



The bridge to possible

Designing Highly Available Networks Using Catalyst 9000 Switches

Minhaj Uddin, Leader Technical Marketing

Cisco Webex App


Questions?

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

How

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

Webex spaces will be moderated until February 24, 2023.



Agenda

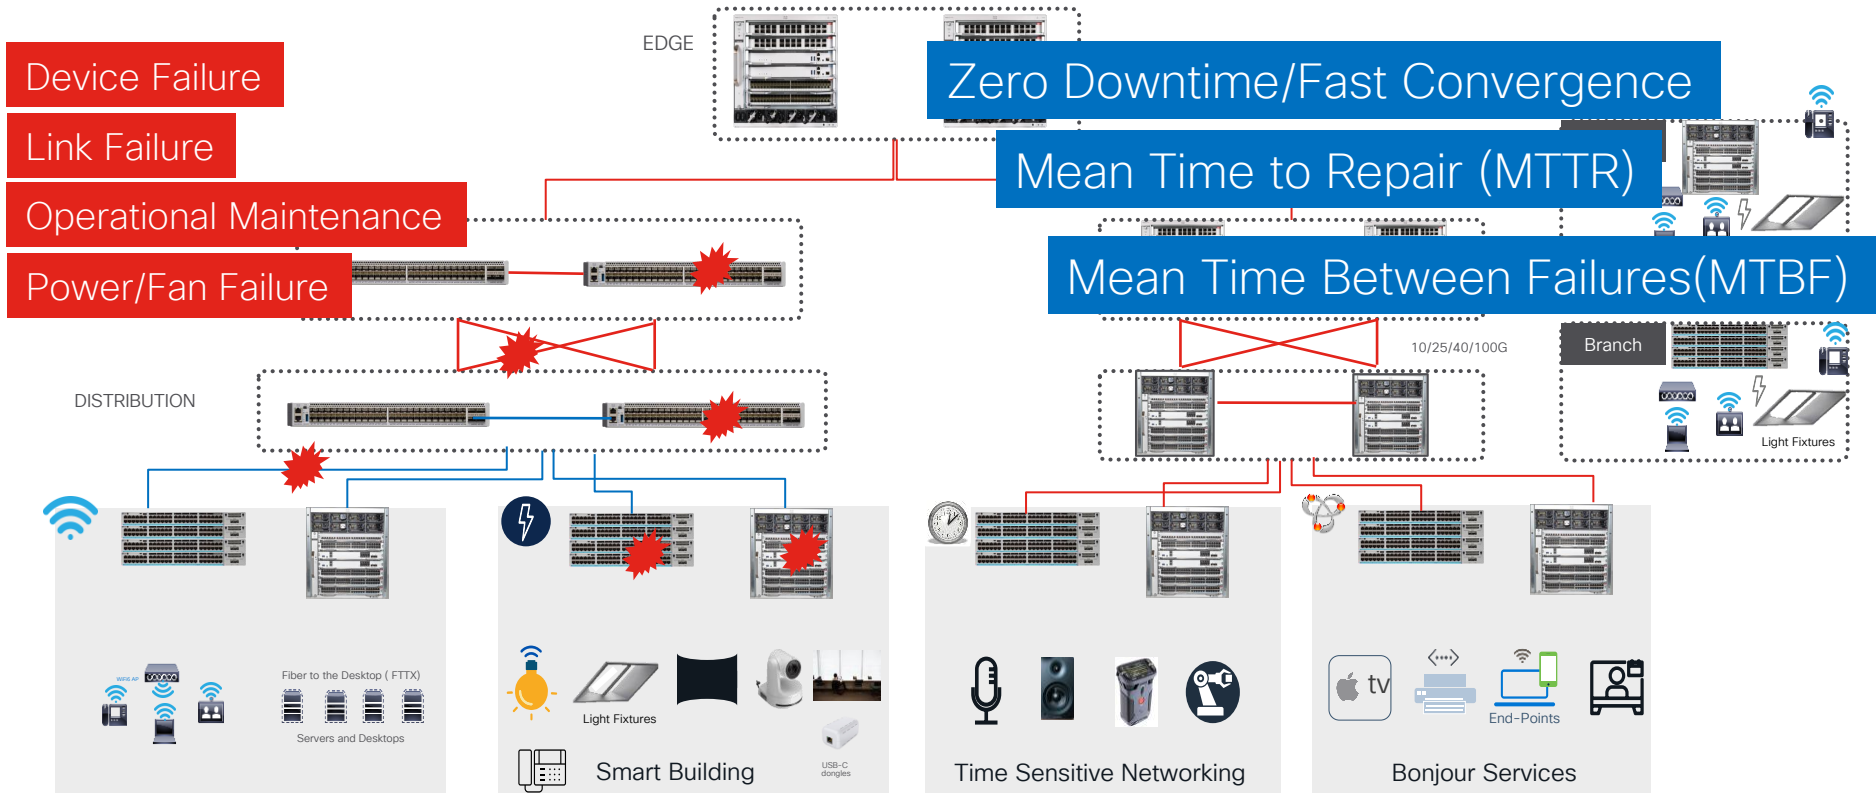
- High Availability Overview and Evolution
- High Availability Architecture and Designs
- High Availability Solution on the Campus Access
 - Stackable High Availability Solution
 - Modular High Availability Solution
- High Availability Solution on the Campus Distribution/Core
- Summary/Q&A

Goals

- Efficiently utilize available bandwidth
- Dynamically respond to all types of disruptions
- Leverage most effective design techniques that meet the design requirements



Typical Campus Network Architecture



Typical Campus Network Architecture

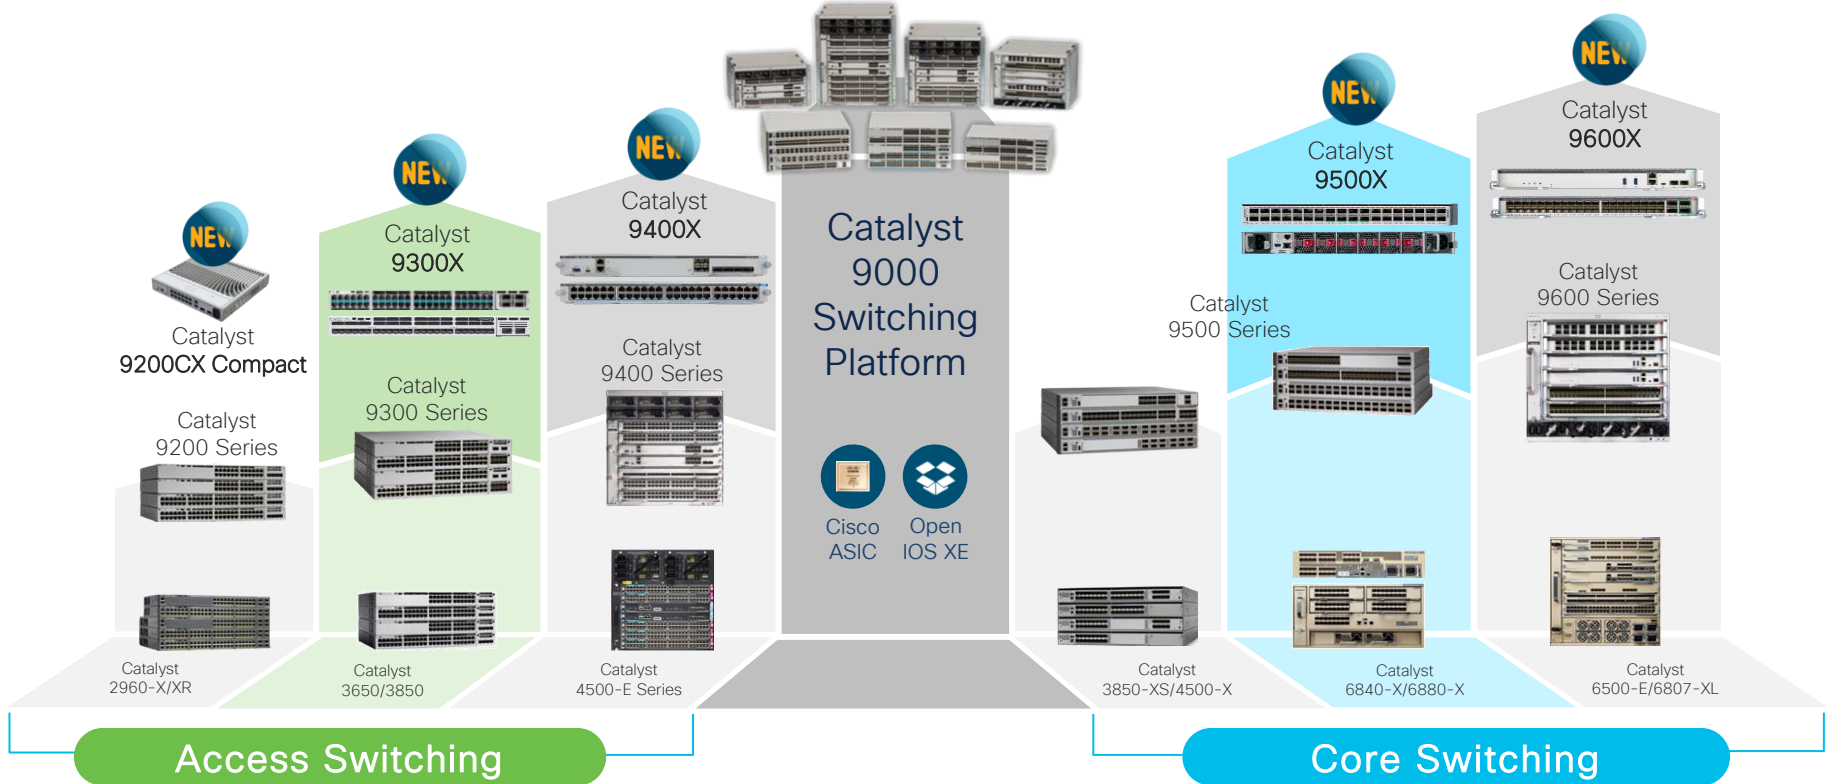
- Add more links
- Add more devices
- Leverage FHRP like HSRP and VRRP
- Change the timers
- Tune the application performance
- Etc...

Convergence Time?
Failover Detection?

What is the best way?

Cisco Catalyst 9000 Switching Portfolio

Adding the “X factor” to the industry’s leading switching family



Catalyst 9000: Pinnacle of Resiliency




Platform Resiliency

- **Redundant Switches and Supervisors** for Standalone and Modular Switches with Non-Stop Forwarding and Routing (NSF/NSR)
- **Redundant Fan & Power Supply** in case of any hardware failure



Design Resiliency

- **StackWise Virtual:** Redundant System for high availability, simplified design and configuration
- **GIR:** No downtime when device removed for maintenance

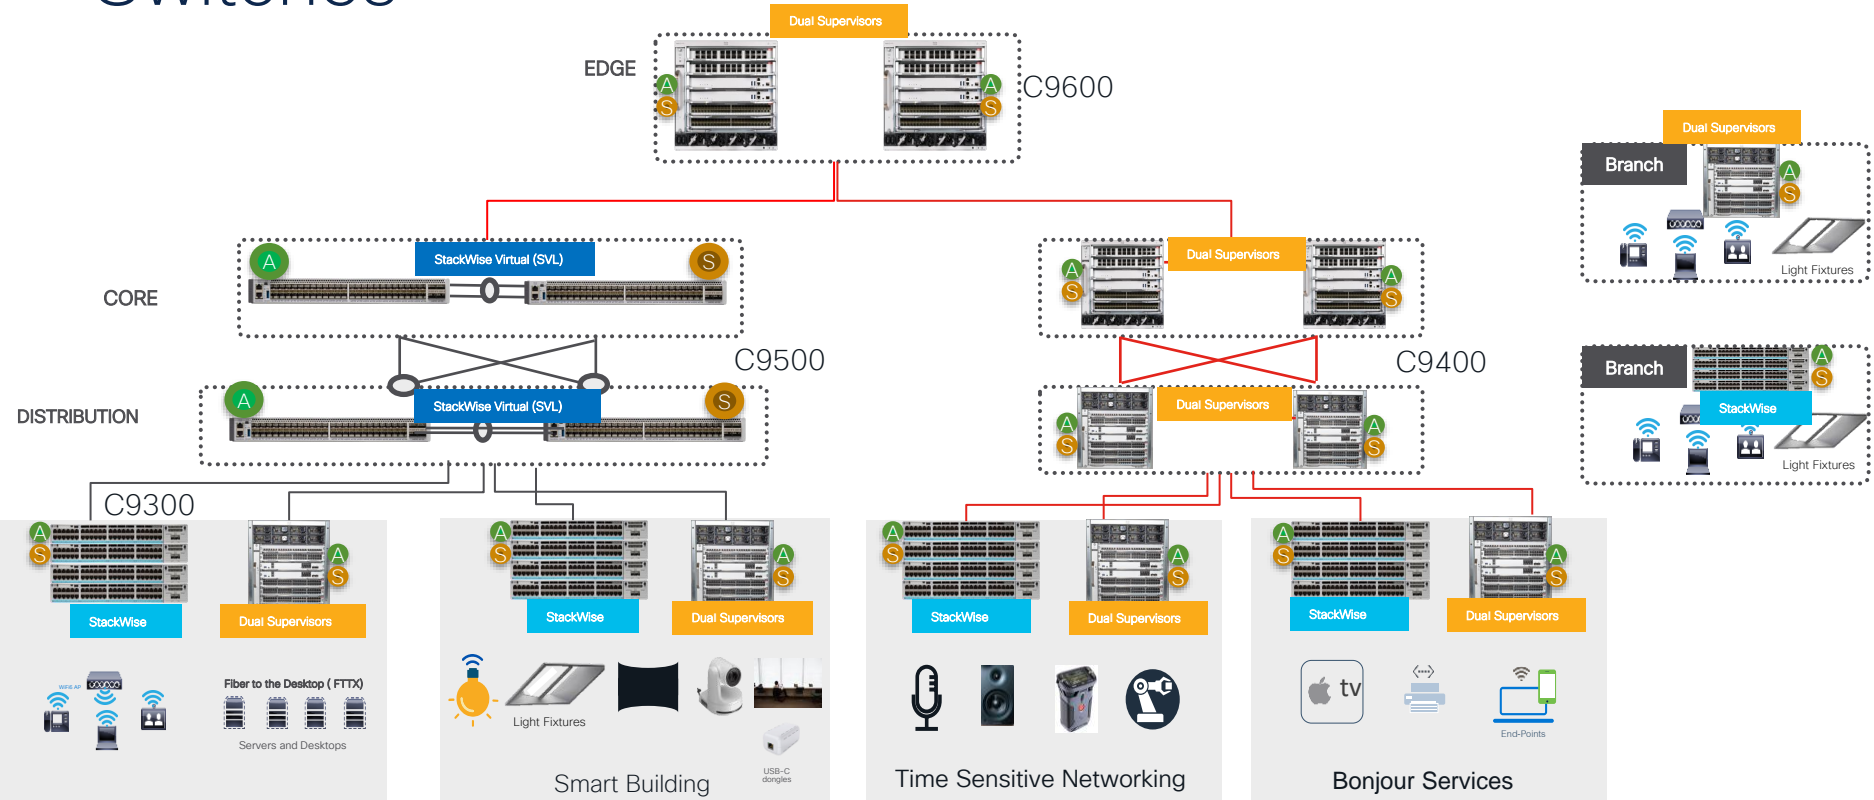


Operational Resiliency

- **ISSU:** Upgrade software with minimal to no traffic loss
- **xFSU :** Upgrade or Reload the Catalyst 9300 with very minimal traffic loss
- **Hot Patching:** No downtime for bug fixes (no reboot)

Eliminate downtime with High Availability designed at every level

Highly Available Architecture with Catalyst 9000 Switches



High Availability Components in Catalyst 9000

Stateful Switchover (SSO)

SSO Aware Applications

Forwarding Information Base

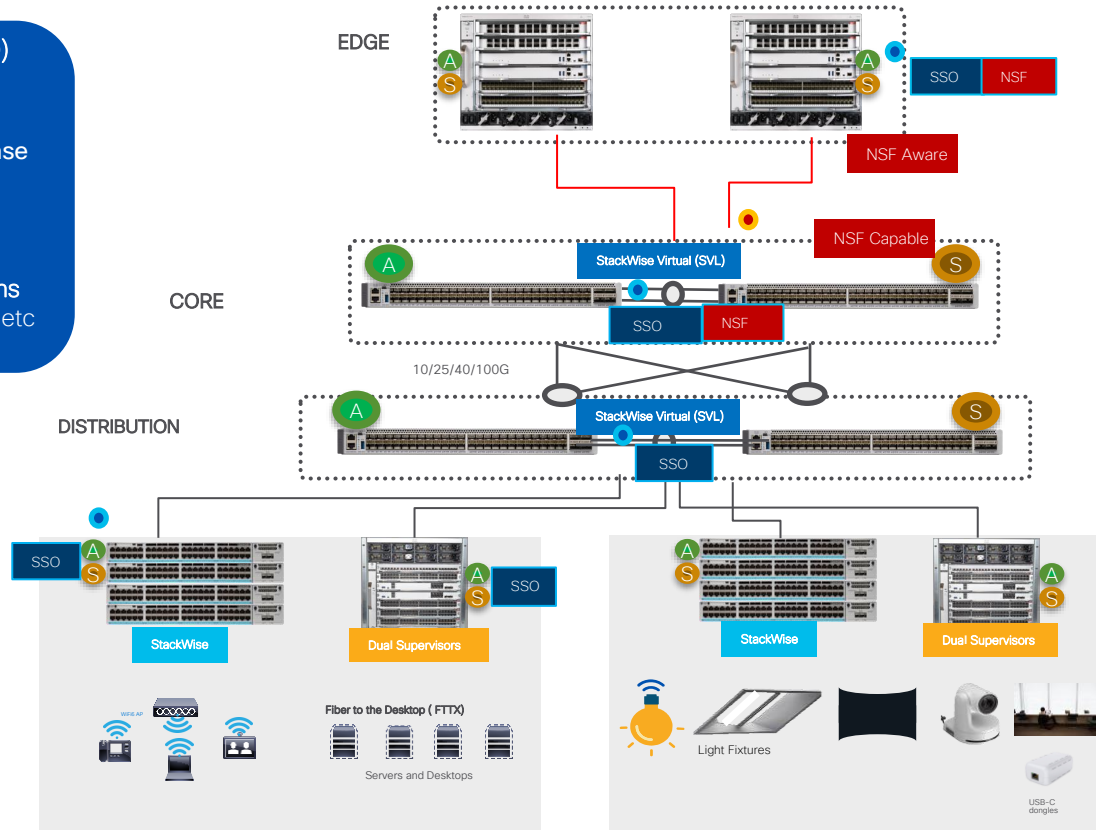
IEEE 802.1x
PAgP / LACP
...and more

SSO Compliant Applications

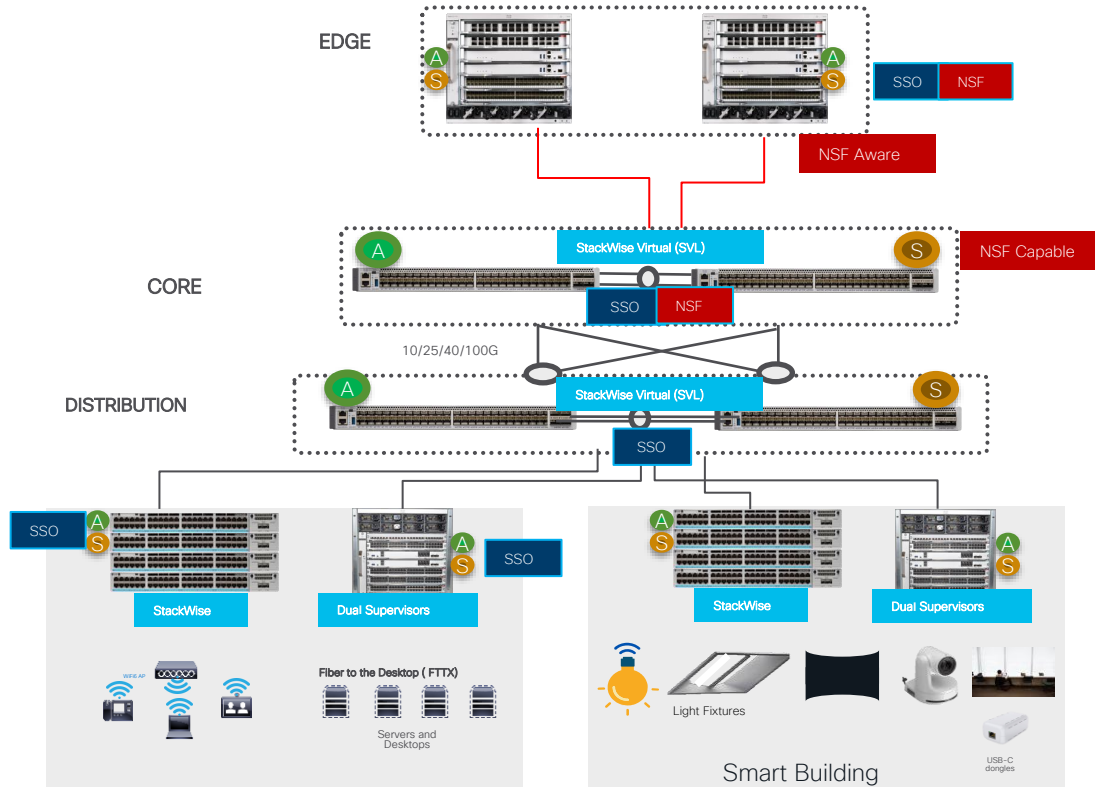
Routing Protocols, Netflow, etc

Non Stop Forwarding(NSF) or
Graceful Restart

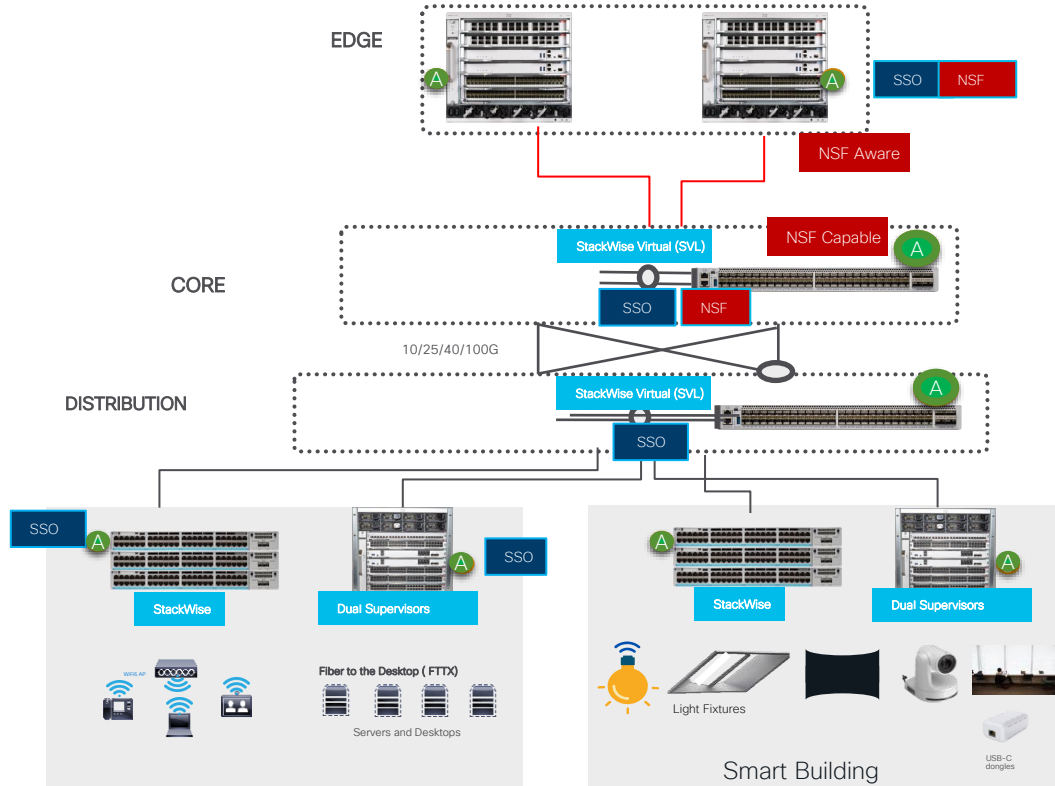
OSPF,BGP,LDP, etc



Typical Campus Network Architecture Failures



Typical Campus Network Architecture Recovery



<200 msec
Sub-Second Traffic Convergence

- ❖ SSO is Enabled by default
- ❖ NSF should be explicitly enabled

Catalyst 9000: Platform Resiliency



C9300 Fixed Platform

- **StackWise: Redundant System for high availability** with NSF/SSO
- **StackPower:** Redundant Power Supplies providing 1+ N redundancy
- **Redundant Fan & Power Supply** in case of any hardware failure



C9400/C9600 Modular Chassis

- **Redundant Supervisor:** Redundant System for high availability, simplified configuration
- **Redundant Fan & Power Supply** in case of any hardware failure




C9500 Fixed Platform

- **StackWise Virtual:** Redundant System for high availability with NSF/SSO
- **Redundant Fan & Power Supply** in case of any hardware failure

Sub-second Traffic impact during failures across C9k Family

Device Redundancy with StackWise



Stackwise-
160/320480/1T

Catalyst 9300 Platform Resiliency for Power Supplies

Normal

PS failure

Combined
(default)



Load sharing on all PSs



Load on Single PS

Redundant



System and PoE Load sharing on active PS
Standby PS shares the load of a system only



Standby PS becomes active
System enters alarm state

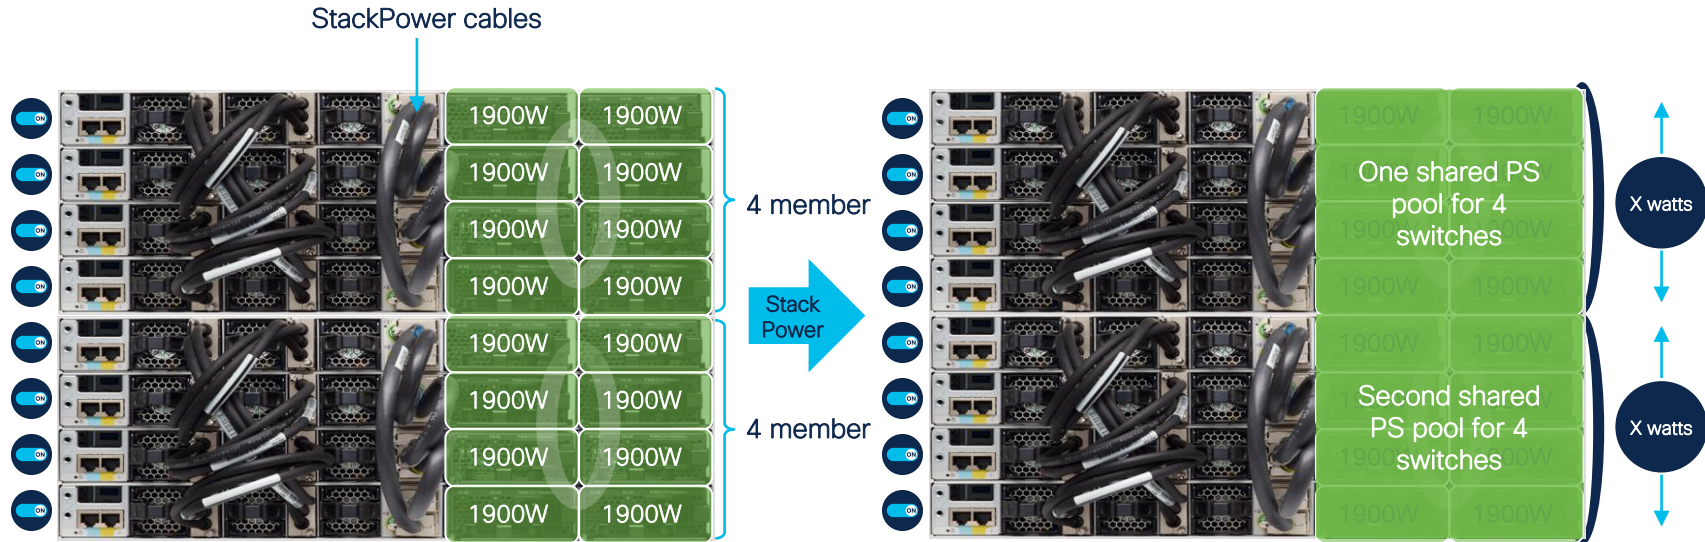
Failed PS

Active

Standby

Power Redundancy with Stacking Power Supplies

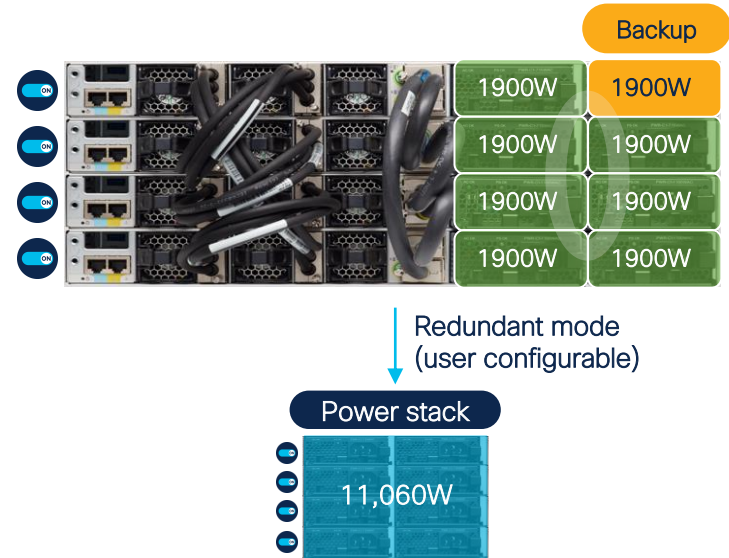
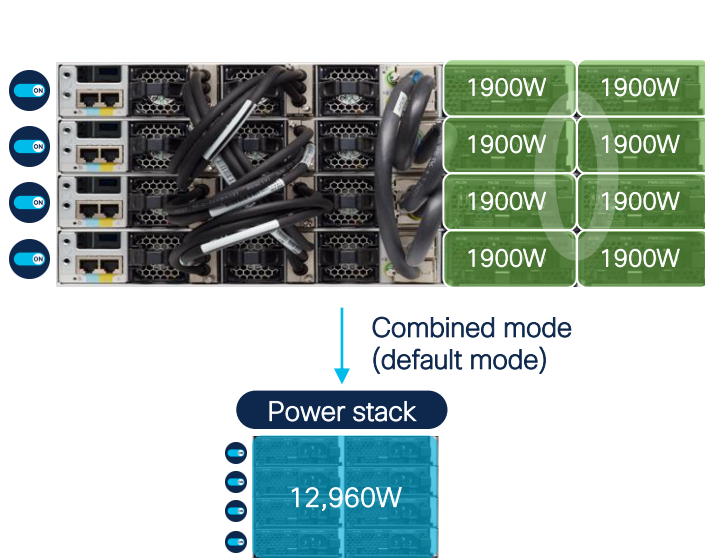
StackPower



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

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

Power stack modes




- All power supplies contribute to common output power budget

- One of the highest valued power supplies is not used in the power budget calculation


Power stack deployment models

Ring model (default)



Power sharing and redundancy are available only within individual pools

Star model (requires XPS 2200 device)



Power sharing and redundancy extend across all eight members in the stack

PoE Redundancy at Port Level

2-event classification


- Fast power negotiation without Link Layer Discovery Protocol (LLDP)
- Physical layer negotiation < 1 second

Perpetual PoE

- Uninterrupted PoE power during control plane reboot

Fast PoE

- Bypasses Cisco IOS® control plane boot
- Restores power to Powered Device (PD) within 30 seconds of power resumption




Power priority

Load shedding

Standalone mode



Stack mode



Load shedding based on configured priority

1. Low-priority ports
2. High-priority ports
3. Switch priority – Highest priority

```
9300-GIR-Access-144#show stack-power load-shedding switch 1
Power Stack      Stack  Stack  Total  Rsvd  Alloc  Unused  Num  Num
Name            Mode  Topolgy Pwr(W) Pwr(W) Pwr(W)  Pwr(W) SW  PS
-----
Powerstack-1    SP-PS Stdaln 1100   30    240    830    1   1

Power Stack      Priority  Consumd  Consumd  Consumd  Alloc  Alloc
SW Name         Sw-Hi-Lo Sw(W)    Hi(W)    Lo(W)    Hi(W)  Lo(W)
-----
1  Powerstack-1  4-13-22  153     0        0        0        0

Switch 1 High Priority Active (powered) Ports:

Switch 1 High Priority Inactive (unused) Ports:

Switch 1 Low Priority Active (powered) Ports:

Switch 1 Low Priority Inactive (unused) Ports:
Gi1/0/1, Gi1/0/2, Gi1/0/3, Gi1/0/4, Gi1/0/5, Gi1/0/6,
Gi1/0/7, Gi1/0/8, Gi1/0/9, Gi1/0/10, Gi1/0/11, Gi1/0/12,
Gi1/0/13, Gi1/0/14, Gi1/0/15, Gi1/0/16, Gi1/0/17, Gi1/0/18,
Gi1/0/19, Gi1/0/20, Gi1/0/21, Gi1/0/22, Gi1/0/23, Gi1/0/24,
```

Operational Resiliency with Extended Fast Software Upgrade

C9300- 17.3.2

C9300X- 17.7.1

Catalyst® 9300/9300X standalone



#Install add file image activate reloadfast commit

Control plane

Data plane

< 30 seconds of traffic impact

Catalyst 9300/9300X stack



#Install add file image activate reloadfast commit

Active control plane

Data plane

< 30 seconds of traffic impact for all ports in the stack

Catalyst 9000: Platform Resiliency



C9300 Fixed Platform

- **StackWise: Redundant System for high availability** with NSF/SSO
- **StackPower:** Redundant Power Supplies providing 1+ N redundancy
- **Redundant Fan & Power Supply** in case of any hardware failure



C9400/C9600 Modular Chassis

- **Redundant Supervisor:** Redundant System for high availability, simplified configuration
- **Redundant Fan & Power Supply** in case of any hardware failure



C9500 Fixed Platform

- **StackWise Virtual:** Redundant System for high availability with NSF/SSO
- **Redundant Fan & Power Supply** in case of any hardware failure


Sub-second Traffic impact during failures across C9k Family

Catalyst 9400/9600 Platform - Supervisor Redundancy

Non Stop Forwarding(NSF) or Graceful Restart

OSPF,BGP,LDP, etc

NSF Capable



Stateful Switchover (SSO)

SSO Aware Applications

Forwarding Information Base

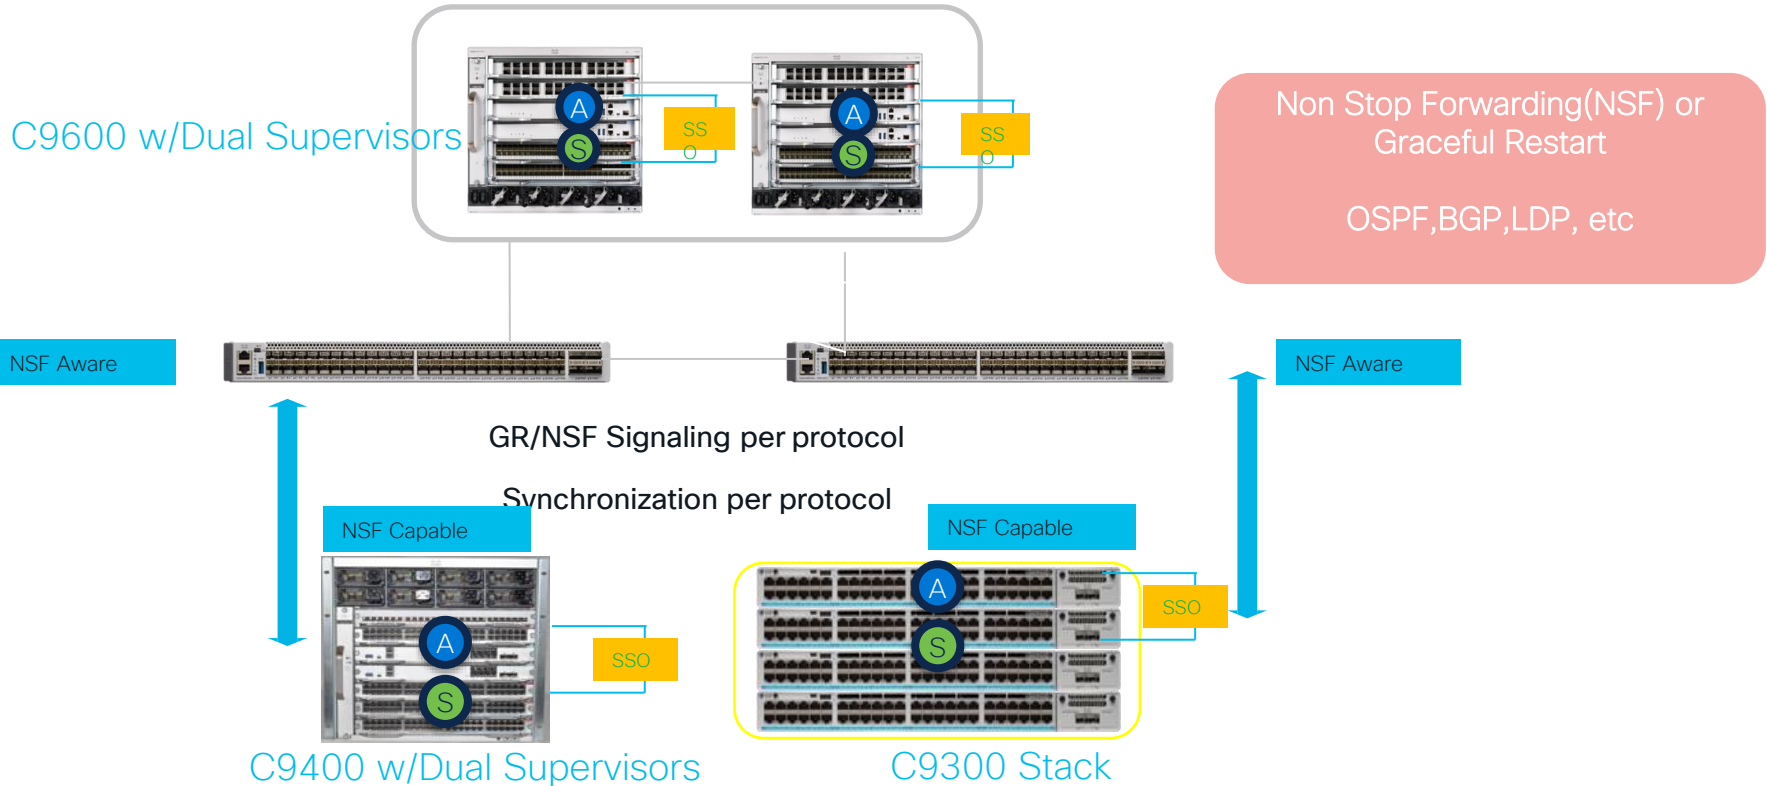
IEEE 802.1x
PAgP / LACP

...and more

SSO Compliant Applications
Routing Protocols, Netflow, etc

Dual Supervisors Required to provide both Control and Data Plane Redundancy

High Availability Architecture with Modular Chassis




Catalyst 9400 Power Redundancy


Normal

PS failure

Combined
(default)



Load sharing on all PSs



Load sharing on functional PSs

Redundant



Load sharing on active PSs
Standby PS in output disabled



Standby PS becomes active
System enters alarm state

Power redundancy: N+1 and N+N

- Default active is PS1 through 4, and standby is PS5 through 8
- Standby power slots are configurable



- Default active is PS1 through 7, and standby is PS8
- Standby power slot is configurable



```
SW(config)#power redundancy-mode redundant ?  
  N+N  Redundant N+N (N is active, N is standby)  
  N+1  Redundant N+N (N is active, 1 is standby)  
SW(config)#power redundancy-mode redundant N+1 ?  
 <1-8> standby slot in N+N mode  
SWR(config)#
```

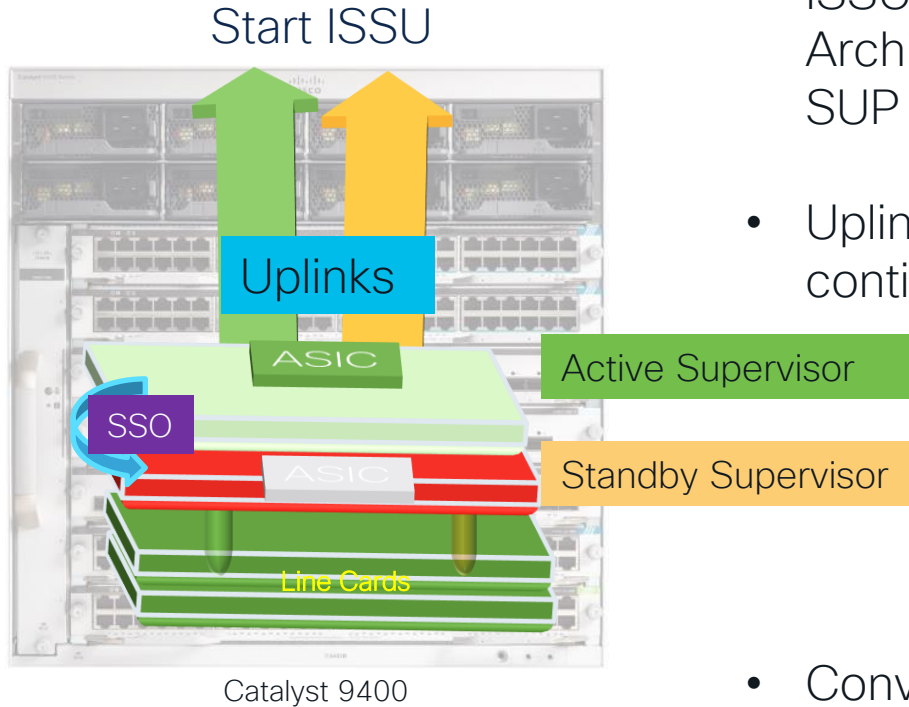
```
SW(config)#power redundancy-mode redundant ?  
  N+N  Redundant N+N (N is active, N is standby)  
  N+1  Redundant N+N (N is active, 1 is standby)  
SW(config)#power redundancy-mode redundant N+1 ?  
 <1-8> standby slot in N+1 mode  
SWR(config)#
```

Active

Standby

Operational Resiliency with ISSU

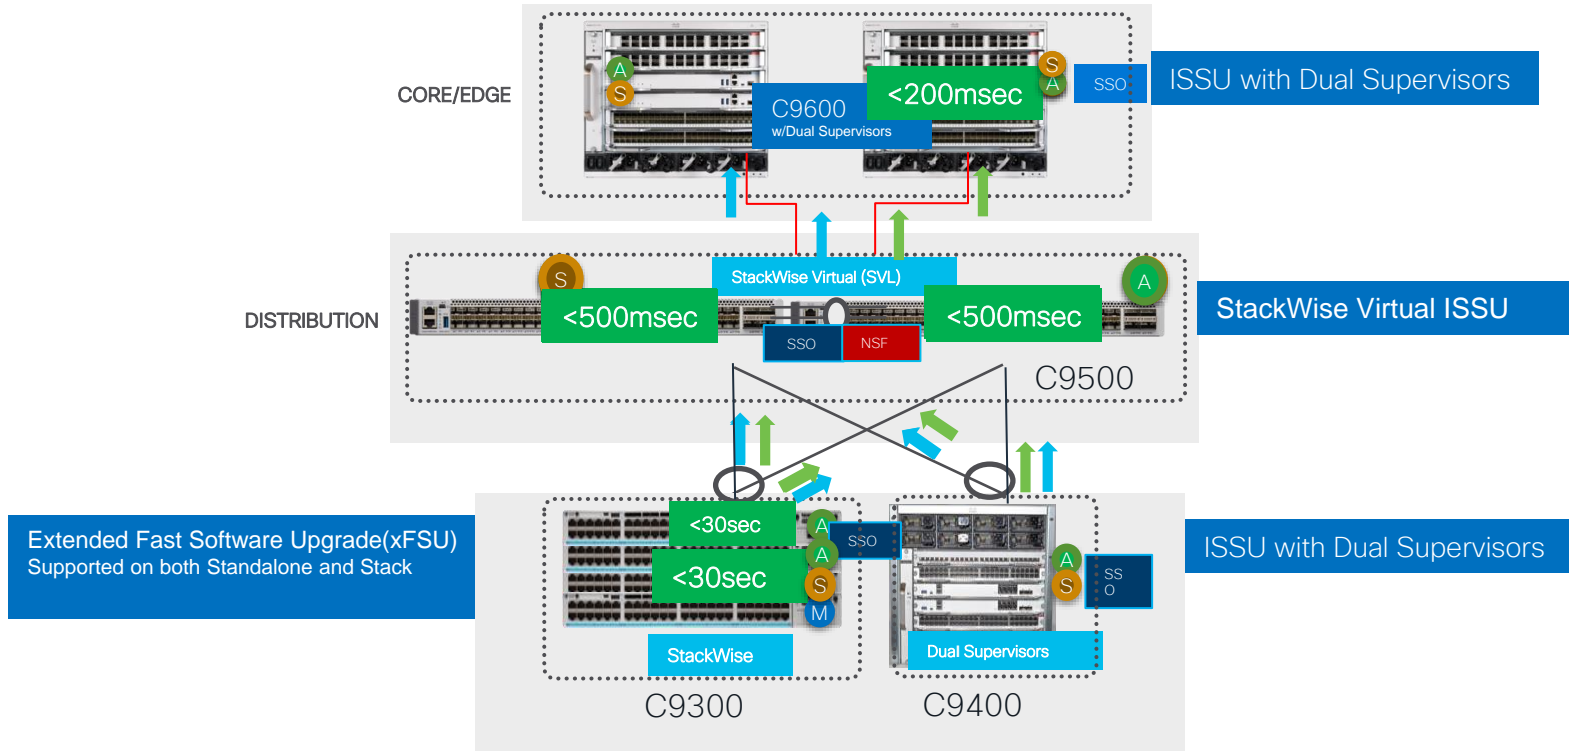
Dual Supervisors



- ISSU Process leverages SSO/NSF Architecture with upgrading the Standby SUP first and then the Active SUP
- Uplinks on both active and standby SUP continue to forward traffic on C9400
- Convergence is less than 200 msec

Campus Network Architecture – Operational Resiliency

Software Upgrades



Campus Network Architecture – Operational Resiliency

ISSU Software Upgrades via DNAC

The image displays four overlapping screenshots of the Cisco DNA Center Provision/Inventory page, illustrating the steps of an ISSU software upgrade. The steps are: 1. Analyze Selection, 2. Distribute, 3. Activate, 4. Schedule and Cleanup, and 5. Summary.

Supervisors (blue box) points to the 'Supervisors' tab in the top right of the first screenshot.

Scheduling Software Activation (blue box) points to the 'Schedule' section in the fourth screenshot, which includes options for 'Now' or 'Later' and fields for 'Task Name*', 'Start Date/Time', and 'Time Zone'.

Software Distribution Checks (green box) highlights the 'Software Distribution Checks' section in the second screenshot, which includes 'Fabric Device Upgrade Check' and 'Flash check'.

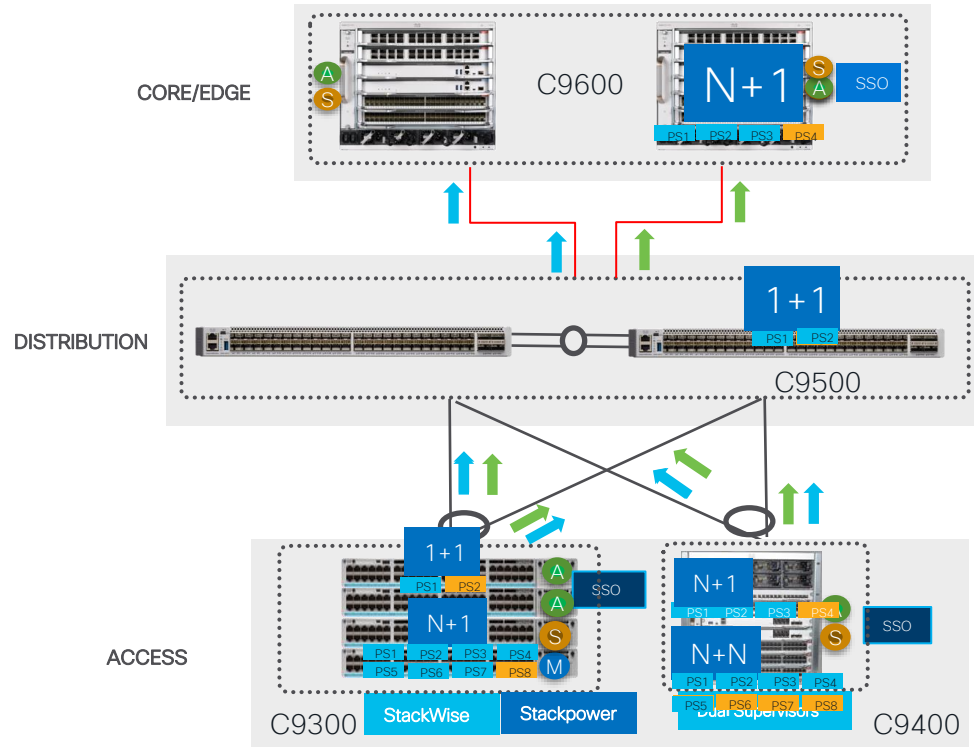
Software Activation Checks (green box) highlights the 'Software Activation Checks' section in the third screenshot, which includes 'Spanning Tree Summary Check', 'CDP neighbors Check', 'Interface Check', 'Fabric Device Upgrade Check', 'Config register check', and 'Startup config check'.

Not able to see the check you would (green box) highlights a note at the bottom of the second screenshot: "Not able to see the check you would like to run? You".

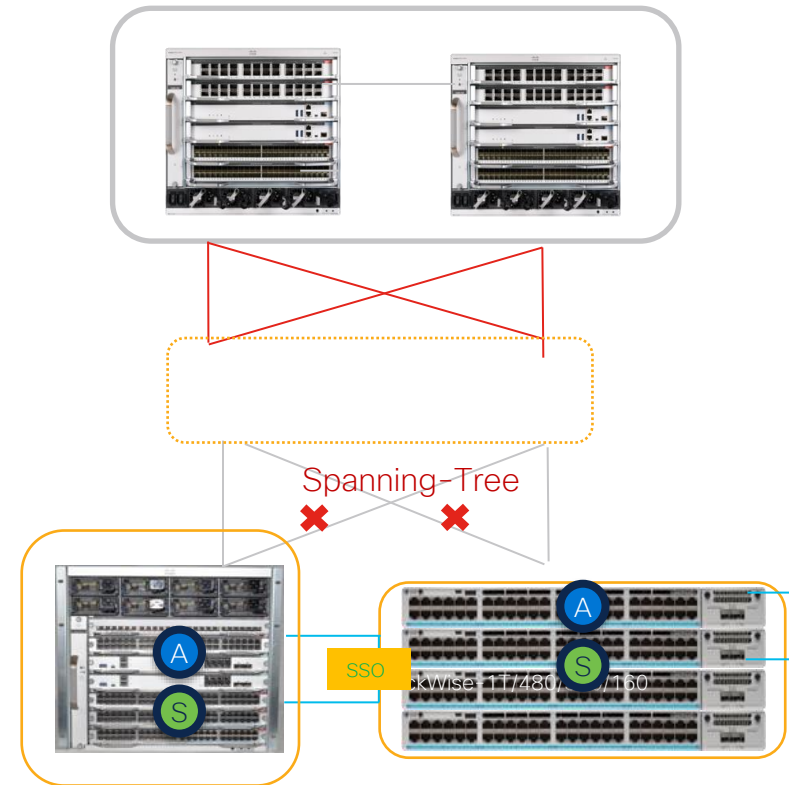
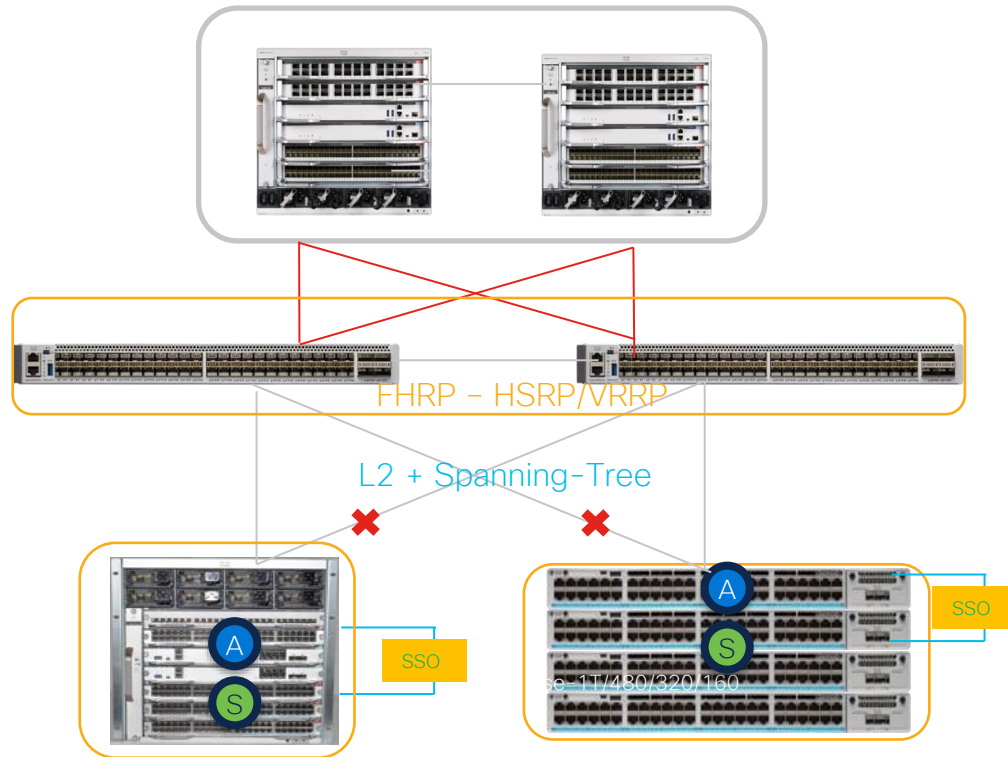
Not able to see the check you would like to run? You (green box) highlights a note at the bottom of the third screenshot: "Not able to see the check you would like to run? You".

Campus Network Architecture – Operational Resiliency


Power Supply Resiliency



High Availability Architecture for Design Resiliency



High Availability Architecture



Simplify Operations by Eliminating STP, FHRP and Multiple Touch-Points

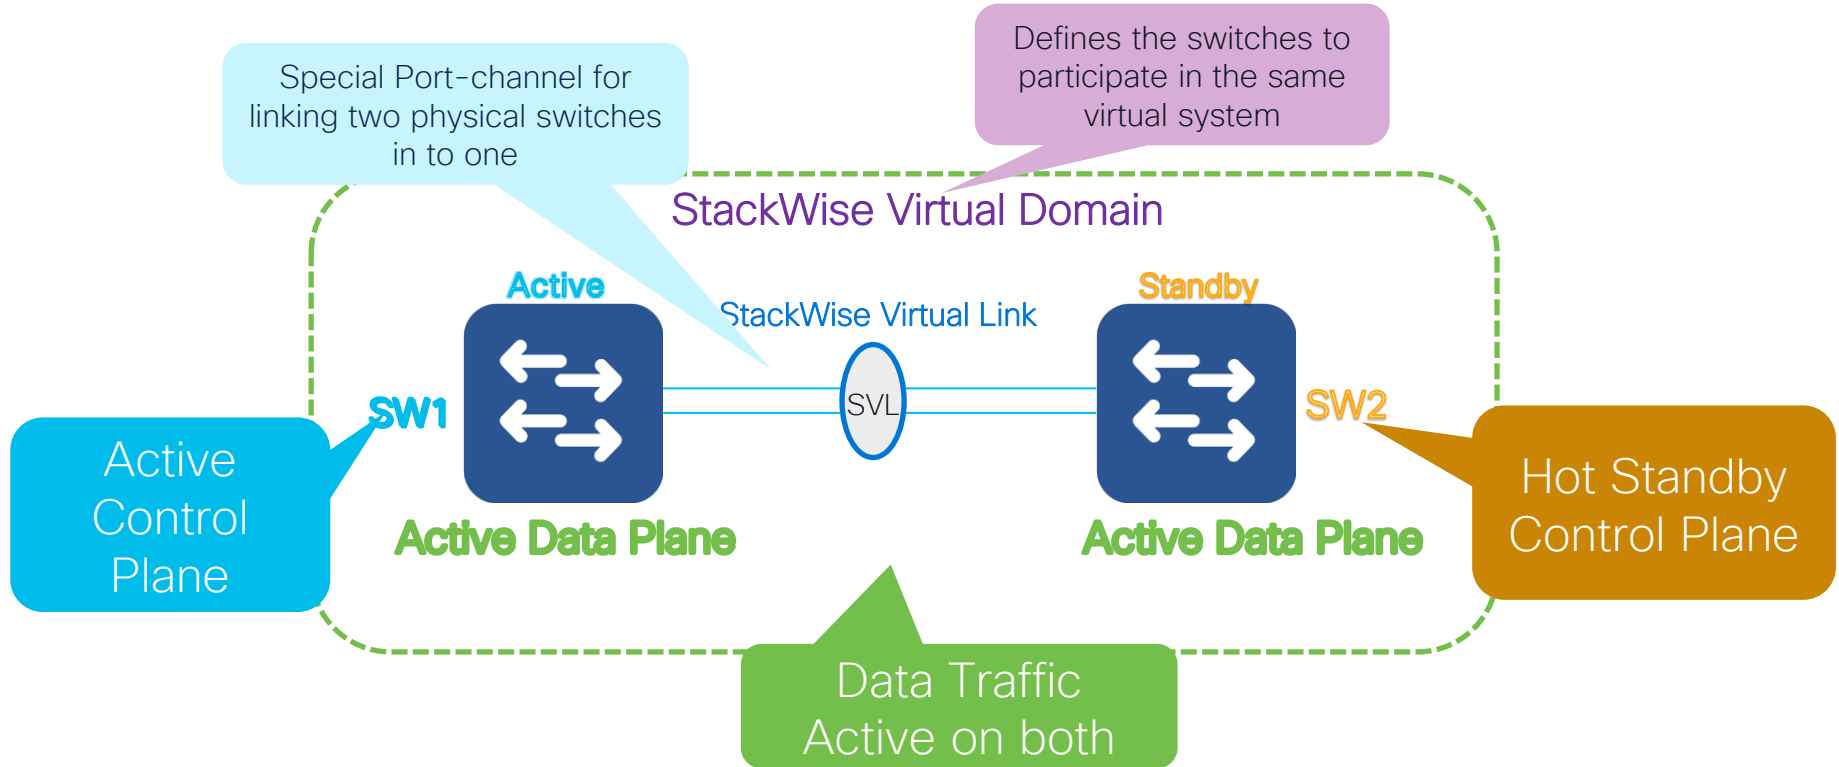
Double Bandwidth & Reduce Latency with Active-Active Multi-chassis EtherChannel (MEC)

Minimizes Convergence with Sub-second Stateful and Graceful Recovery (SSO/NSF)

StackWise Virtual



StackWise Virtual – Key Concepts



Stackwise Virtual Architecture

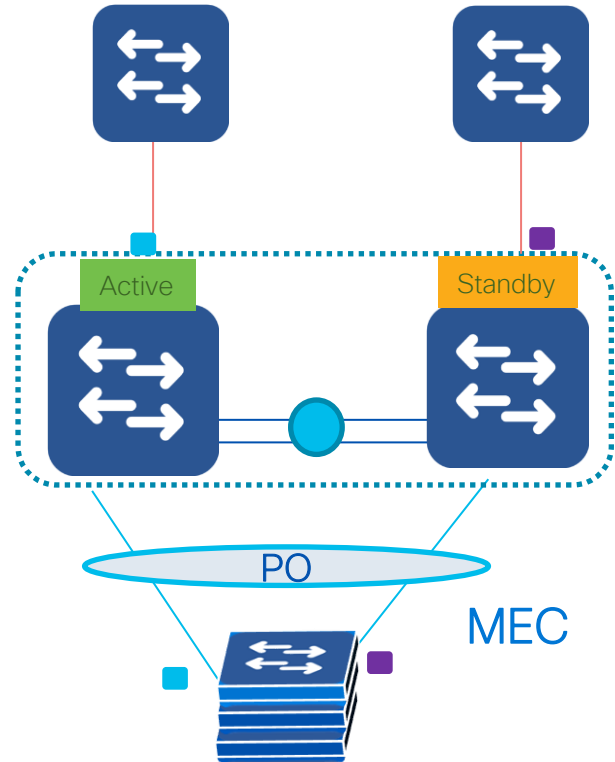
Active/Active Data Plane and Multi-Chassis EtherChannel

Active/Active Data Plane

- Both the switches are capable of forwarding the traffic locally without sending it over Interconnected-Link


Multi-Chassis EtherChannel

- Port-Channel Spanning across multiple Chassis



StackWise Virtual – Multi-Chassis EtherChannel

- Multi-Chassis EtherChannel (MEC) in StackWise Virtual enables cross stack-member link bundling into single logical L2/L3 Interface
- MECs can be deployed in three modes – Cisco PAgP, LACP and Static (ON)
- StackWise Virtual EtherChannel Support
 - C9400: Support 252 port-channel with IOS-XE 16.12 or later
 - Port channel ID 127 and 128 are reserved for SVL
 - Port channel ID 1-126 and 129-252 are available for L2/L3 network configuration
 - All other C9K platforms: Support 128 port-channel
 - Port channel ID 128 is reserved for SVL
 - Port channel ID 1-127 are available for L2/L3 network configurations



Stackwise Virtual Failure and Recovery

High Availability


Failure Scenario's

In a SVL Domain, one switch is elected as Active and the other as Standby

All Neighbors view SVL as a single Entity, single MAC, single IP

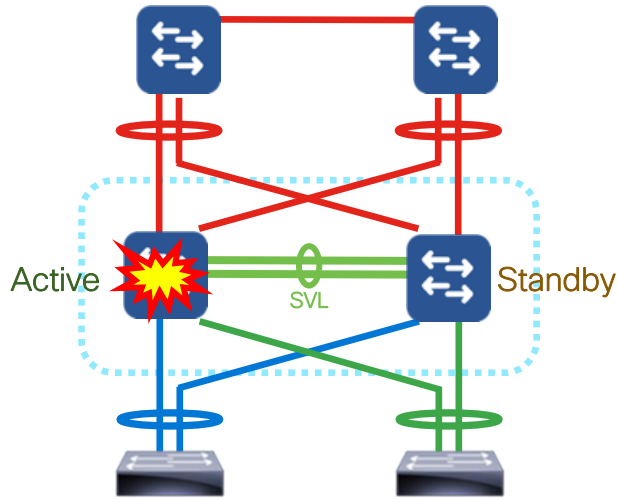
Since the SVL is always configured as a Port Channel, the chance of the entire SVL going down is remote...

However... IT IS POSSIBLE! ☹️



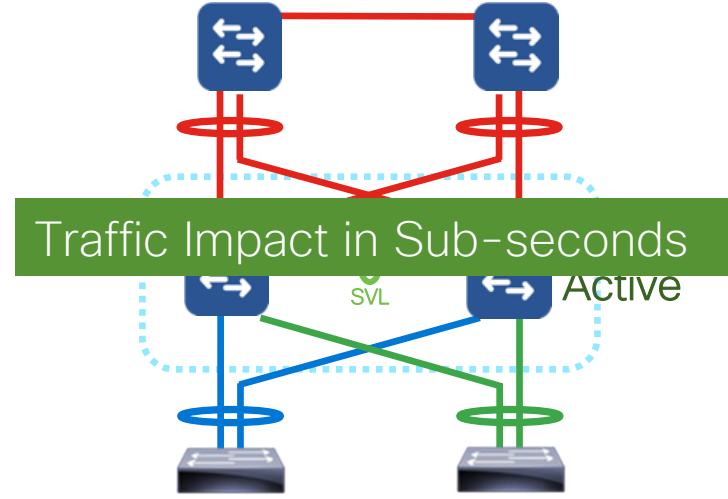
Recommend to deploy the SVL with 2 or more links, distributed across entities for highest redundancy

StackWise Virtual – Inter Chassis SSO/NSF



Virtual Switch incurs a failure of the (SSO) Active Switch 1

The Standby Switch detects failure by loss of all SVL ports, or no replies to SSO keep-alive packets



The original Standby Switch now takes over as the new Virtual Switch Active

Virtual Switch initiates Graceful Restart (NSF)

Non Stop Forwarding of packets continues using hardware entries synced to Switch 2

NSF Aware neighbors exchange protocol updates with the new Virtual Switch Active

High Availability

Dual-Active Detection

If the entire SWV bundle fails, the SWV Domain will enter into a “Dual Active” scenario without DAD


Both switches transition to SSO Active state, and share the same network configuration

- IP addresses, MAC address, Router IDs, etc.

This can cause communication problems in the network!

3 Step Process


- 1** Dual-Active Detection - using any detection method enabled in the system.
- 2** Previous SWV Active shuts down ALL interfaces, and enters “Recovery Mode” ... preventing further network disruption
- 3** Dual-Active Recovery - when the SWV recovers, the switch in Recovery Mode will reload to boot into a preferred standby state



High Availability

Dual-Active Protocols


Fast Hello



- ❖ Direct L2 Point-to-Point Connection

- ❖ Sub-Second Convergence
 - ❖ Typically ~50-100ms

Enhanced PAGP

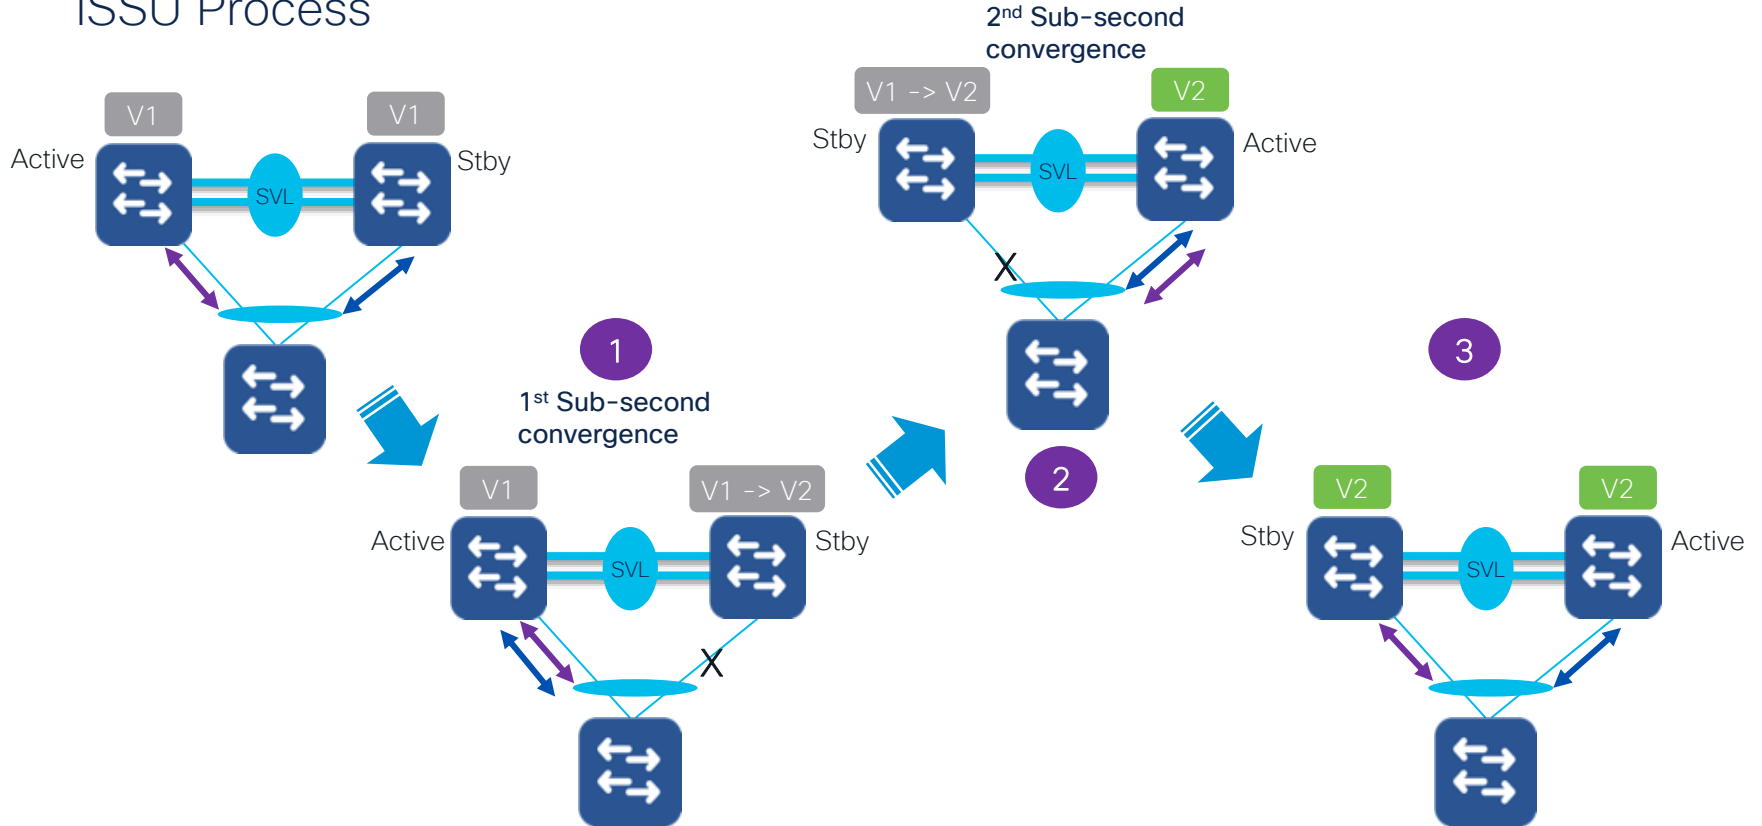


- ❖ Requires ePAGP capable neighbor:

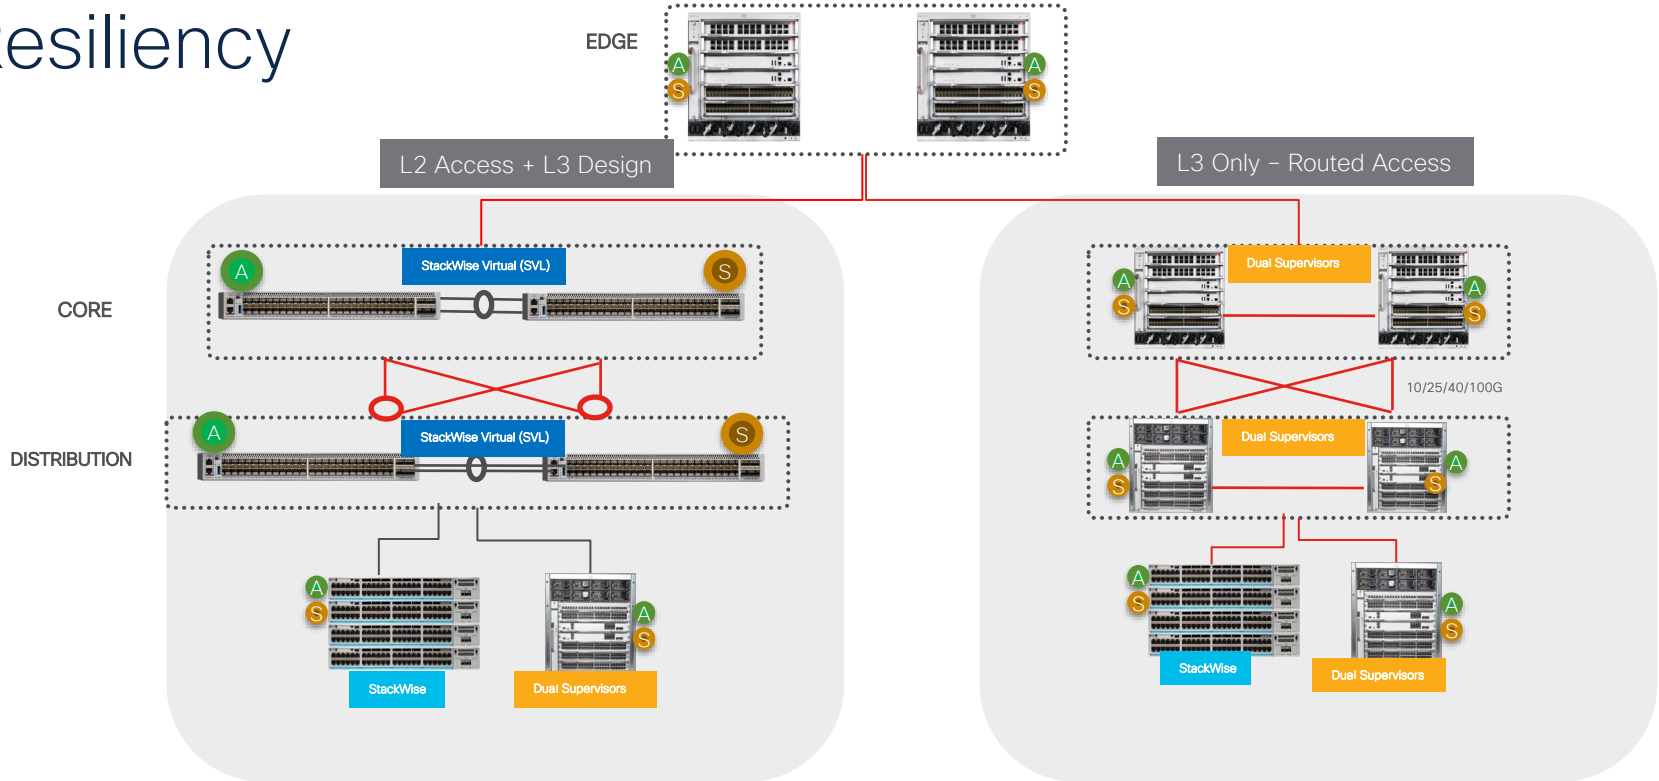
- ❖ Sub-Second Convergence
 - ❖ Typically ~200-250ms

Stackwise Virtual ISSU

ISSU Process

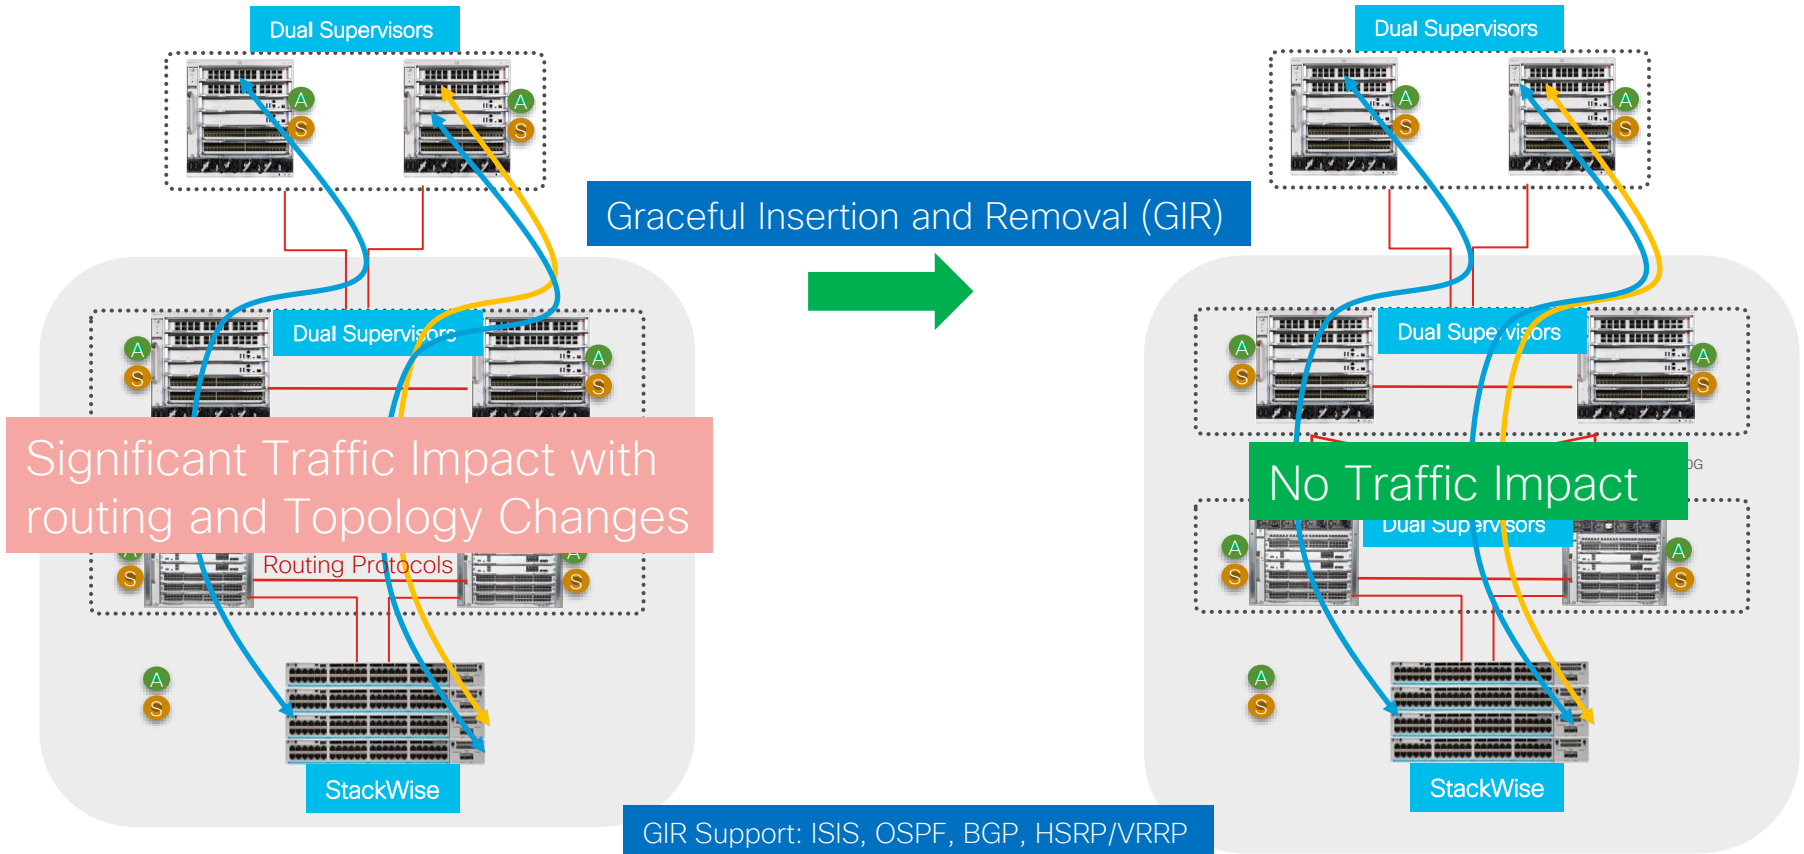


Typical Campus Network Architecture – Design Resiliency



Typical Campus Network Architecture – Design Resiliency

Routed Access Designs

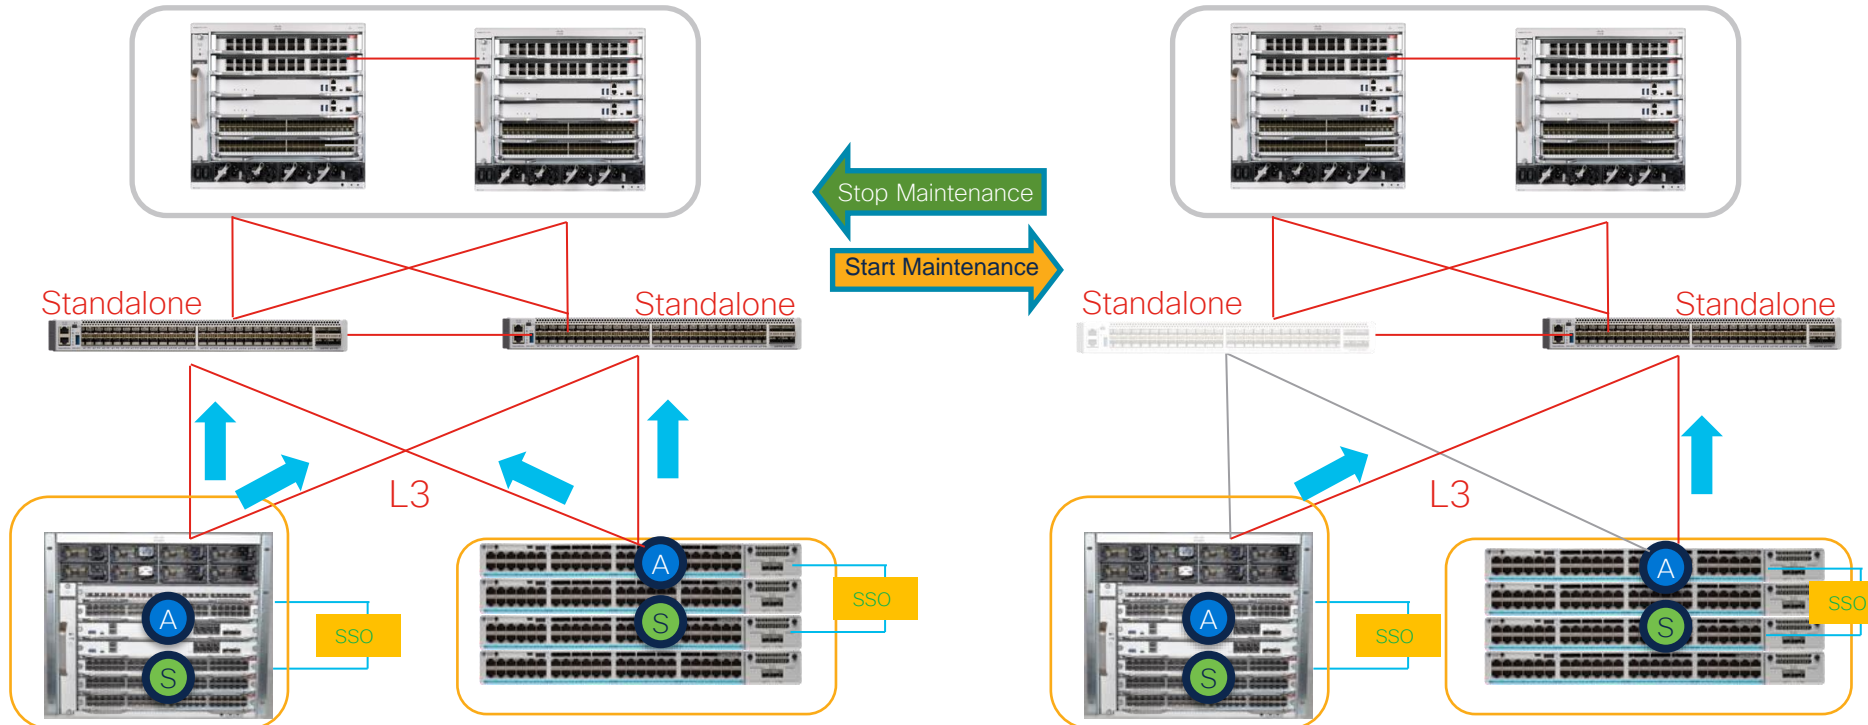


Graceful Insertion and Removal (GIR)



High Availability Architecture

Routed Access Multi-Tier Architecture




L2 and L3 Topology with GIR Isolation

```
9300#start maintenance
Template default will be applied.
Do you want to continue?[confirm]
*Mar 25 17:43:20.162: %MMODE-6-
MMODE_CLIENT_TRANSITION_START: Maintenance Isolate
start for router isis 1
*Mar 25 17:43:50.213: %MMODE-6-
MMODE_CLIENT_TRANSITION_COMPLETE: Maintenance Isolate
complete for router isis 1
*Mar 25 17:43:50.213: MMODE-6-
MMODE_CLIENT_TRANSITION%_START: Maintenance Isolate
start for shutdown l2
*Mar 25 17:44:20.214: %MMODE-6-
MMODE_CLIENT_TRANSITION_COMPLETE: Maintenance Isolate
complete for shutdown l2
*Mar 25 17:44:20.214: %MMODE-6-MMODE_ISOLATED: System
is in Maintenance
```

Order for Maintenance:

BGP -> IGP in parallel (ISIS) -> L2




L2 and L3 Topology with GIR Isolation

```
9300#stop maintenance
```

```
*Mar 25 19:15:40.235: %MMODE-6-  
MMODE_CLIENT_TRANSITION_START: Maintenance  
Insert start for shutdown l2  
*Mar 25 19:16:10.237: %MMODE-6-  
MMODE_CLIENT_TRANSITION_COMPLETE: Maintenance  
Insert complete for shutdown l2  
*Mar 25 19:16:10.237: %MMODE-6-  
MMODE_CLIENT_TRANSITION_START: Maintenance  
Insert start for router isis 1  
*Mar 25 19:16:40.288: %MMODE-6-  
MMODE_CLIENT_TRANSITION_COMPLETE: Maintenance  
Insert complete for router isis 1  
*Mar 25 19:16:40.612: %MMODE-6-MMODE_INSERTED:  
System is in Normal Mode
```

Order for Maintenance:

L2 → IGP in parallel (ISIS) → BGP



Open IOS-XE Patchability



Ready for Software Patching

SMU is an emergency point fix positioned for expedited delivery to a customer in case of a network down or revenue affecting scenario.

Cold Patching: Install of a SMU will require a system reload in the first release. It is traffic impacting.

Hot Patching: Install of a SMU does not require a reload.



Catalyst 9000: Platform Resiliency



C9300 Fixed Platform

- **StackWise: Redundant System for high availability** with NSF/SSO
- **StackPower:** Redundant Power Supplies providing 1+ N redundancy
- **Redundant Fan & Power** Supply in case of any hardware failure



C9400/C9600 Modular Chassis

- **Redundant Supervisor:** Redundant System for high availability, simplified configuration
- **Redundant Fan & Power** Supply in case of any hardware failure



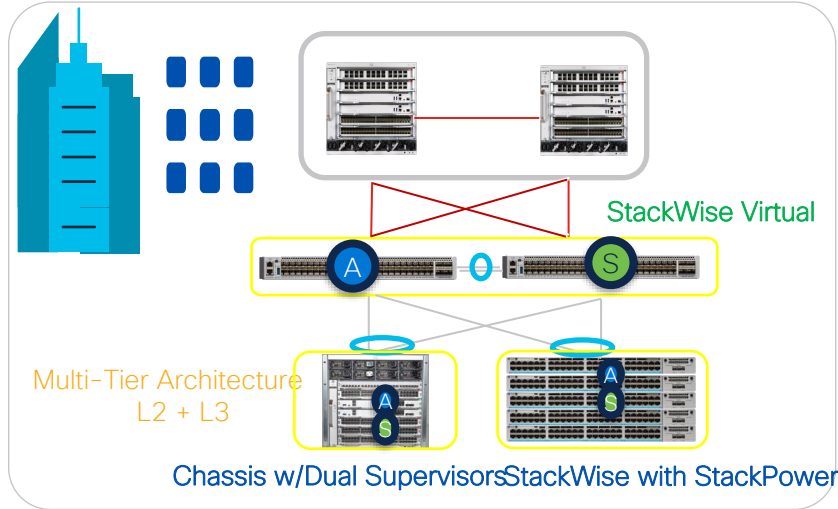
C9500 Fixed Platform

- **StackWise Virtual:** Redundant System for high availability with NSF/SSO
- **Redundant Fan & Power** Supply in case of any hardware failure

Sub-second Traffic impact during failures across C9k Family

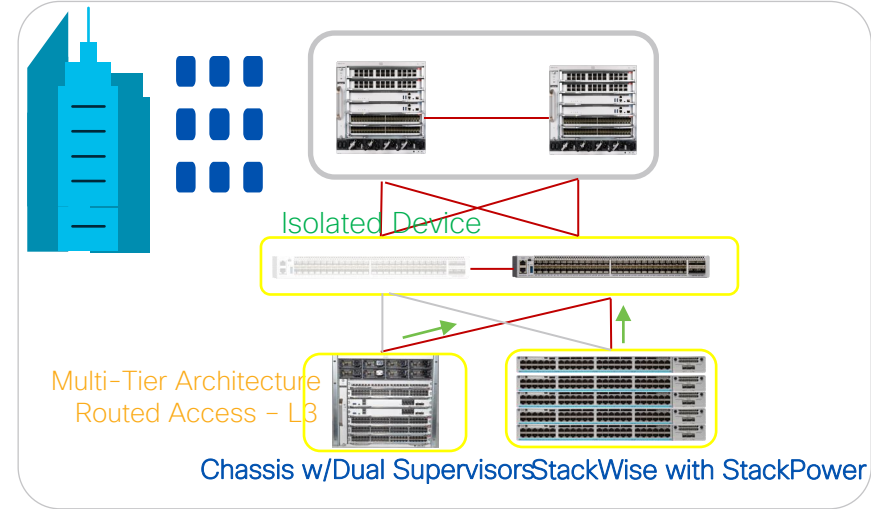
Catalyst 9000: Features for Design Resiliency

StackWise Virtual with Backward StackWise



Device Resiliency with Simplified Design

Graceful Insertion and Removal (GIR)




Device Isolation without Traffic Impact

Device Resiliency and Sub-second failover for Every Design

Catalyst 9000: Operational Resiliency

C9400/C9600 Modular Chassis

Standalone Mode



In-Chassis Dual Supervisor ISSU

< 200 msec Traffic Impact

Requires Dual Supervisor

C9400/C9500/C9600 Platform

StackWise Virtual Mode



StackWise Virtual ISSU

< 1 sec Traffic Impact

C9300/C9300L/C9300X

Standalone and StackWise Mode



Extended Fast Software Upgrade (xFSU)

< 30 sec Traffic Impact

Hot and Cold Patching for IOS Resiliency

Sub-second Traffic impact during software upgrades across C9k Family

Platform/Solution Support

Features	C9200	C9300/X	C9400/X	C9500	C9500X	C9600	C9600X
StackWise	✓	✓	NA	NA	NA	NA	NA
StackWise Virtual (SVL)	NA	NA	✓	✓	✓	✓	✓
SVL ISSU	NA	NA	✓	✓	✓	✓	Roadmap
Dual Supervisor ISSU	NA	NA (xFSU with stacking)	✓	NA	NA	✓	
Cold/Hot Patching	Cold Patching	✓	✓	✓	✓	✓	✓
GIR	NA	✓	✓	✓	✓	✓	✓
Power Redundancy	✓	✓	✓	✓	✓	✓	✓

Complete your Session Survey

- Please complete your session survey after each session. Your feedback is important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (open from Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Session Catalog and clicking the "Attendee Dashboard" at



<https://www.ciscolive.com/emea/learn/sessions/session-catalog.html>

Continue Your Education



Visit the Cisco Showcase for related demos.



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



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



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



The bridge to possible

Thank you

CISCO *Live!*

CISCO *Live!*

ALL IN