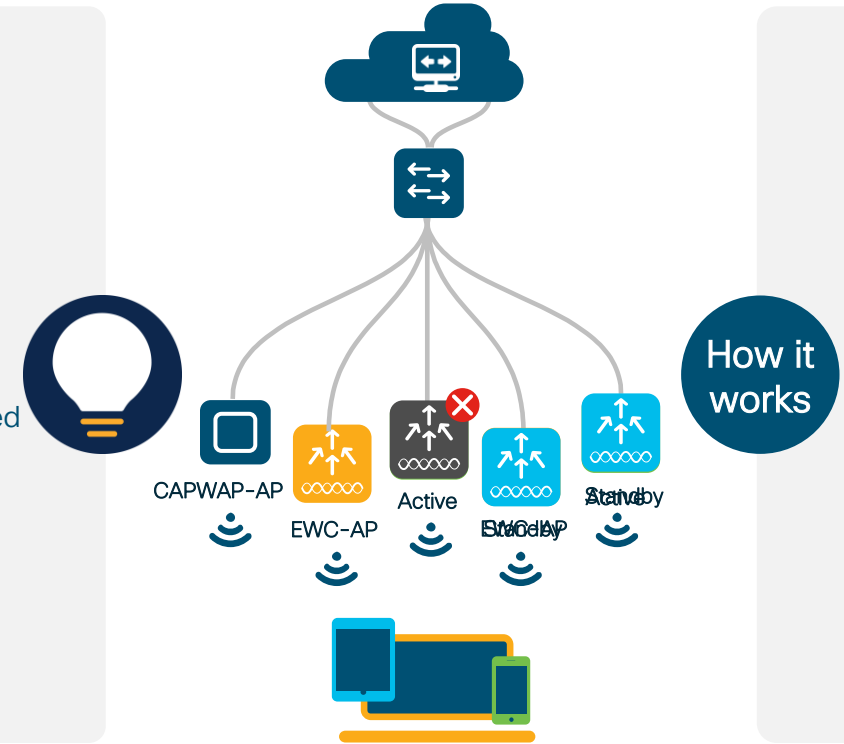

How it works

- Failure of active controller triggers a switchover to standby
- Standby controller is active in less than 10 seconds, and another EWC-AP is elected as a standby, providing redundancy
- APs fail over to the new controller

EWC and Mobility Express Summary

“Quick and Easy setup, no additional Hardware, WAN Independency”

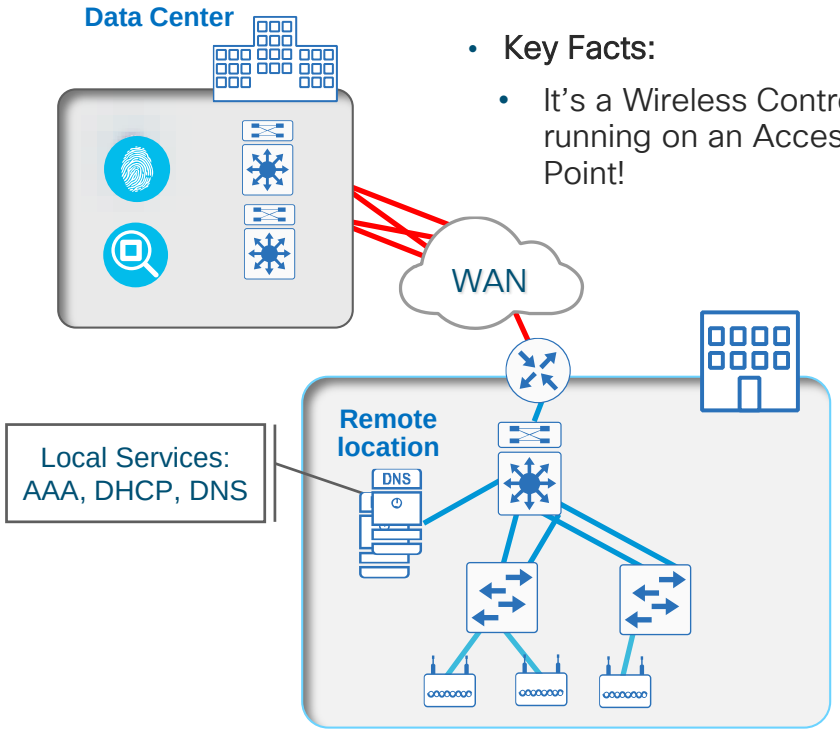

- Key Facts:
  - It’s a Wireless Controller running on an Access Point!

When to use:

- WAN independency is required and low hardware footprint is desired.
- Ideal for new deployments using 18xx/28xx/38xx Series Access Points or Catalyst Access Points

High Availability:

- “Self-Healing” redundancy
- Independent from WAN
- Local AAA, DHCP, DNS for full WAN independency

Keep in Mind:

- Switchport as Trunk if SSID/VLAN separation needed
- Per branch configuration and management
- consider adding Prime Infrastructure or Cisco DNA Center for central management

cisco Live!

HA Deployment Best Practices: Remote Site/Small Site Key Design Questions

Local Controller

Controller (Appliance/virtual)

- Specific per branch configuration
- Independency from WAN quality
- Reduced configuration on switches
- Full feature support
- L3 roaming supported

Mobility Express

- Specific per branch configuration
- Independency from WAN quality
- low hardware footprint (Controller running on Access Point)

FlexConnect

- Single pane of Mgmt. & Troubleshooting
- Reduced branch footprint
- Built-in resiliency
- Perfect fit for centralized IT Team

FlexConnect quick recap...

- CAPWAP management and data plane are split:
  - Central Switching (SSID data traffic sent to WLC)
  - Local Switching (SSID data traffic sent to local VLAN)
- Two modes of operation from AP perspective:
  - **Connected** (when WLC is reachable)
  - **Standalone** (when WLC is not reachable)

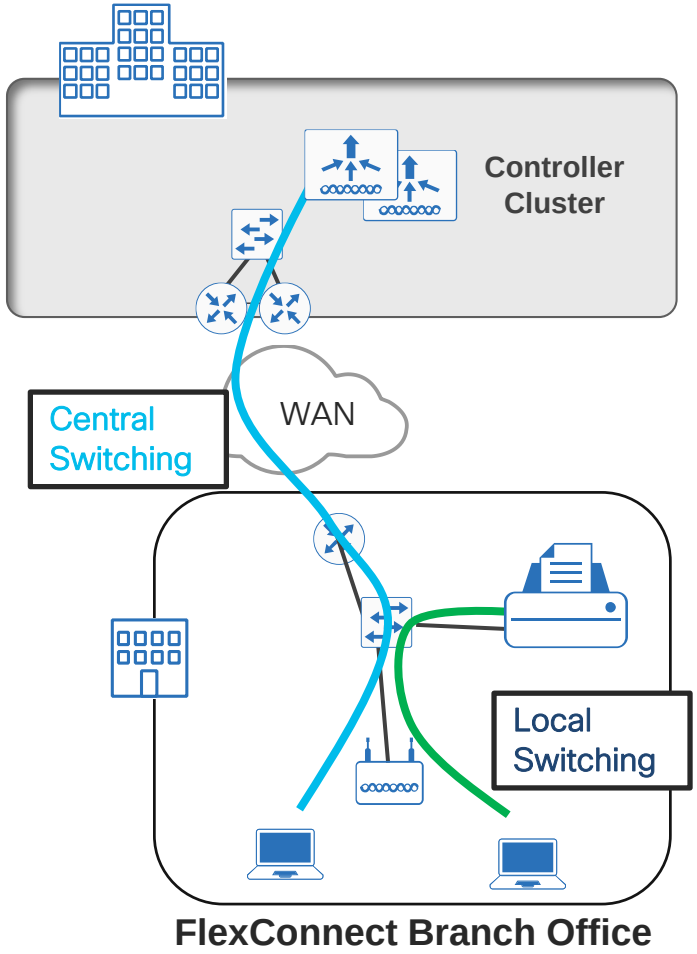

FlexConnect HA

	Limitations	Benefits
FlexConnect Local Switching	L2 roaming Flex Groups for AAA Local Auth. Fault Tolerance: identical configuration on N+1 controllers	Upon WLC failure AP stays up and clients are not disconnected Equivalent to Client SSO AAA survivability available
FlexConnect Central Switching	Same as Centralized mode	Same as Centralized mode

Clients at locally switched SSIDs stay connected at Controller/WAN outage

Local Switching SSIDs → all connected Clients stay connected!

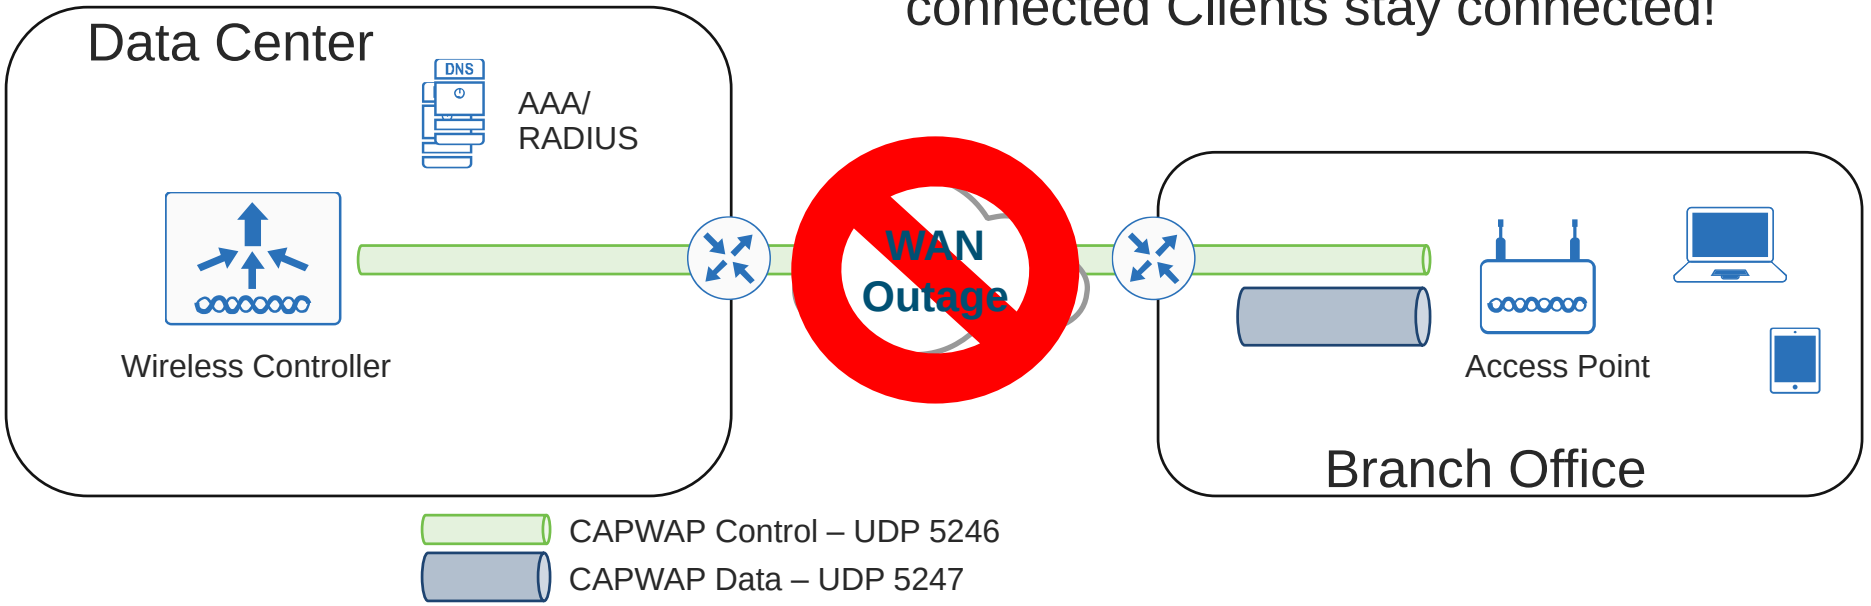

Impact of WAN Outage or Controller Failure

Controller failure :

- N+1 HA Design:
- No Impact for locally switched SSIDs
- FlexConnect AP will search for backup WLC and resume client sessions with centrally switched SSIDs.
- 1:1 HA Design with Client SSO:
- No impact for centrally switched SSIDs: Centrally and locally switched SSIDs stay up.

WAN Failure/ Controller not reachable:

- Access Point will continue to transmit/receive Data on locally switched SSIDs.
- Connected Clients stay connected
- Fast roaming is possible for Clients with CCKM/OKC/802.11r support
- New Clients can connect if local RADIUS or Authentication provided.
- Lost features: RRM, wIDS, location, WebAuth, NAC

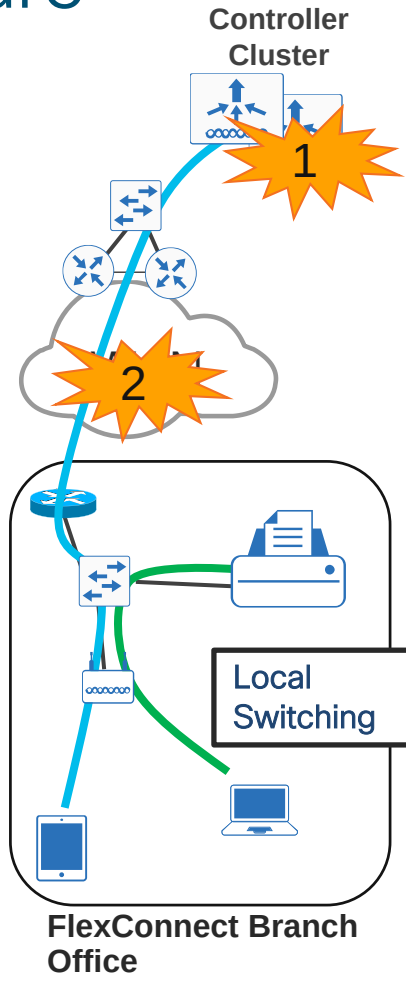

FlexConnect Summary

"Central Controller Cluster for thousands of Sites and Access Points"

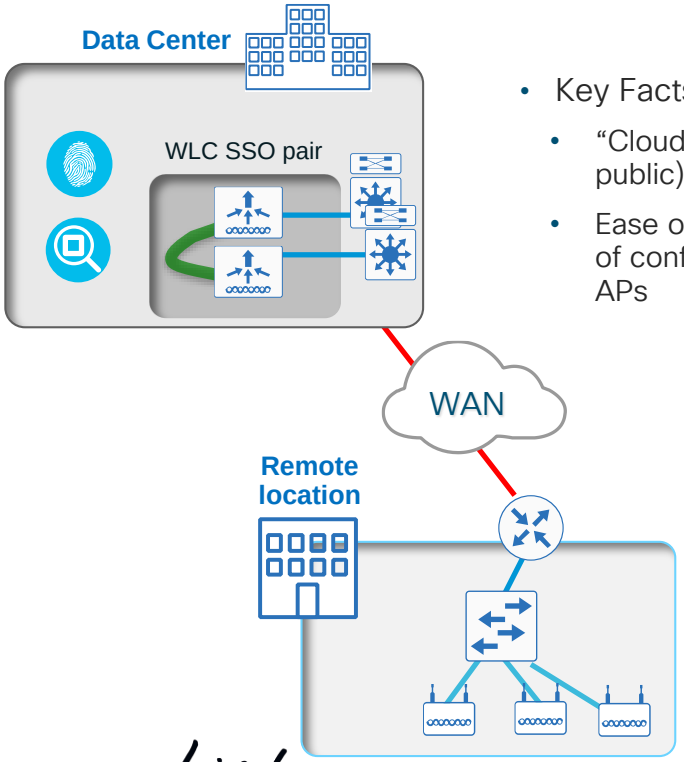

- Key Facts
  - "Cloud Controller" (private or public)
  - Ease of Operations: single point of configuration for up to 6000 APs

When to use:

- Perfect for centralized IT Team

High Availability:

- If controller not reachable:
- Local Data path stays UP and Clients stay connected, you can use AAA survivability
- SSO at central site provides control plane survivability

Keep in Mind:

- Switchport as Trunk if SSID/VLAN separation needed
- WAN Performance
- Some feature limitations (compared with local Controller)

Agenda

- What to do at the Radio Frequency layer?
- HA Design and Deployment Practices
  - Central/Large Site Deployments
  - Remote/Small Site Deployments
- **Wireless Controller Features for Planned Outages**
- Key takeaways

Next-generation Cisco Catalyst wireless access

Cisco Catalyst 9800 Series Wireless Controllers

Powered by Cisco IOS® XE
Open and programmable

Cisco Catalyst 9100 Access Points

Powered by Wi-Fi 6 technology
Superior RF experience

Resilient

- Deterministic capacity at scale
- Superior battery life for IoT and mobile devices
- Software updates with minimal disruption

Leadership in RF innovation

Secure

- Detect encrypted threats with Encrypted Traffic Analytics (ETA)
- Multi-lingual AP with RF snapshots
- WPA3, Trustworthy systems

Extending Cisco's intent-based network

Intelligent

- Enhanced analytics with Cisco DNA
- IOx infra support to host IOT applications
- Deploy in infrastructure of choice and cloud of choice

Innovation Beyond the Standard

Resilient: High Availability Summary

Reducing downtime for upgrades and unplanned events

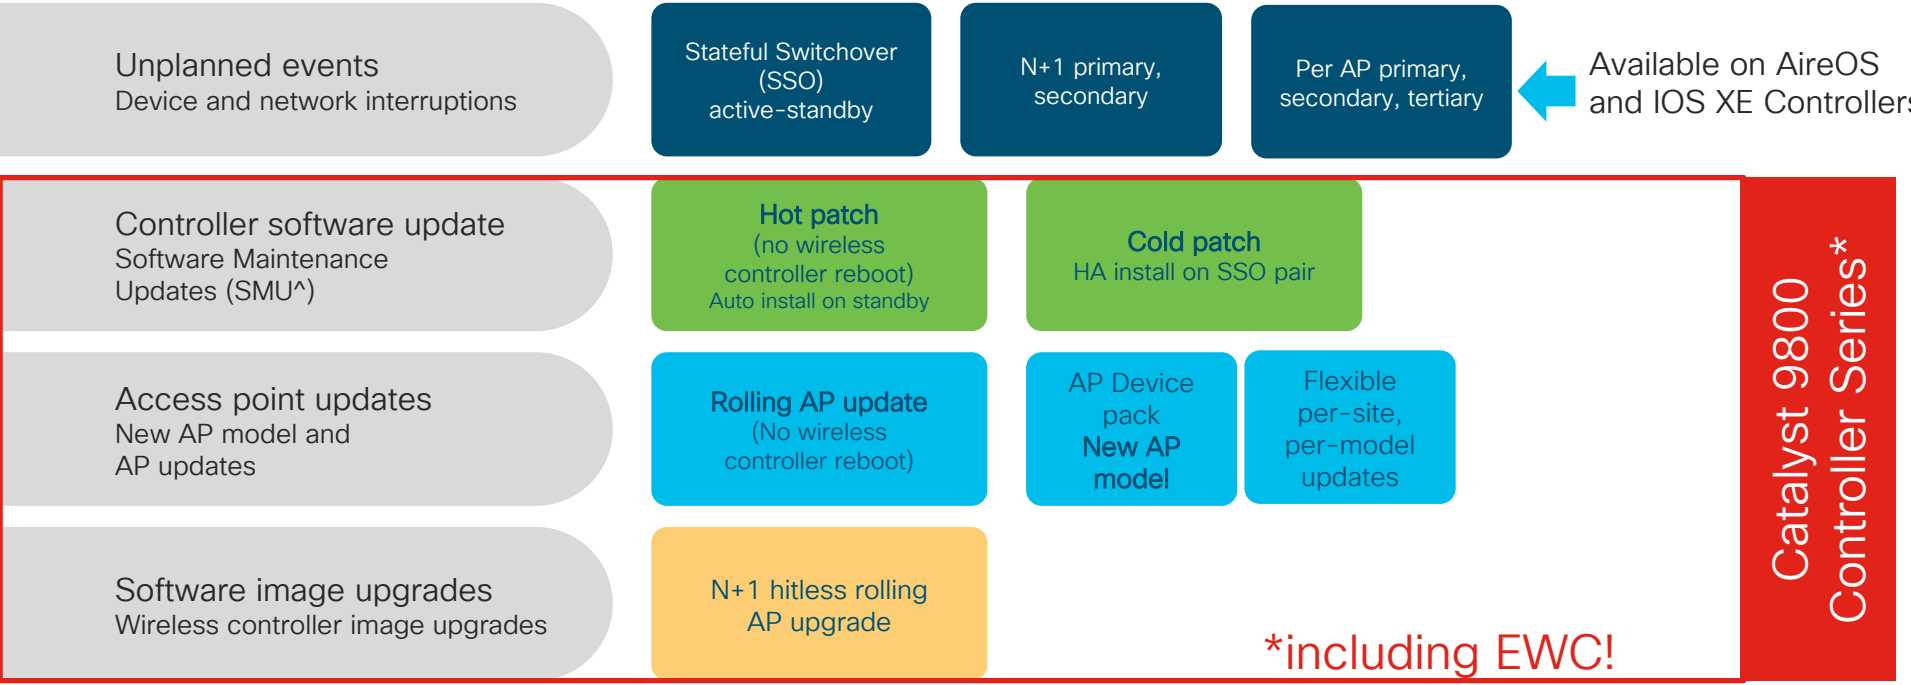

Resilient: Seamless software update infrastructure

Seamless SW Updates

Update (patch) controllers without client downtime. Update specific model APs with AP Service Pack

Flexible Per-Site Updates

AP Device Pack

Introduce new AP models in your network without any downtime and without impacting other APs

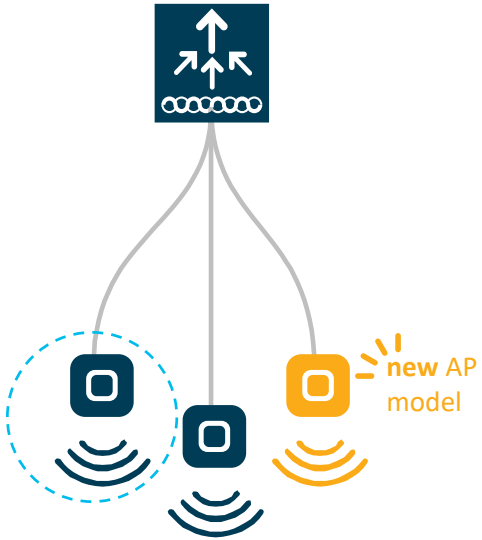

How it Works

- ✓ Install controller specific updates (patches) without client downtime to fix issues seamlessly
- ✓ Service updates for specific Access Point models without impacting other models
- ✓ New Access Points can join the controller with an AP device pack without impacting other APs

Resilient: N+1 Rolling AP Upgrade

Wireless Controller image upgrade using N+1 staging Controller

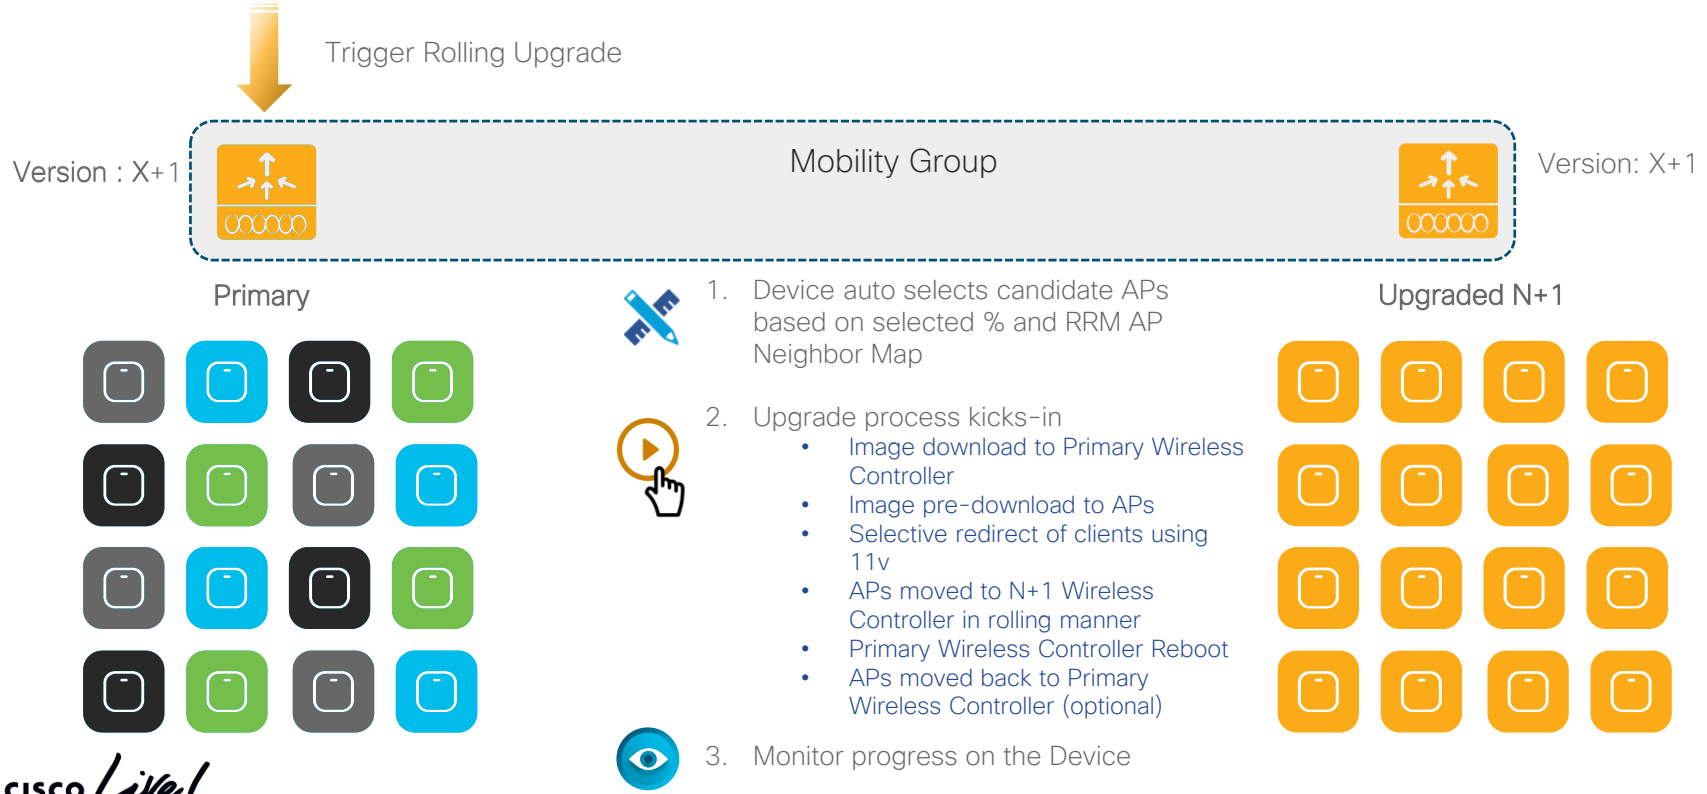

Example: Apply AP Service Pack (per AP Model/ Site) using Rolling AP Upgrade

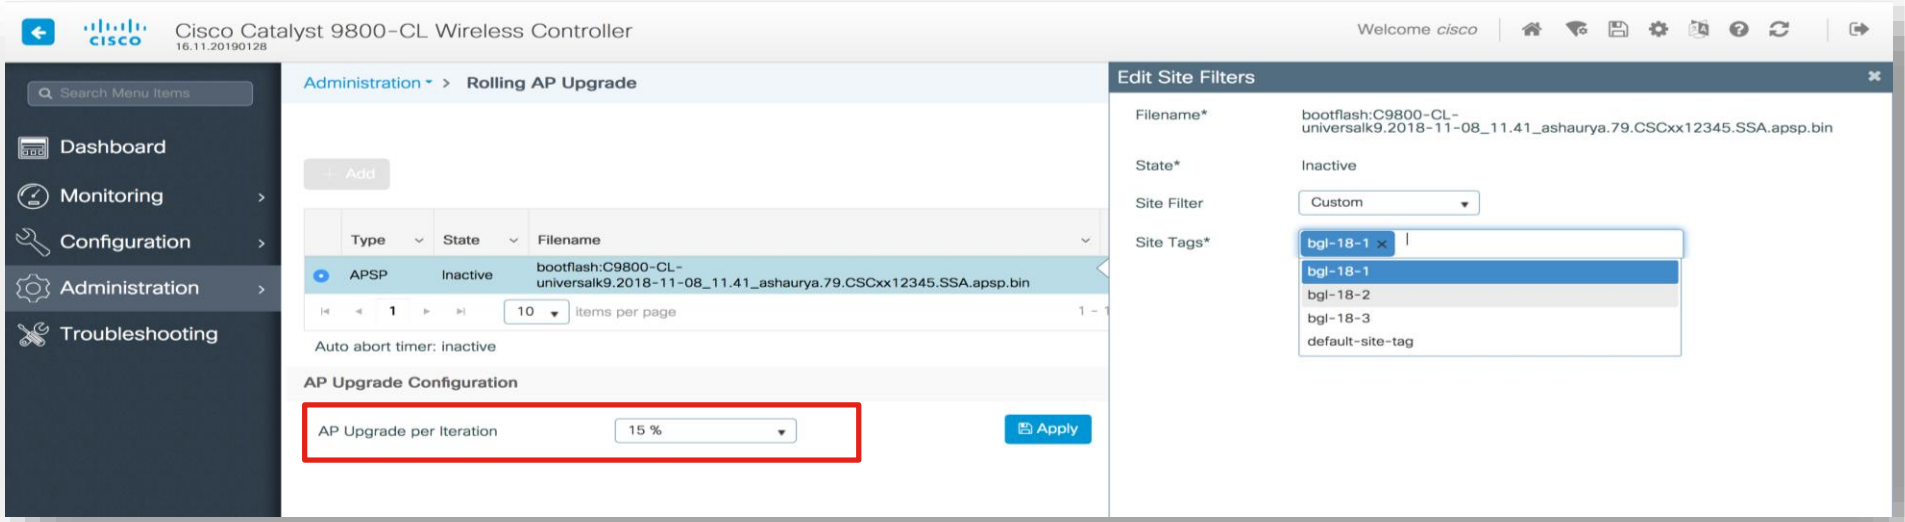

Rolling AP Upgrade – Client Steering

- Clients steered from candidate APs to non-candidate APs
- 802.11v BSS Transition Request
- Dissociation imminent
- If clients do not honor this, they will be de-authenticated before AP reload

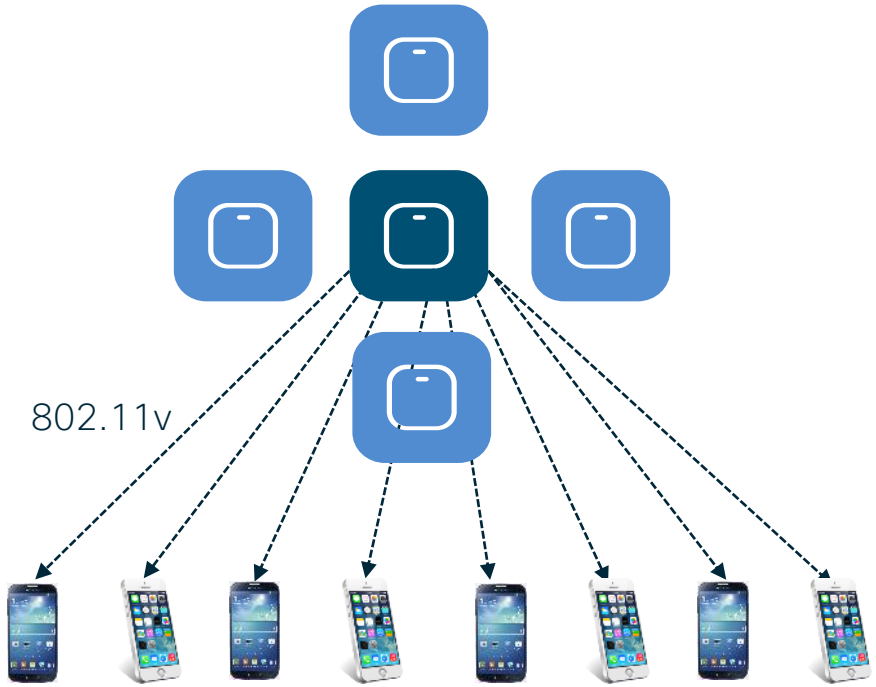

Summary of HA Options and Evolution

How long can my network be down?

Catalyst 9800 controller differentiation

	Controller Fault	Controller and AP s/w update	Image Upgrade
Standalone 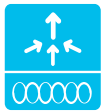	10s of minutes for AP and client recovery	Zero-downtime with SMU and APSP	Tens of minutes for AP and client recovery
N+1 HA 	Noticeable Outage to clients and APs	Zero-downtime with SMU and APSP	No Outage to APs and Client Automated Orchestration from Cisco DNA Center
SSO Pair 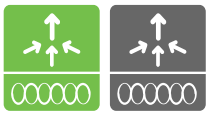	Sub-second AP and client recovery	Zero-downtime with SMU and APSP	Outage to APs and Client Need for extra WLC Manual orchestration

Summary of HA Options and Evolution

How long can my network be down?

Catalyst 9800 controller differentiation

	Controller Fault	Controller and AP s/w update	Image Upgrade
Standalone	10s of minutes for AP and client recovery	Zero-downtime with SMU and APSP	Tens of minutes for AP and client recovery
N+1 HA	Noticeable Outage to clients and APs	Zero-downtime with SMU and APSP	• No Outage to APs and Client • Automated Orchestration • from Cisco DNA Center
SSO Pair	Sub-second AP and client recovery	Zero-downtime with SMU and APSP	Outage to APs and Client Need for extra WLC Manual orchestration

Agenda

- What to do at the Radio Frequency layer?
- HA Design and Deployment Practices
  - Central/Large Site Deployments
  - Remote/Small Site Deployments
- Wireless Controller Features for Planned Outages
- Key takeaways

Key Takeaways

High Availability for Wireless is a multi level approach, starting from Level 1 (RF)

You have different solutions to chose based on the downtime that is acceptable for your business application

Cisco Controller Client SSO eliminates network downtime upon controller failure

Hot-Patches and Rolling AP Upgrades reduce/eliminate downtime for software updates/patches (Catalyst 9800 Controller only)

Agenda

- What is high availability?
- Campus network foundations and structured design
- Campus wired LAN design and high availability
- Campus wireless LAN design and high availability
- **Summary and conclusions**

Summary and conclusions

Design and deployment guidance available

https://cisco.com/go/cvd and https://cs.co/en-cvds

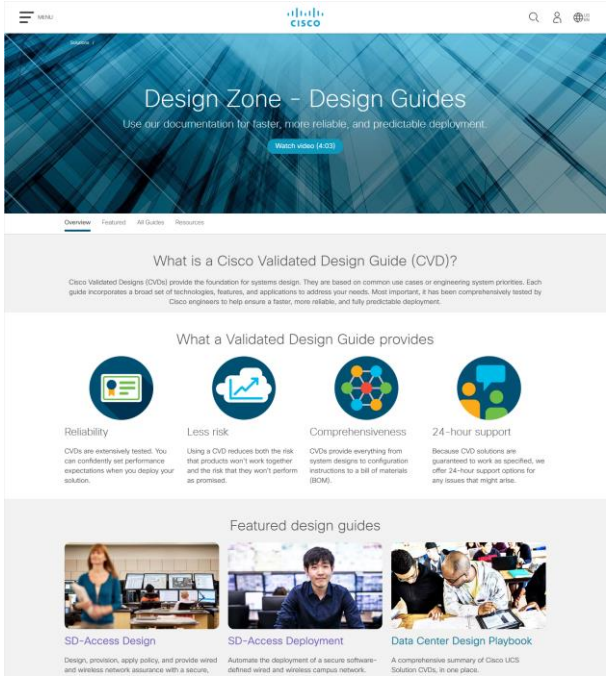

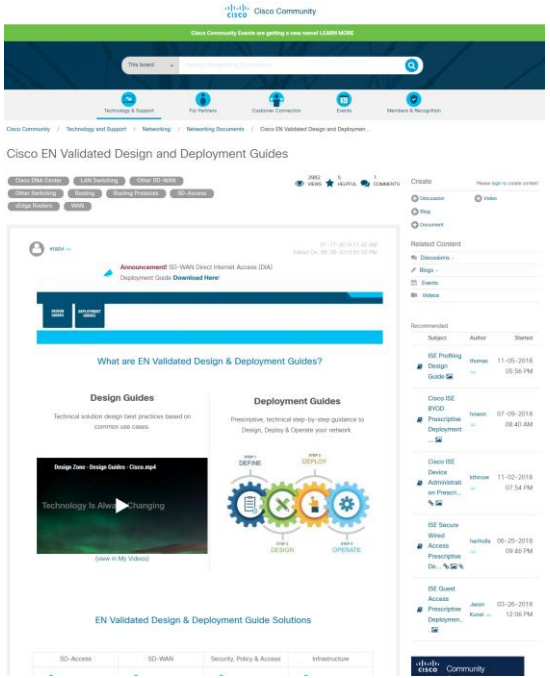

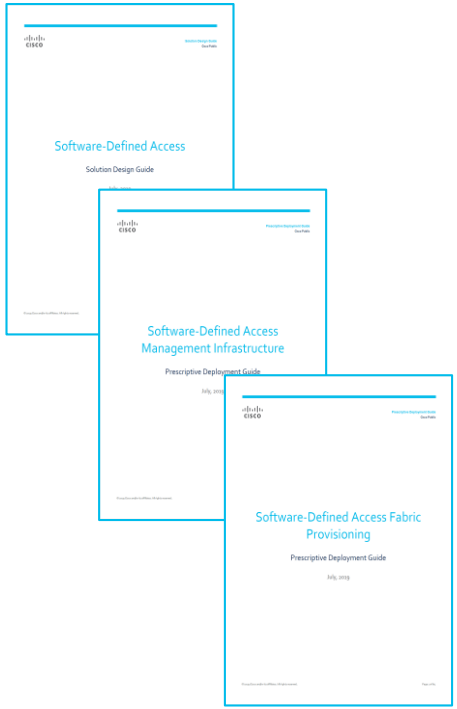

Reconvergence

Effect on “mission-critical”, real-time operations

- First step on the Moon – July 20, 1969 … how it really happened …

Reconvergence

Effect on “mission-critical”, real-time operations

- And how it would have looked with … standard HSRP timers …

Reconvergence

Effect on “mission-critical”, real-time operations

- And how it would have looked with … 500 millisecond reconvergence …

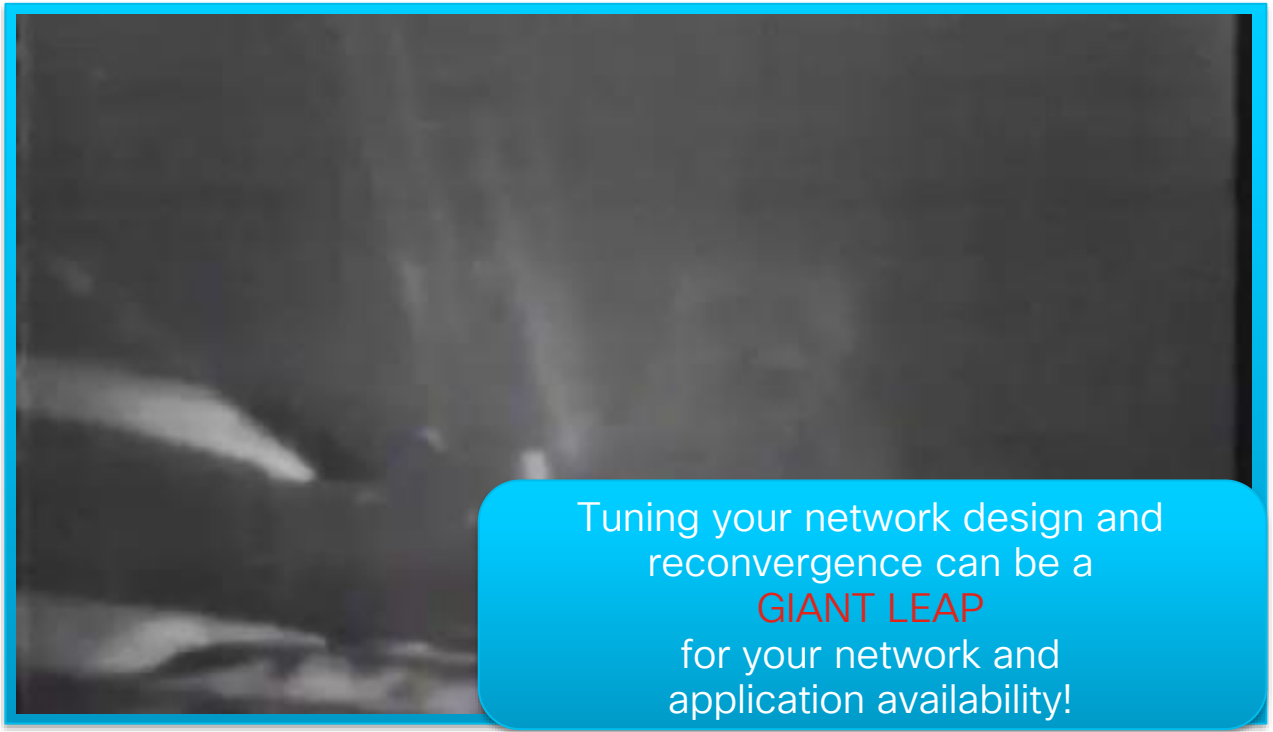

CL Barcelona 2020 Learning Maps

cisco Live!

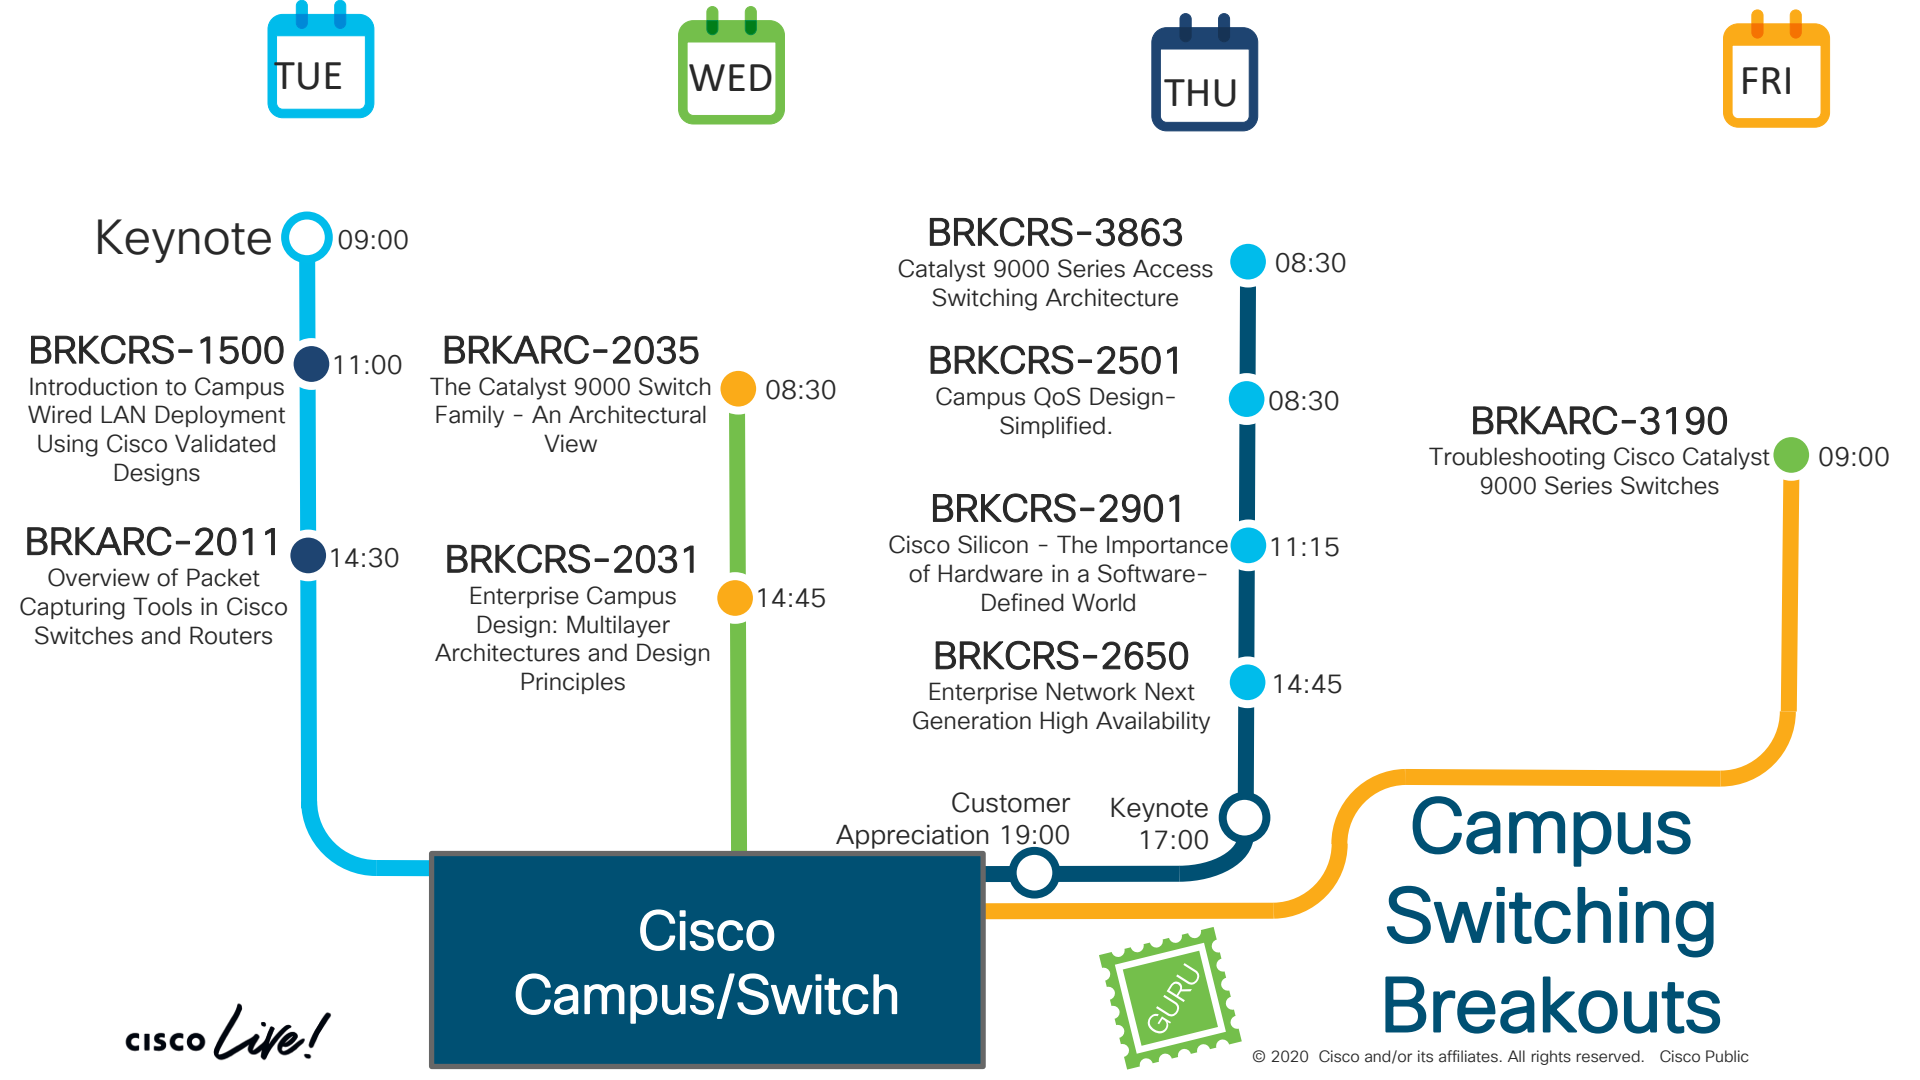

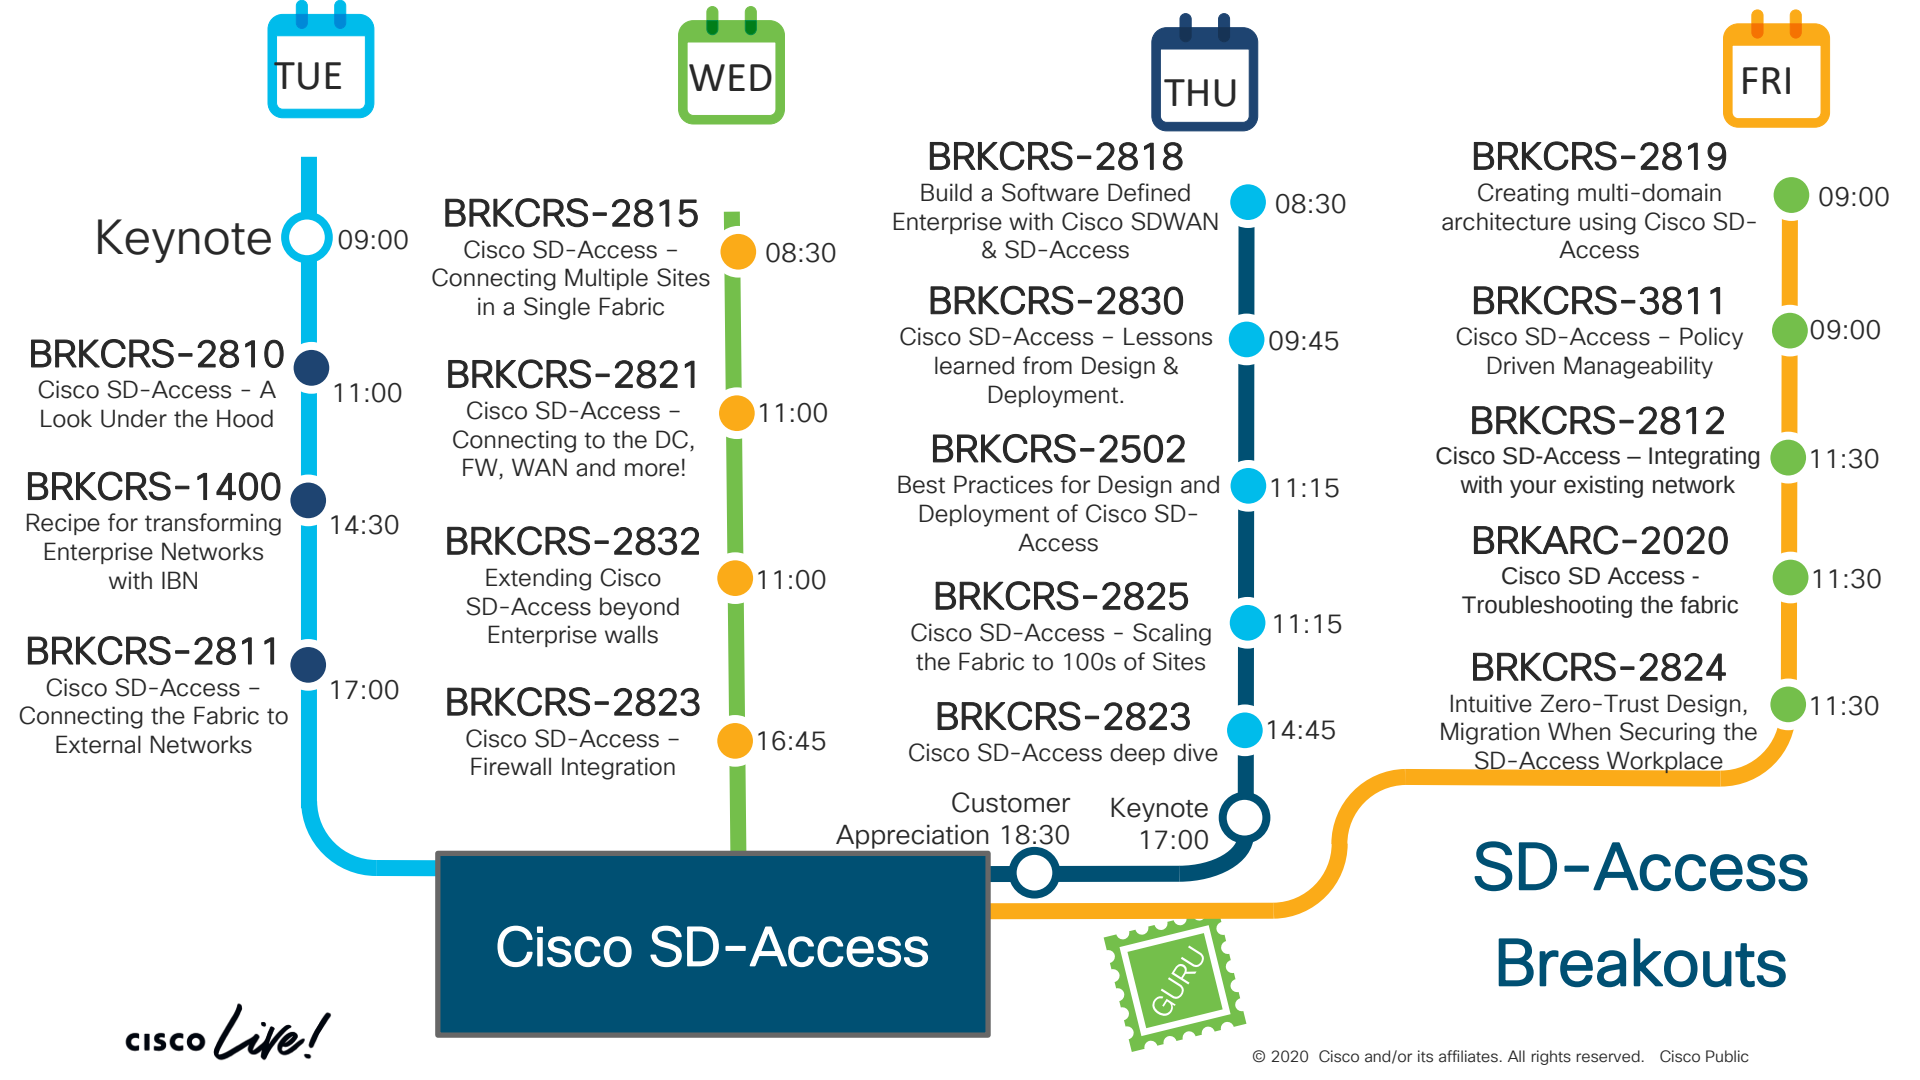

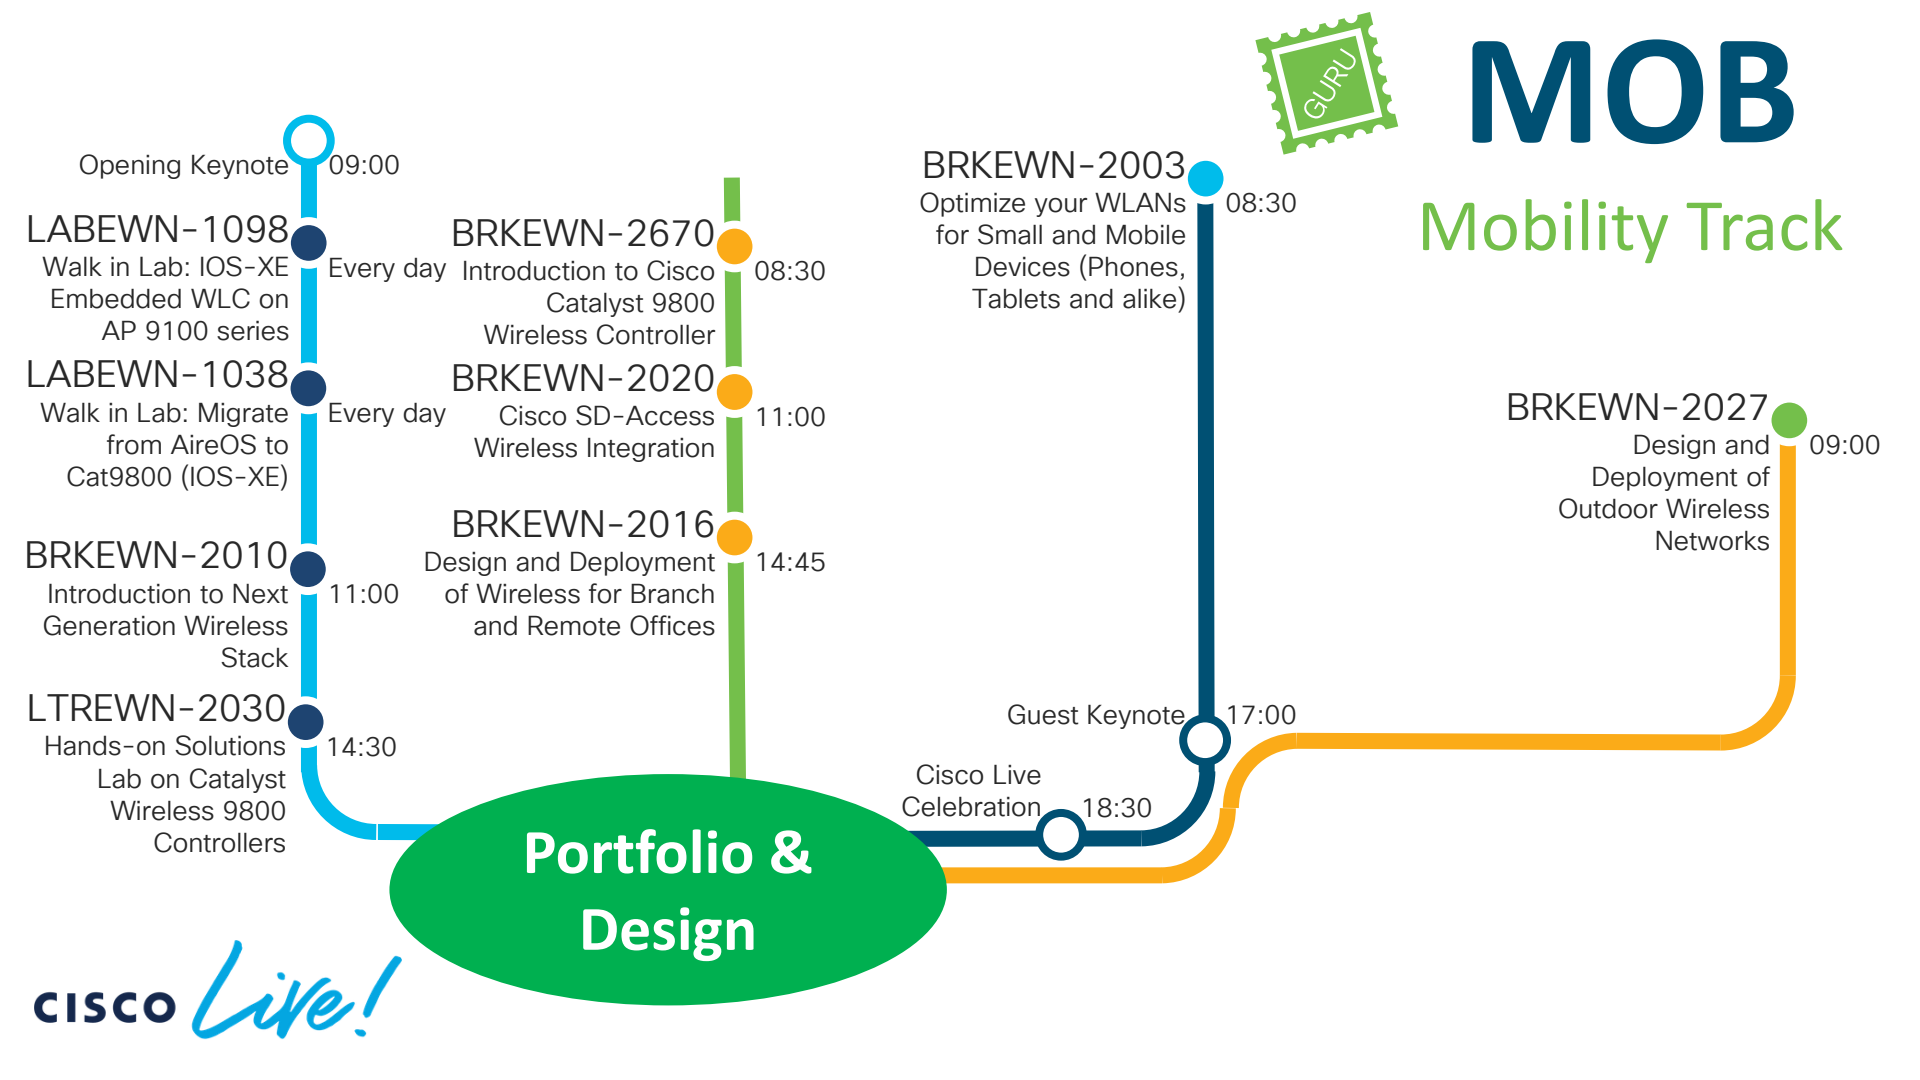
GURU
MOB
Mobility Track
Opening Keynote 09:00
LABEWN-1098
Walk in Lab: IOS-XE Embedded WLC on AP 9100 series
Every day
LABEWN-1038
Walk in Lab: Migrate from AireOS to Cat9800 (IOS-XE)
Every day
BRKEWN-2010
Introduction to Next Generation Wireless Stack
11:00
LTREWN-2030
Hands-on Solutions Lab on Catalyst Wireless 9800 Controllers
14:30
BRKEWN-2670
Introduction to Cisco Catalyst 9800 Wireless Controller
08:30
BRKEWN-2020
Cisco SD-Access Wireless Integration
11:00
BRKEWN-2016
Design and Deployment of Wireless for Branch and Remote Offices
14:45
BRKEWN-2003
Optimize your WLANs for Small and Mobile Devices (Phones, Tablets and alike)
08:30
BRKEWN-2027
Design and Deployment of Outdoor Wireless Networks
09:00
Guest Keynote 17:00
Cisco Live Celebration 18:30
Portfolio & Design
cisco Live!

MOB

Mobility Track

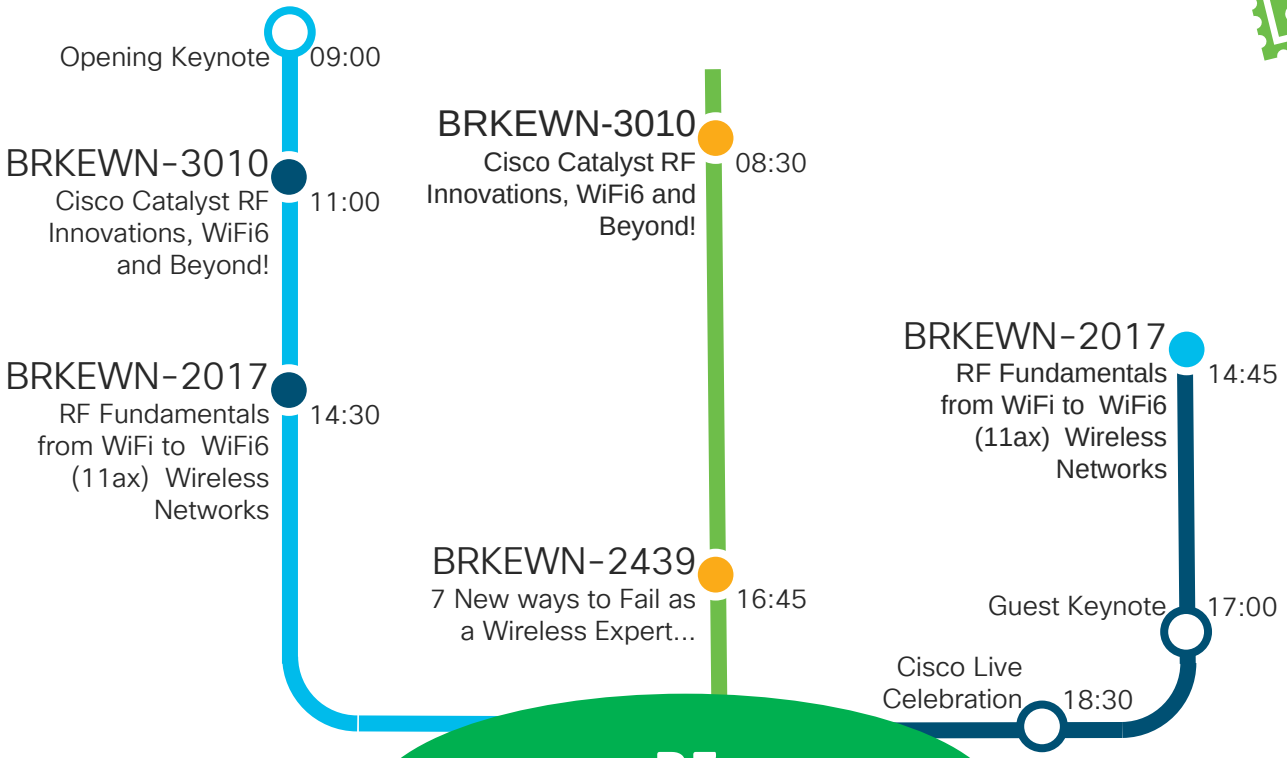

cisco Live!

OPS

Operations Track

www.ciscolive.com/emea/learn/technology-tracks/operations.html

Opening Keynote 09:00

LTRNMS-2500
Lab: A Practical Look at Cisco DNA Center 09:30

BRKOPS-2131
Cisco DNA Analytics and Assurance - The Shortest Path to Network Innocence 14:30

BRKOPS-2562
Data is the new Oil: The Nuts & Bolts of leveraging Cisco DNA Assurance data for creating value added services 17:00

LTRNMS-2043
Cisco DNA Assurance and Analytics Lab 09:00

BRKOPS-2991
Machine Learning in Network Operations: Lessons Learned 11:00

TCRNMS-2100
TechCircle: Cisco DNA Center Innovations 13:15

BRKOPS-2024
Wireless Automation & Assurance with Cisco DNA Center using APIs 16:45

BRKOPS-3825
Interpreting streaming telemetry data using ML/AI 11:15

BRKNMS-2031
Cisco DNA Center: The evolution from traditional Management to Intent-Based Automation & Assurance

BRKOPS-2100
Resolving Network Faults Faster through Automating Entire Fault Management Process. 14:45

Guest Keynote 17:00

Cisco Live Celebration 18:30

BRKSDN-2295
Controlling the wild wild west of applications in your network using Cisco DNAC QoS Policies 09:00

BRKOPS-2826
Cisco DNA Center Maintenance and Troubleshooting 11:30

DNA Assurance

GURU

cisco Live!

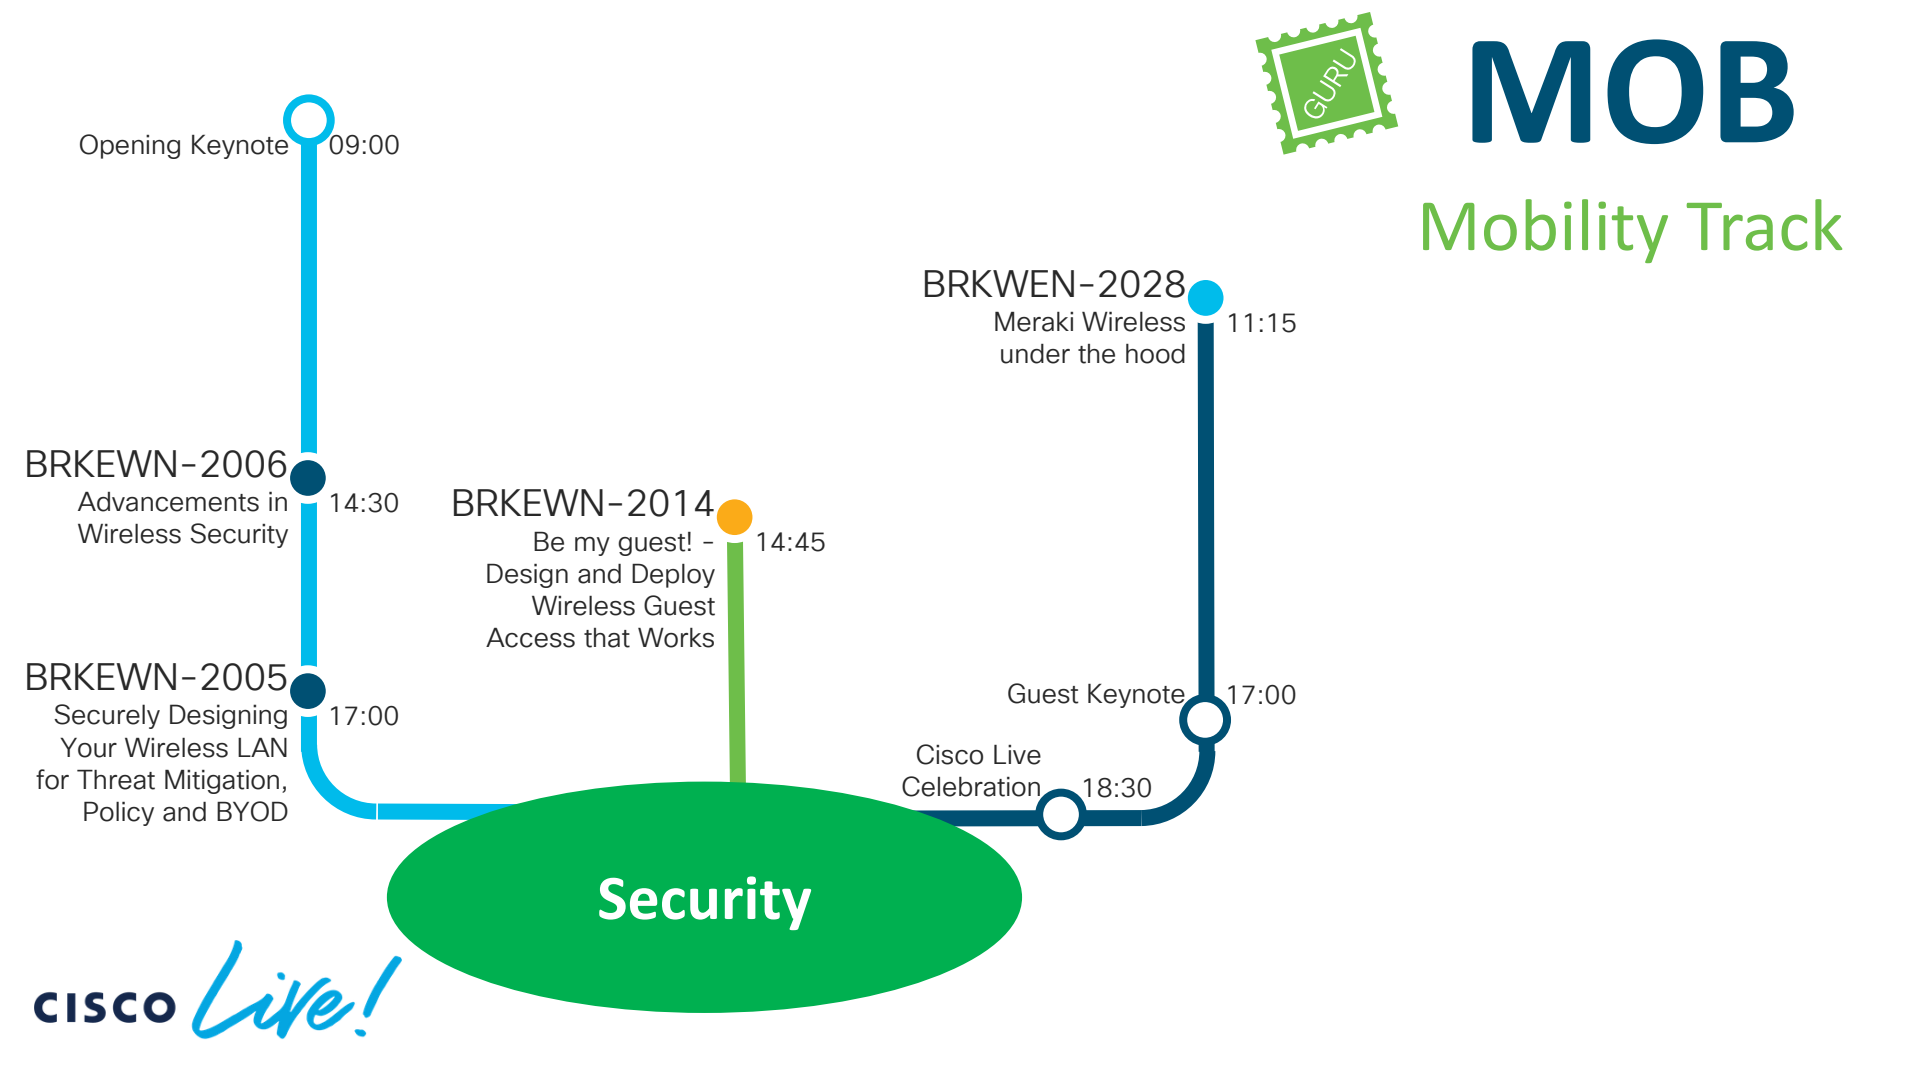
MOB
Mobility Track
Opening Keynote 09:00
BRKEWN-2006
Advancements in Wireless Security
14:30
BRKEWN-2005
Securely Designing Your Wireless LAN for Threat Mitigation, Policy and BYOD
17:00
BRKEWN-2014
Be my guest! - Design and Deploy Wireless Guest Access that Works
14:45
BRKWEN-2028
Meraki Wireless under the hood
11:15
Guest Keynote 17:00
Cisco Live Celebration 18:30
Security

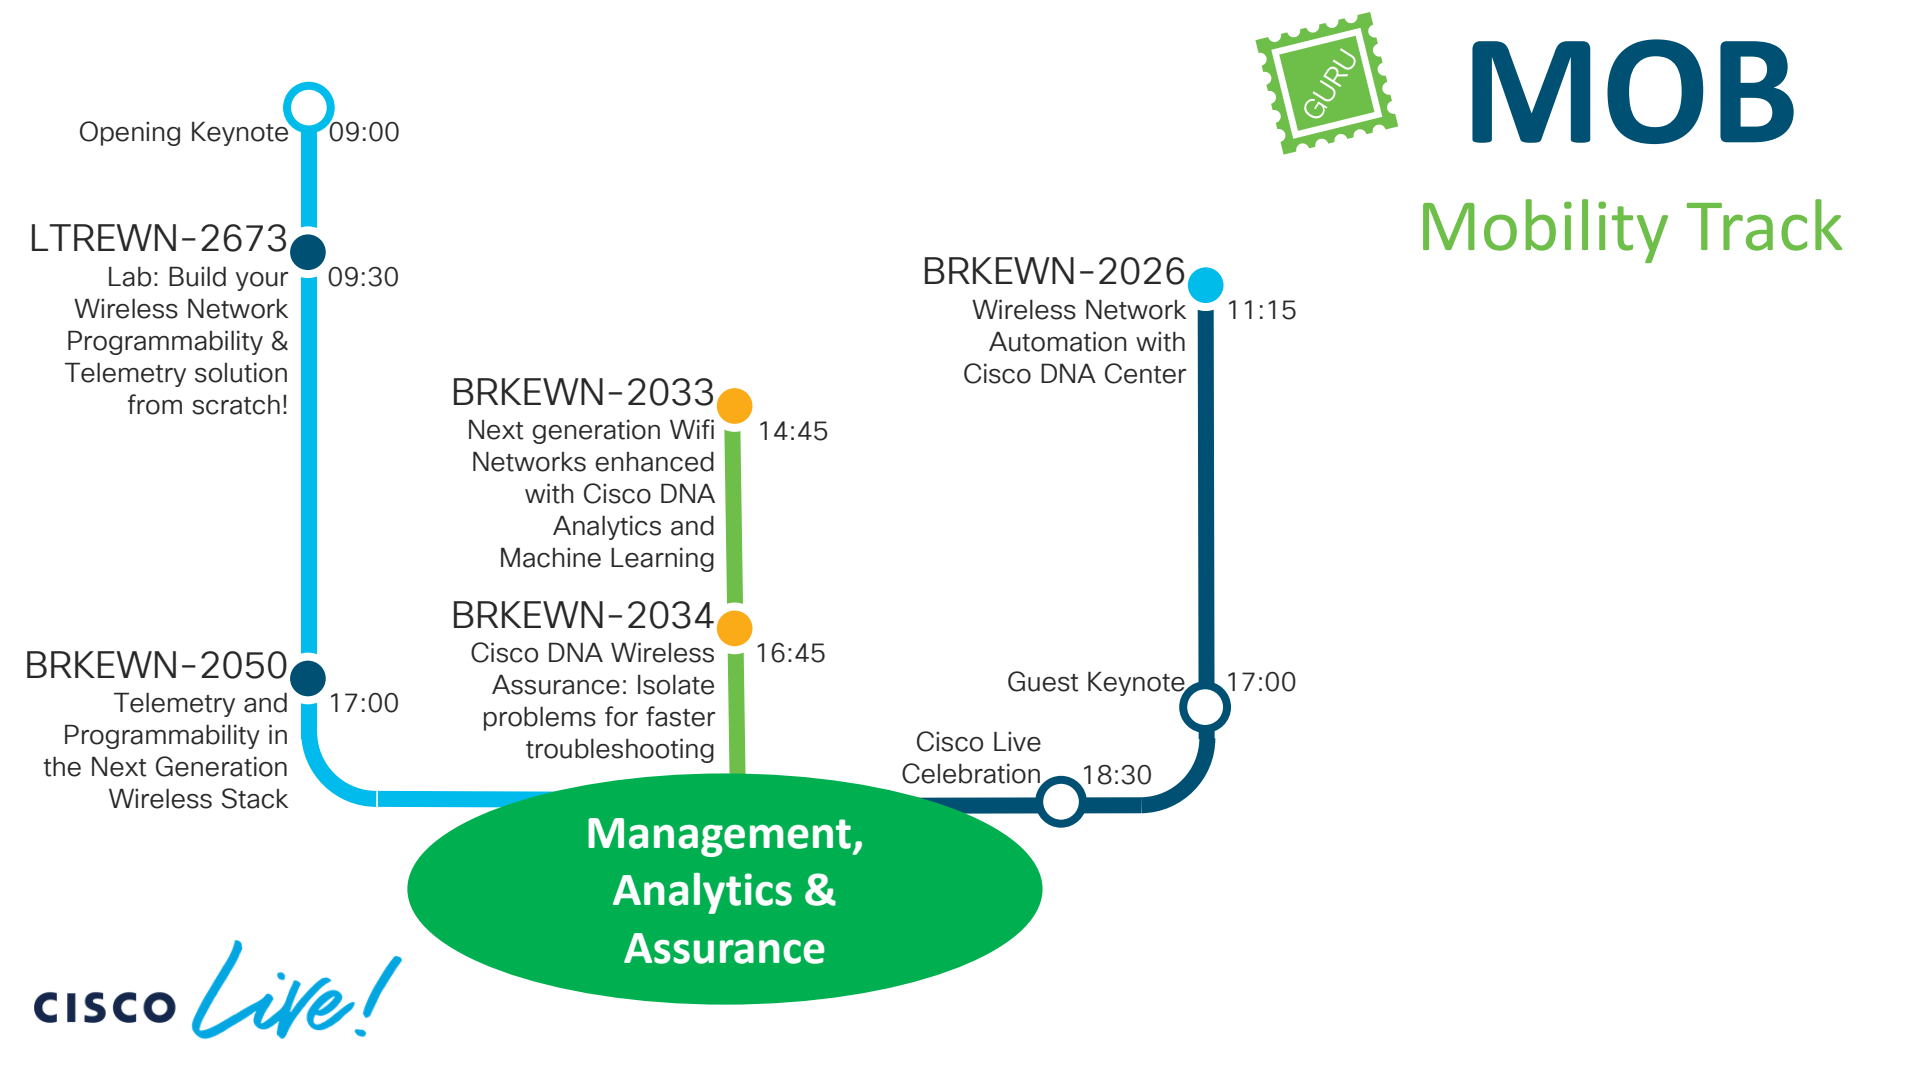
MOB
Mobility Track
Opening Keynote 09:00
LTREWN-2673
Lab: Build your Wireless Network Programmability & Telemetry solution from scratch!
09:30
BRKEWN-2050
Telemetry and Programmability in the Next Generation Wireless Stack
17:00
BRKEWN-2033
Next generation Wifi Networks enhanced with Cisco DNA Analytics and Machine Learning
14:45
BRKEWN-2034
Cisco DNA Wireless Assurance: Isolate problems for faster troubleshooting
16:45
BRKEWN-2026
Wireless Network Automation with Cisco DNA Center
11:15
Guest Keynote 17:00
Cisco Live Celebration 18:30
Management, Analytics & Assurance
cisco Live!

MOB

Mobility Track

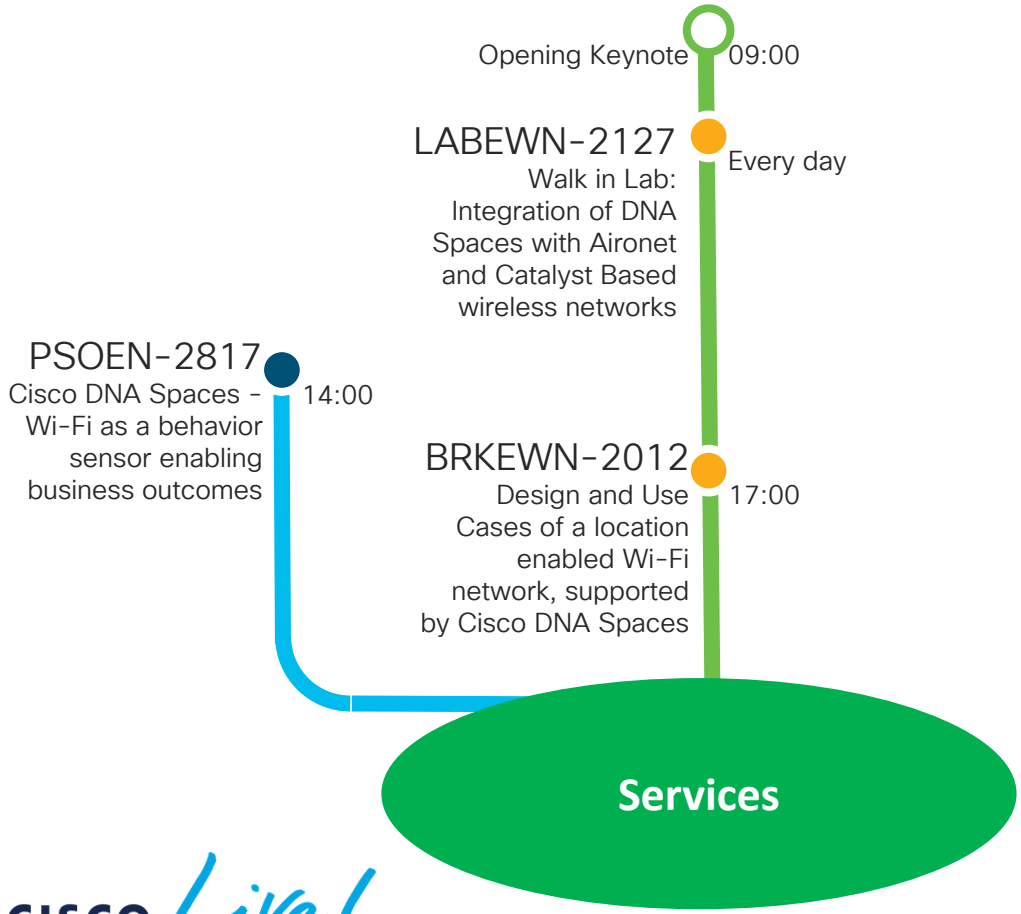

cisco Live!

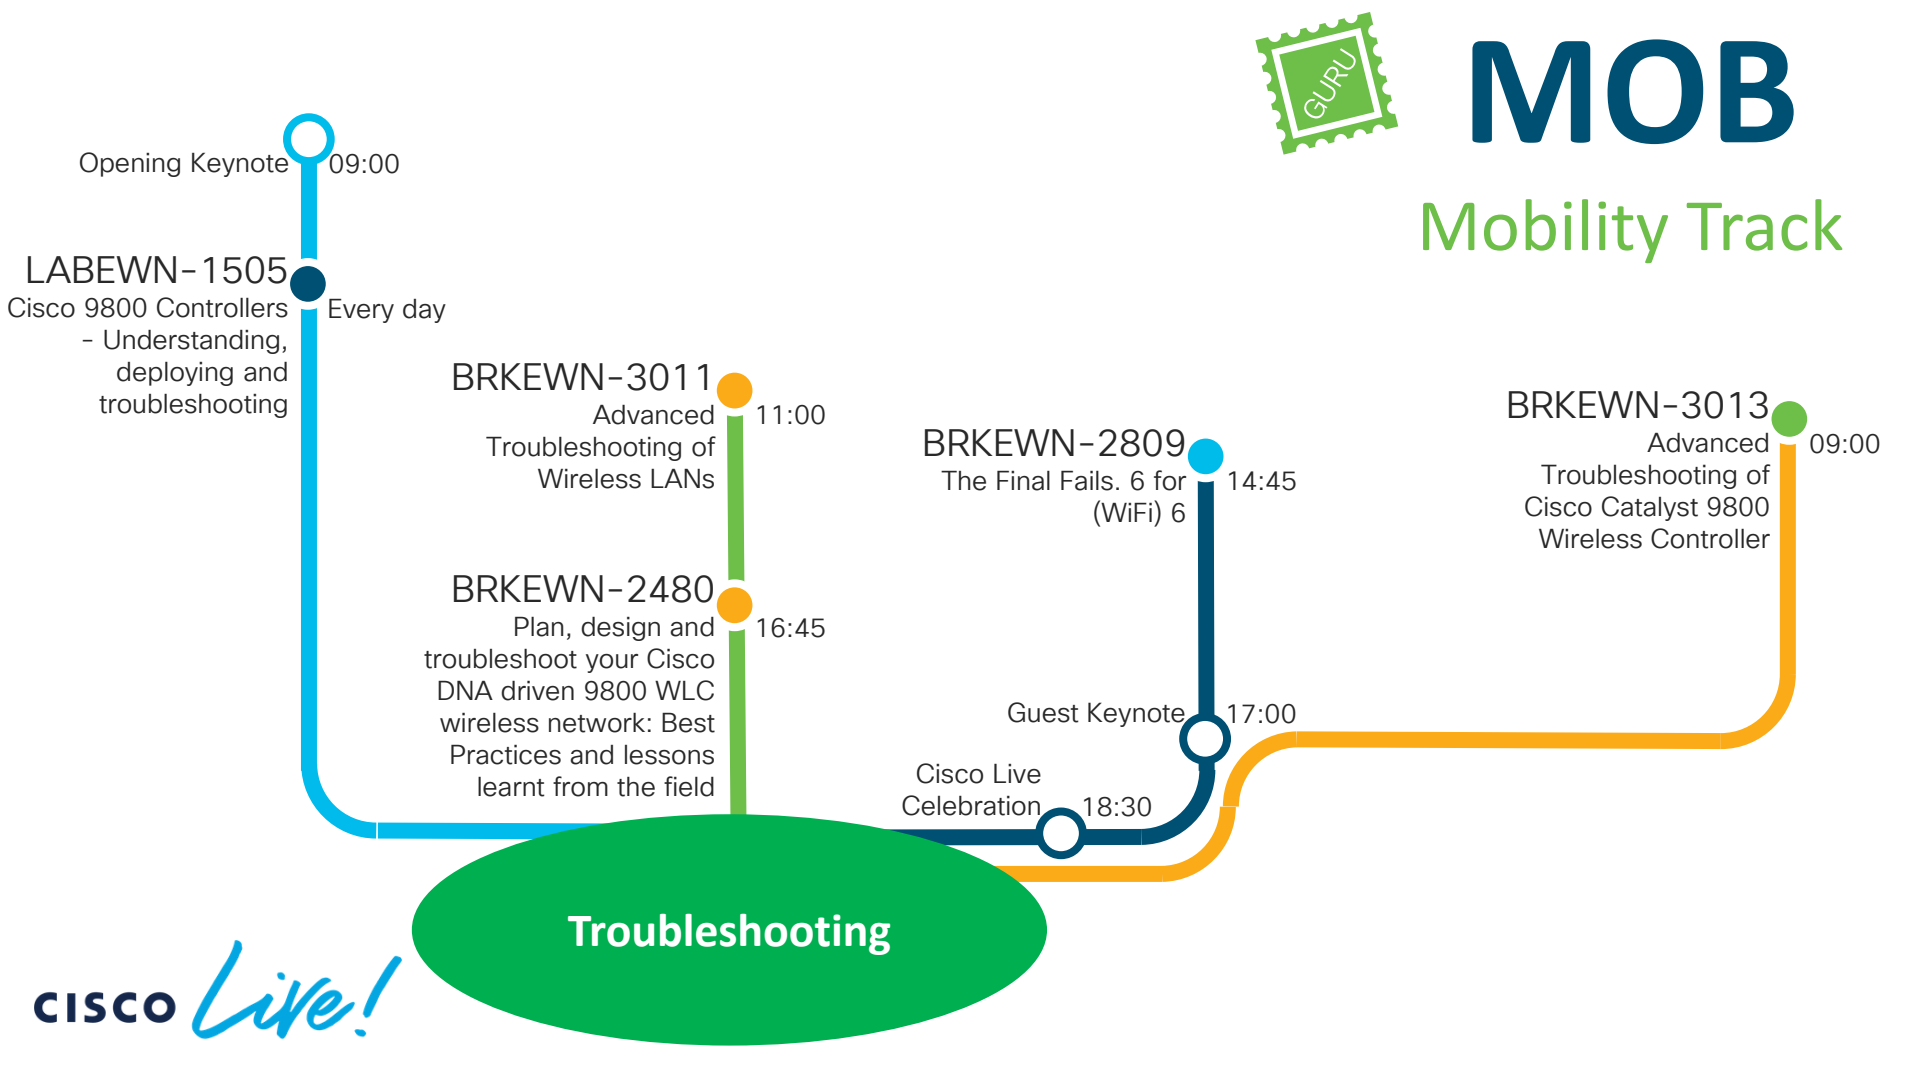
GURU
MOB
Mobility Track
Opening Keynote
09:00
LABEWN-1505
Cisco 9800 Controllers - Understanding, deploying and troubleshooting
Every day
BRKEWN-3011
Advanced Troubleshooting of Wireless LANs
11:00
BRKEWN-2480
Plan, design and troubleshoot your Cisco DNA driven 9800 WLC wireless network: Best Practices and lessons learnt from the field
16:45
BRKEWN-2809
The Final Fails. 6 for (WiFi) 6
14:45
Guest Keynote
17:00
Cisco Live Celebration
18:30
BRKEWN-3013
Advanced Troubleshooting of Cisco Catalyst 9800 Wireless Controller
09:00
Troubleshooting
cisco Live!

OPS

Operations Track

www.ciscolive.com/emea/learn/technology-tracks/operations.html

Opening Keynote 09:00

LTRNMS-2500
Lab: A Practical Look at Cisco DNA Center 09:30

BRKNMS-2285
How to be a hero with Cisco DNA Center Platform APIs 14:30

BRKSDN-2497
Build Your API-Based NW Troubleshooting Kit 17:00

BRKNMS-2426
Cisco DNA Center - From 0 to 100 How to get the network up and running from scratch 08:30

PSOOPS-2236
Unlocking the power of open platform with Cisco DNA Center Platform 11:00

TCRNMS-2100
TechCircle: Cisco DNA Center Innovations 13:15

BRKOPS-2150
Deploying Advanced Network Services using Cisco DNA Center 14:45

BRKOPS-2024
Wireless Automation & Assurance with Cisco DNA Center using APIs 16:45

BRKNMS-2031
Cisco DNA Center: The evolution from traditional Management to Intent-Based Automation & Assurance 11:15

Guest Keynote 17:00

Cisco Live Celebration 18:30

BRKSDN-2295
Controlling the wild wild west of applications in your network using Cisco DNAC QoS Policies 09:00

BRKOPS-2826
Cisco DNA Center Maintenance and Troubleshooting 11:30

DNA Automation

GURU

cisco Live!

OPS

Operations Track

www.ciscolive.com/emea/learn/technology-tracks/operations.html

Opening Keynote 09:00

BRKNMS-2573
From Prime Infrastructure to Software Defined Network (SDN) Management with Cisco DNA Center 11:00

BRKOPS-2131
Cisco DNA Analytics and Assurance - The Shortest Path to Network Innocence 14:30

BRKNMS-2426
Cisco DNA Center - From 0 to 100 How to get the network up and running from scratch 08:30

BRKOPS-2110
End-2-end policy from the Campus to the DC and back, a packet journey with SDA to ACI 11:00

TCRNMS-2100
TechCircle: Cisco DNA Center Innovations 13:15

BRKSDN-2500
Real World Use Cases for Deploying and Operating Cisco SD-Access Using Cisco DNA Center 14:45

BRKNMS-2031
Cisco DNA Center: The evolution from traditional Management to Intent-Based Automation & Assurance 11:15

Guest Keynote 17:00

Cisco Live Celebration 18:30

BRKSDN-2295
Controlling the wild wild west of applications in your network using Cisco DNAC QoS Policies 09:00

BRKOPS-2859
Towards operating a multi-domain network 11:30

Operating Cisco SDA

GURU

cisco Live!

Complete your online session survey

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on 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 Content Catalog on ciscolive.com/emea.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.com.

Continue your education

Demos in the Cisco campus

Walk-in labs

Meet the engineer 1:1 meetings

Related sessions

Thank you

cisco Live!

You make **possible**