Cisco live!

January 29 – February 2, 2018 · Barcelona

BRKCRS-2650

Next Generation High Availability in Campus

Minhaj Uddin, Technical Marketing Engineer

Cisco live!

Cisco Spark

Questions?

Use Cisco Spark to communicate with the speaker after the session

How

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

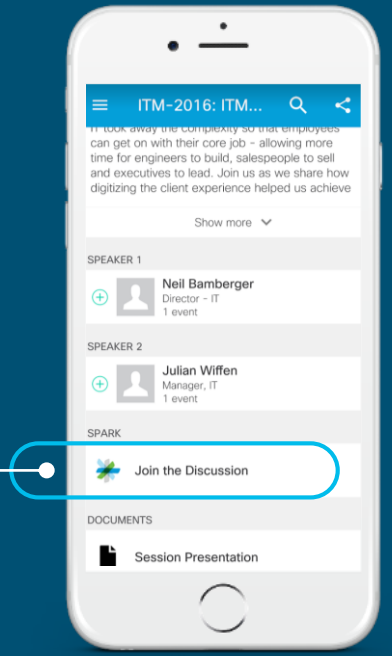

cs.co/ciscolivebot#BRKCRS-2650

High Availability in the Access
Session Overview and Objectives

High Availability has become part of the Cisco DNA and is being deployed on all levels of products

In this session, Our focus will be to learn about the existing and new High Availability features present on the Catalyst 9k Switches. We will also categorize features based on access and Distribution layer in the Enterprise Network. In the end we will see how these features can be leveraged effectively to achieve highly available network. We will also show good design practices across all the features that will help us achieve better service availability.

Agenda

- High Availability Overview and Evolution
- High Availability Solution on the Campus Access
  - Stackable High Availability Solution
  - Modular High Availability Solution
- High Availability Solution on the Campus Distribution
- High Availability Solution in IOS
- Summary/Q&A

Enterprise-Class Availability

Campus Systems Approach to High Availability

- Network-level redundancy
- System-level resiliency
- Enhanced management
- Human ear notices the difference in voice within 150–200 msec
- 10 consecutive G711 packet loss
- Video loss is even more noticeable
- 200-msec end-to-end campus convergence

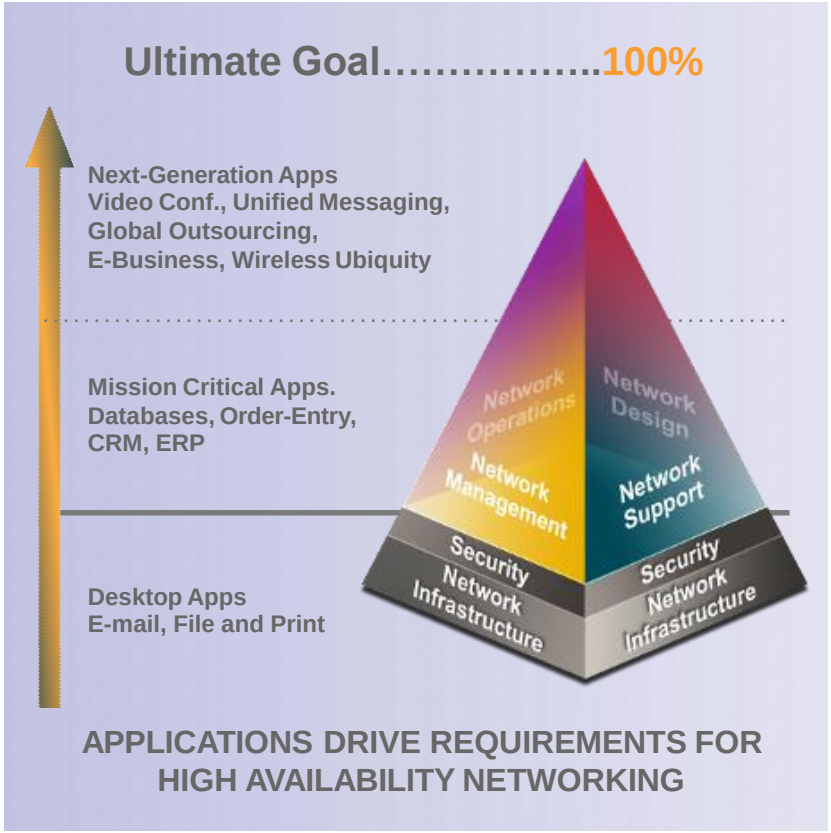

Cisco HA Evolution

No Redundancy	Redundancy with RPR	Redundancy with RPR+	Redundancy with SSO
No Redundant Units	Adding Redundant Units	Adding Redundant Units	Adding Redundant Units
Failure on Supervisor causes reload	Failure on Active Sup causes Switchover	Failure on Active Sup causes Switchover	Failure on Active Sup causes Switchover
Line Cards reload on failure	Standby Unit is in STANDBY_COLD state	Standby Unit is in STANDBY_WARM state	Standby Unit is in STANDBY_HOT state
	Line Cards reload after switchover	Line Cards reload after switchover	Line Cards Stay up after switchover
	Startup Configuration Synchronized to Peer	Startup Configuration Synchronized to Peer	Startup Configuration Synchronized to Peer
		Running Configuration Synchronized to Peer and applied after switchover	Running Configuration Synchronized to Peer and applied.

Cisco HA Evolution

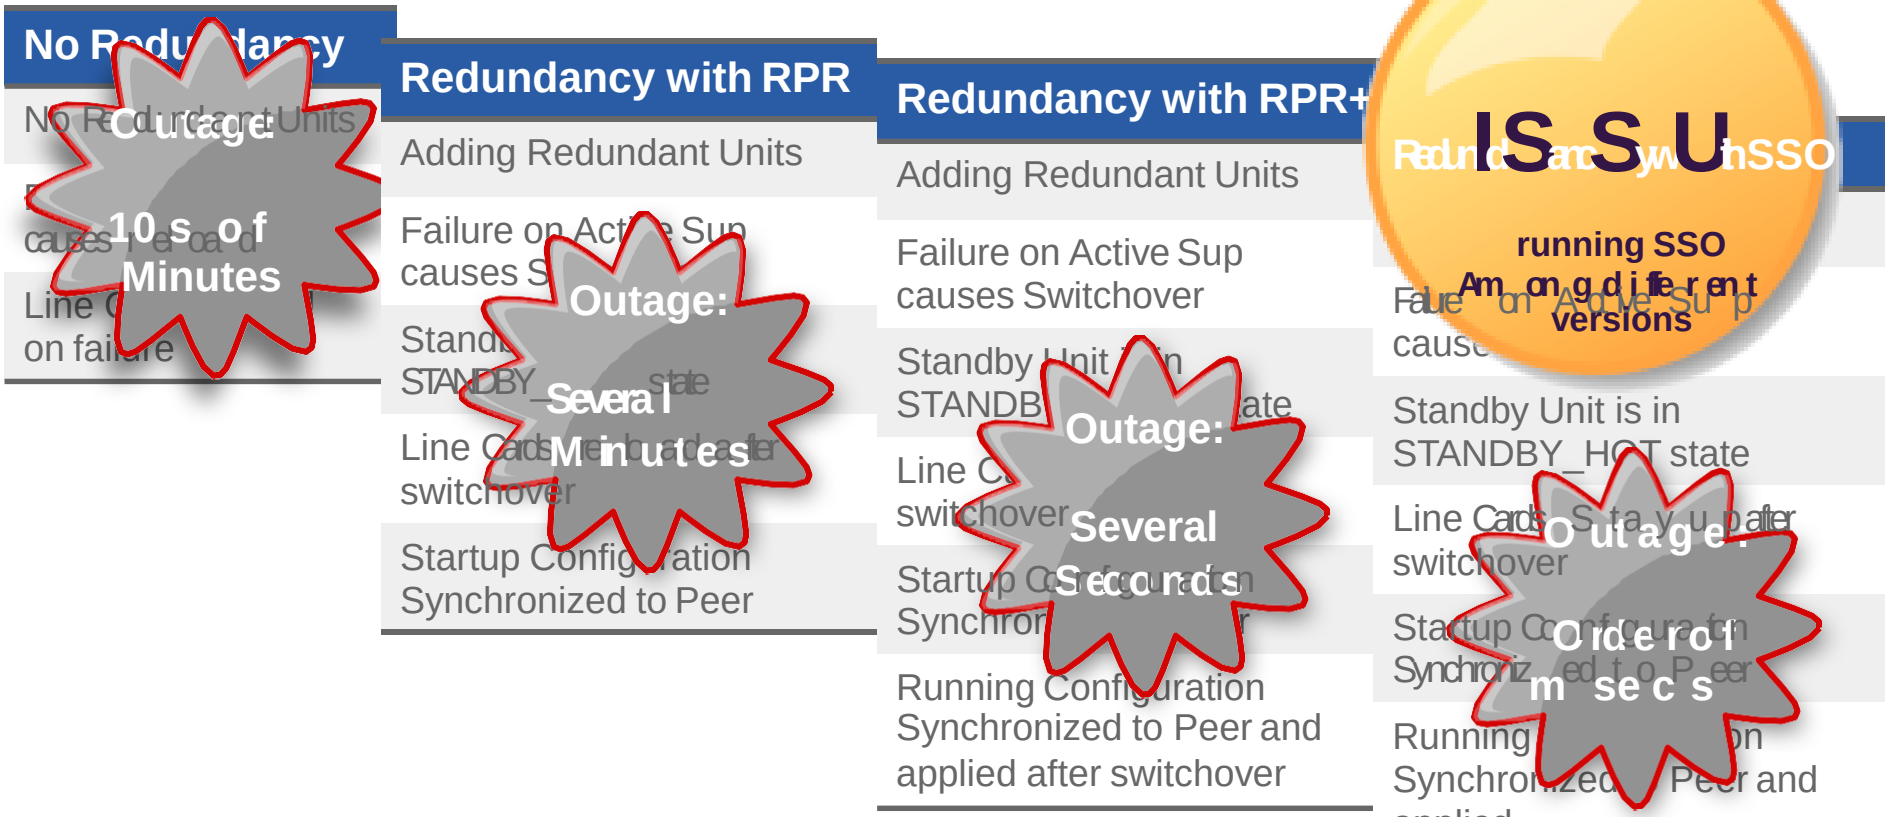

Agenda

- High Availability Overview and Evolution
- High Availability Solution in the Campus Access
  - Stackable High Availability Solution
  - Modular High Availability Solution
- High Availability Solution on the Campus Distribution
- High Availability Solution in IOS
- Summary/Q&A

Importance of High Availability for Access switches

Feature and Device rich layer

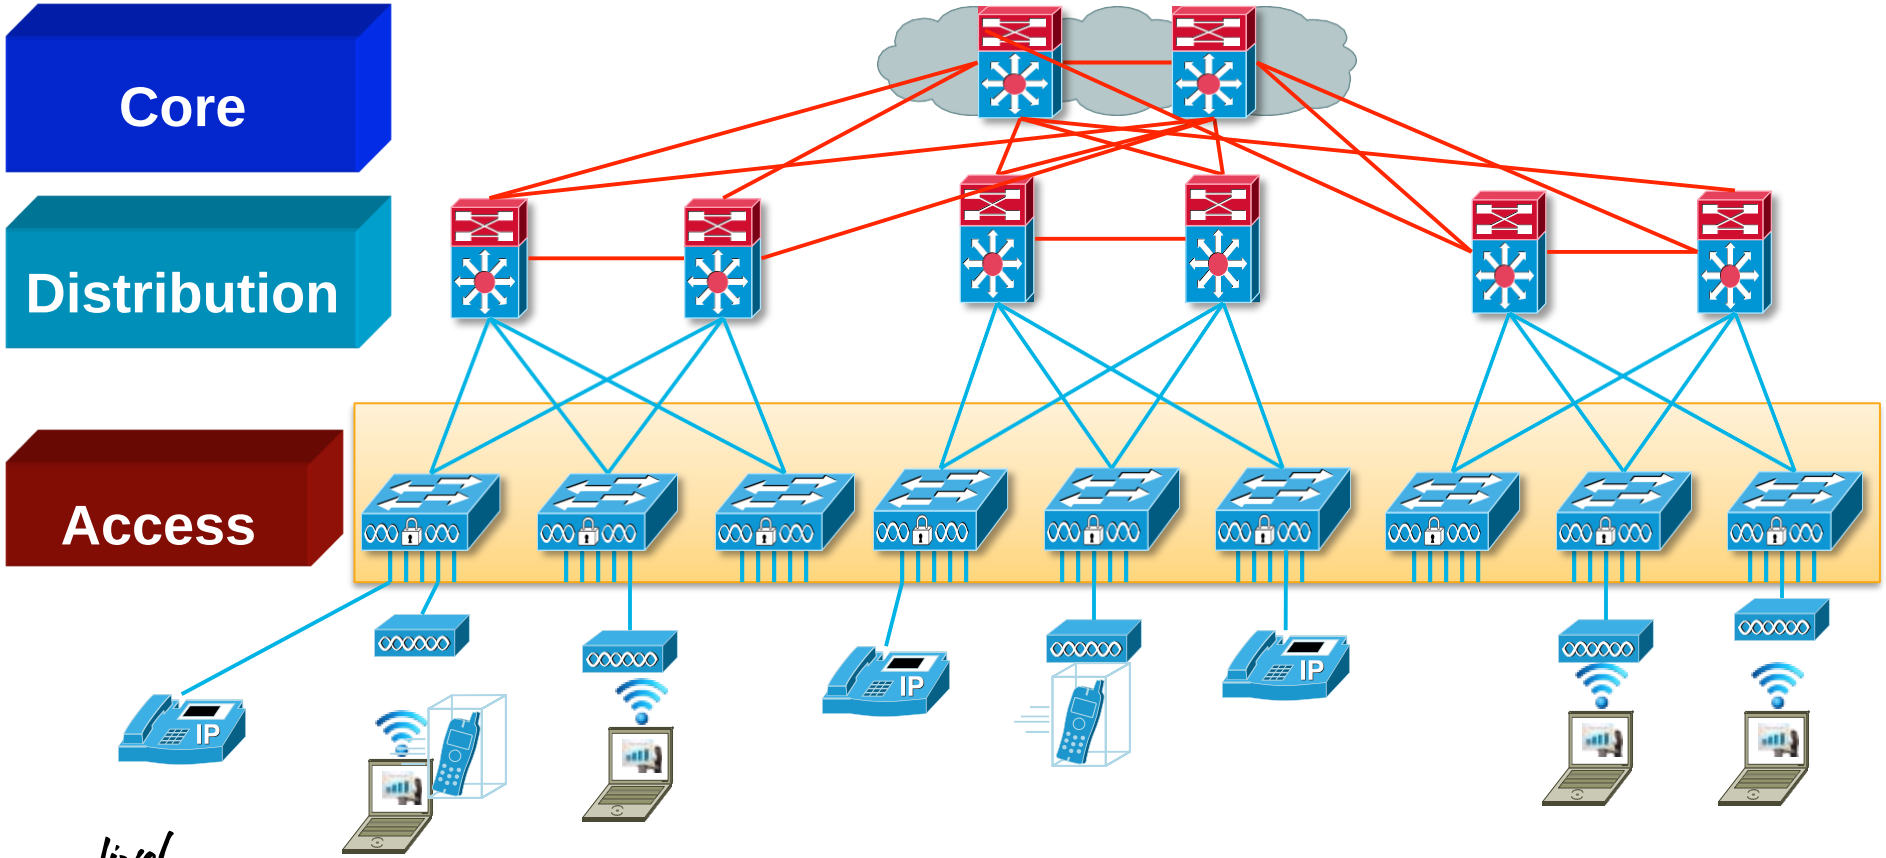

Stacking High Availability Evolution

Access Layer

- One Active and Standby Unit
- Rest are Member Units

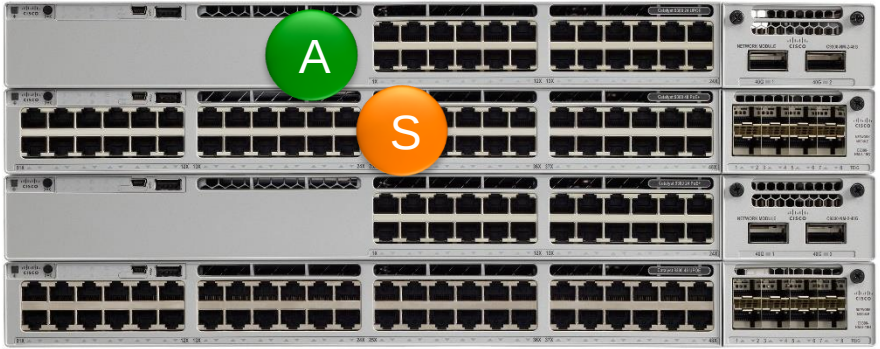

Catalyst 9300

Catalyst 9400

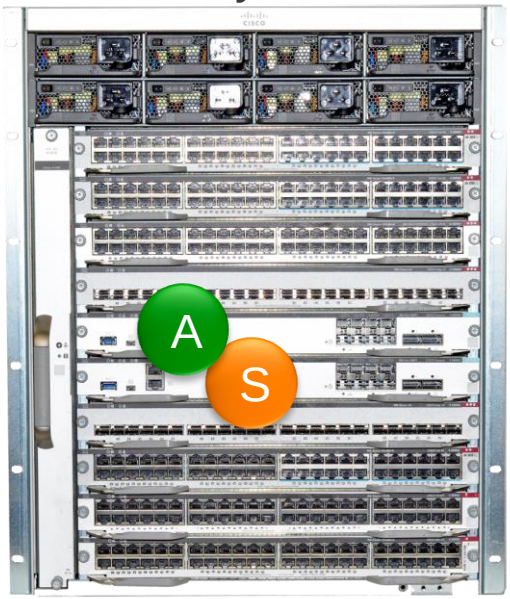

- One Active Supervisor
- One Standby Supervisor

Enterprise Access Layer

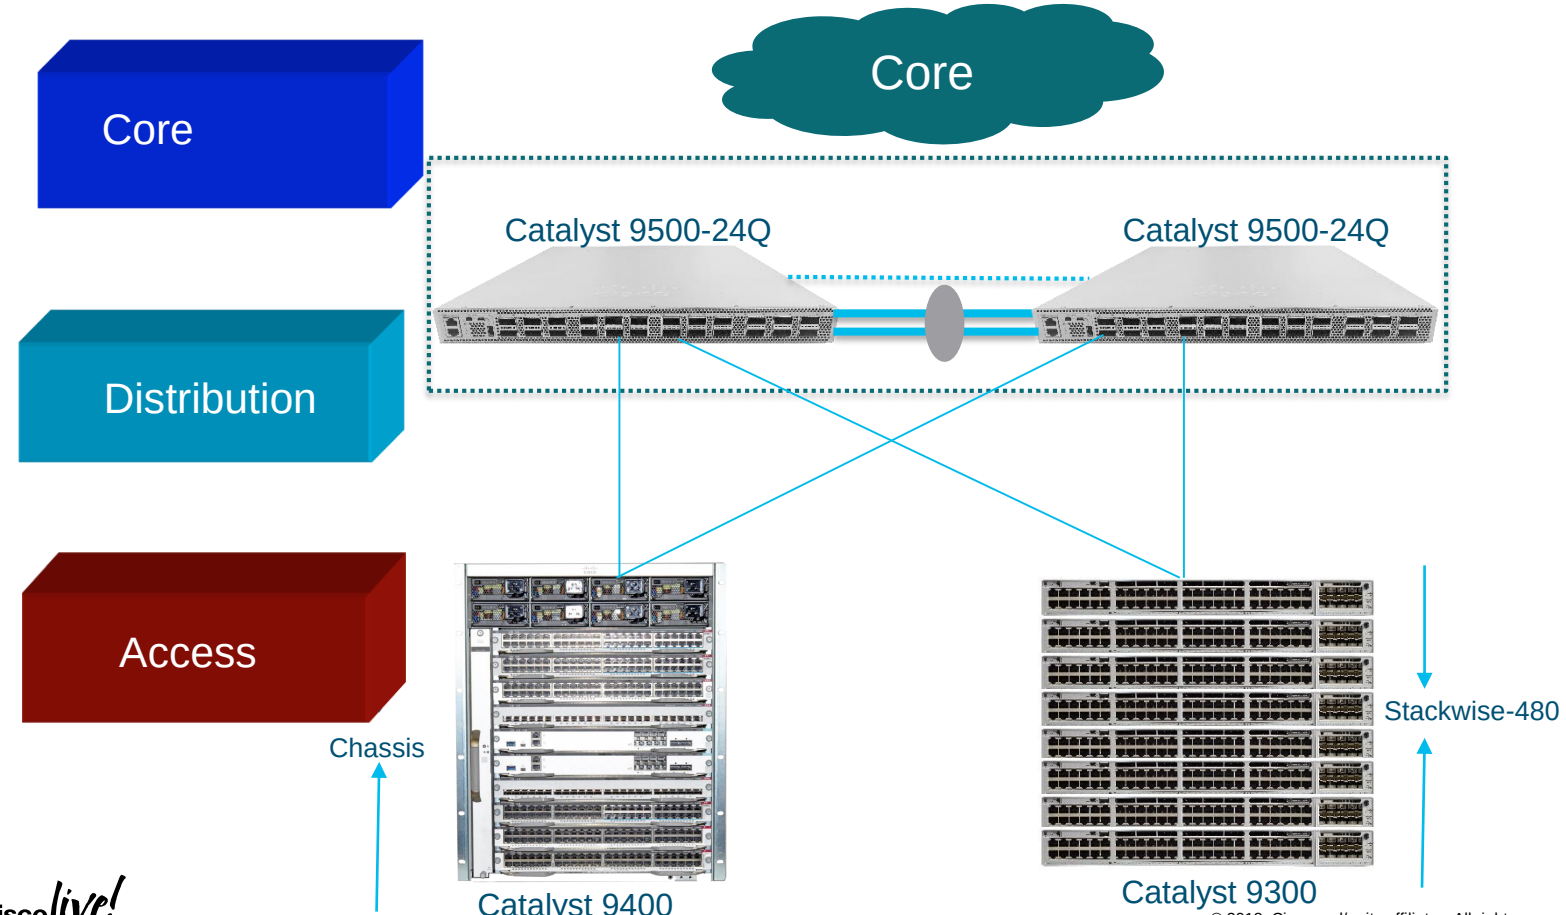

Stackwise-480 Architecture

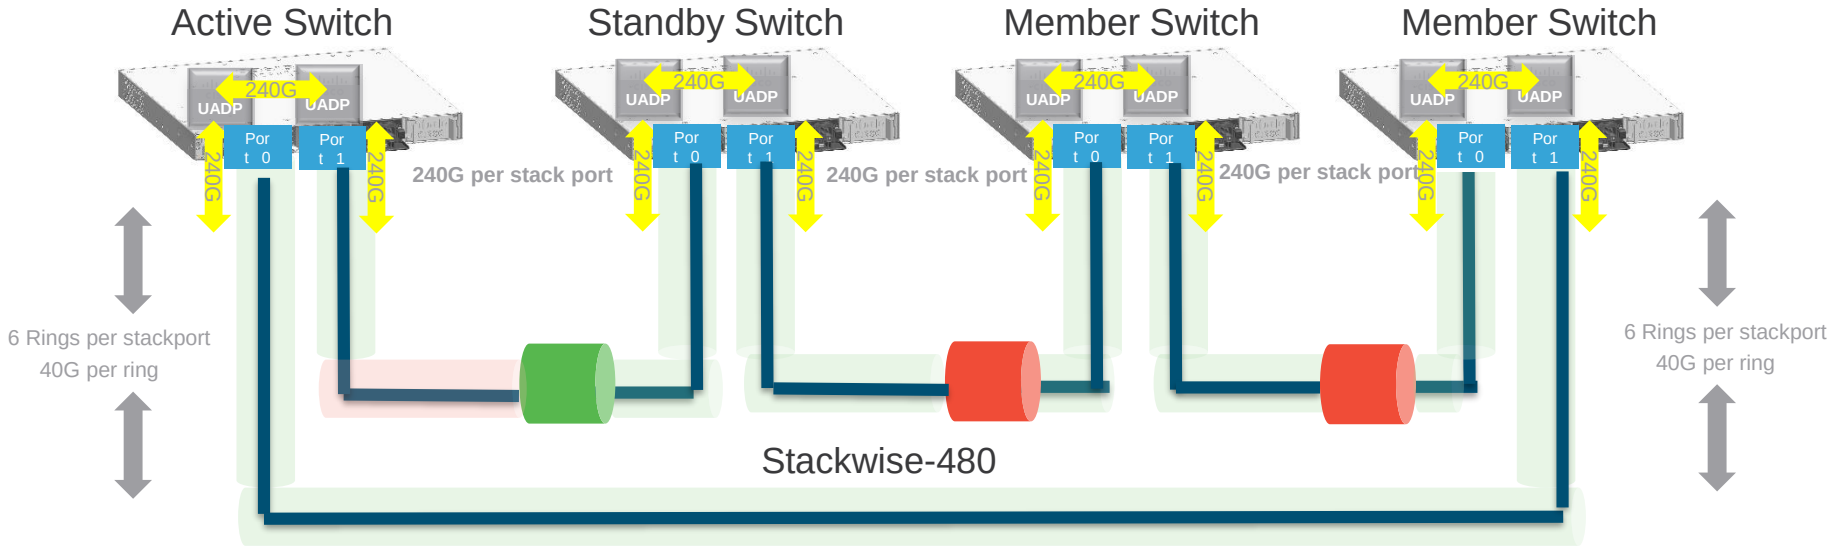

Centralized Control Plane – Scalable Distributed Data Plane

How many Can I stack together?

- Up to 8 Switches can be stacked together using back stacking cables
- All 9300 models are supported in the stack
- All the switches in the stack should run the same IOS and License

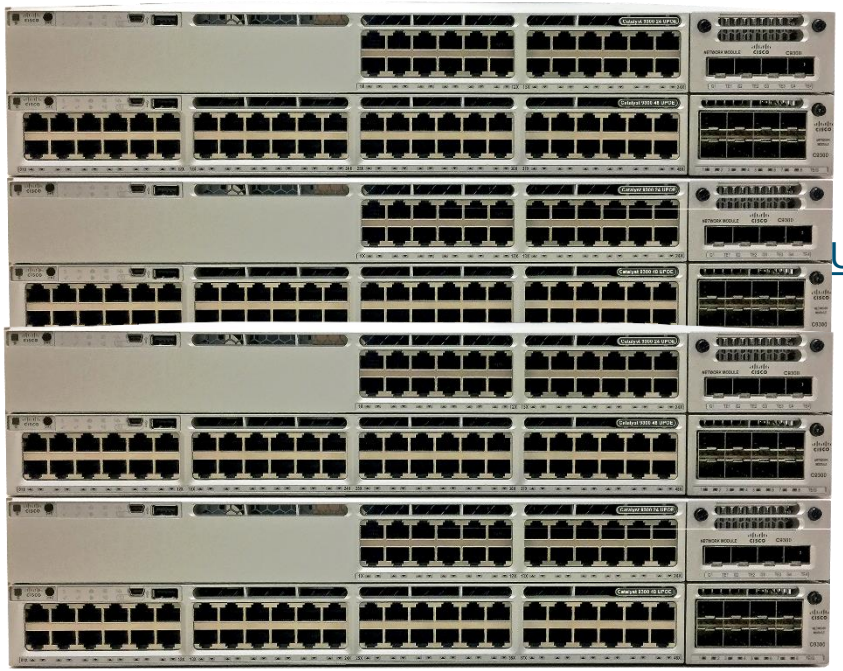

Stack Discovery

- Stack Interfaces brought online
- Infra and LC Domains boot in parallel
- Stack Discovery Protocol discovers Stack topology – broadcast, followed by neighborcast
- In full ring, discovery exits after all members are found.
- In half ring, system waits for 2mins
- Active Election begins after Discovery exits

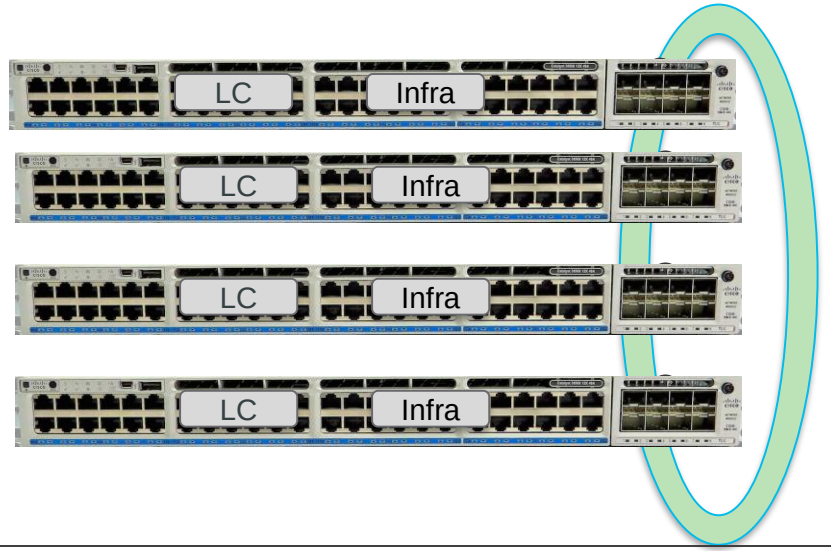


```
Stack port 1 cable is connected and the link is up
Stack port 2 cable is connected and the link is up
Waiting for 120 seconds for other switches to boot
%IOSXE-1-PLATFORM: process stack-mgr: %STACKMGR-1-DISC_START: Switch 3 is starting stack discovery.
##All switches in the stack have been discovered
```

```
Switch number is 3
%IOSXE-1-PLATFORM: process stack-mgr: %STACKMGR-1-DISC_DONE: Switch 3 has finished stack discovery.
%IOSXE-1-PLATFORM: process stack-mgr: %STACKMGR-1-SWITCH_ADDED: Switch 3 has been added to the stack.
```

Stack Active Election

1) The stack (or switch) whose member has the higher user configurable **priority 1–15**

2) The switch or stack whose member has the **lowest MAC address**

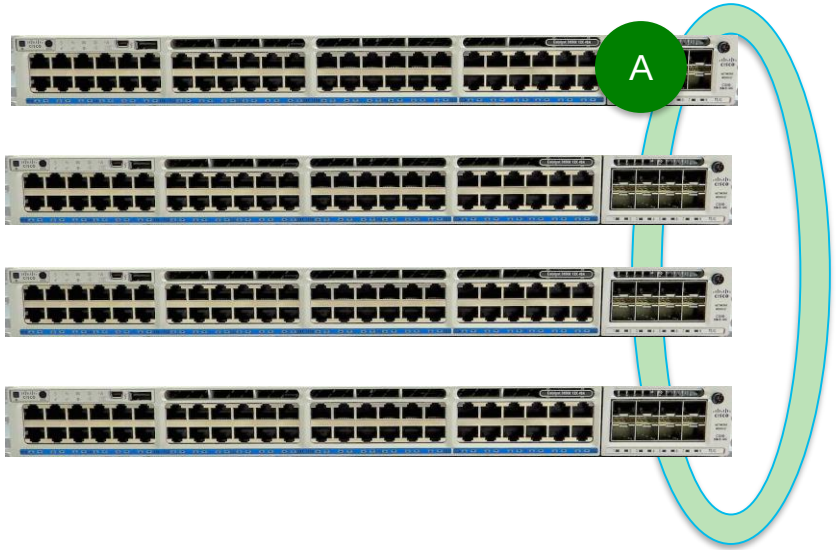


```
%IOSXE-1-PLATFORM: process stack-mgr: %STACKMGR-1-ACTIVE_ELECTED: Switch 3 has been elected ACTIVE.
```

Stack Initialization

- Active starts RP Domain (IOSd, WCM, etc) locally
- Programs hardware on all LC Domains
- Traffic resumes once hardware is programmed
- Starts 2min Timer to elect Standby in parallel
- Active elects Standby
- Standby starts RP Domain locally
- Starts Bulk Sync with Active RP
- Standby reaches “Standby Hot”

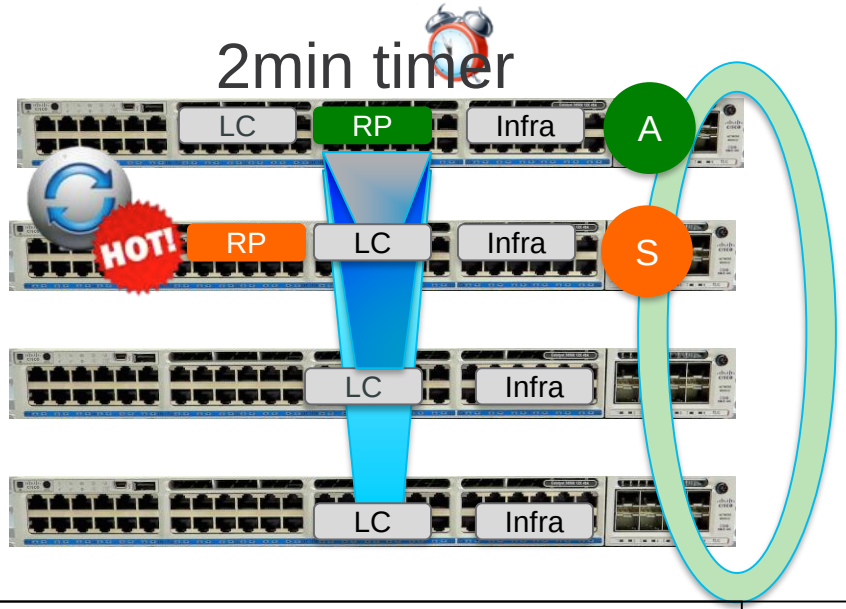


```
%STACKMGR-1-STANDBY_ELECTED: 3 stack-mgr:  Switch 2
has been elected STANDBY.
```

```
GUIDELINE#show switch
Switch/Stack Mac Address : 2037.0652.a580 - Local Mac Address
Mac persistency wait time: Indefinite
                                                     H/W   Current
Switch#   Role    Mac Address     Priority Version  State
-----------------------------------------------------------------
 1       Member   2037.0653.ca80     5       P6A     Ready
 2       Standby  2037.0653.db00     10      P6A     HAsync in progress
*3       Active   2037.0652.a580     15      V01     Ready
```

HA Best Practices & Recommendations

- Power up the first Switch that you want to make it as Active
- Configure Priority of the switch (1-15) – 1 by default – the higher the better
- Power up the second member that you want to make as Standby & then power up rest of the members
- To add a member to an existing stack plug in the stack cable first, then power up the switch
- Avoid stack Merge & Stack split if possible

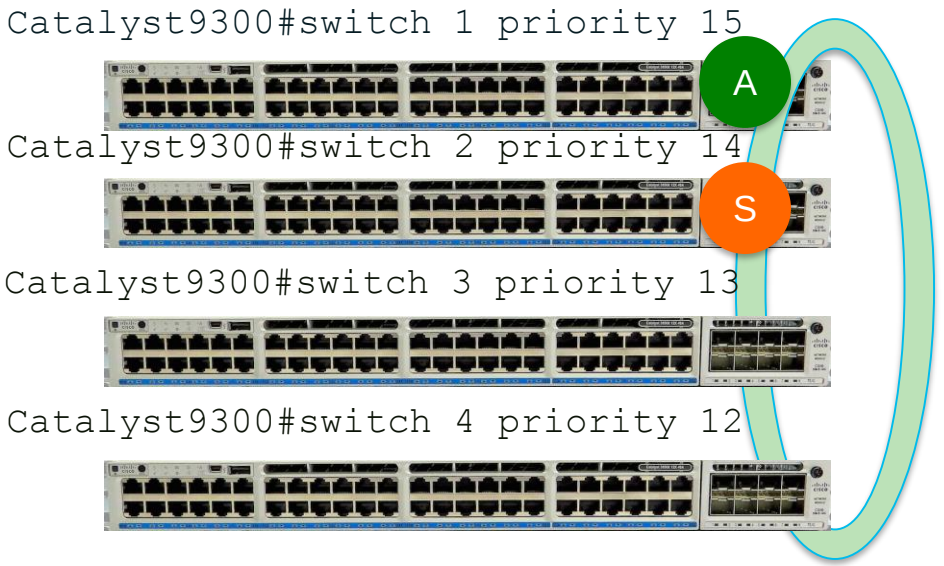

Stack Member Addition

- Stack discovery initiated and completed
- Plug in the member, completing full ring
- Power up the member
- Stack Discovery process runs and completes immediately after discovery happens
- Active detects the new addition, and programs the hardware of the member
- Active is not pre-empted by powering on another member even if it was High Priority

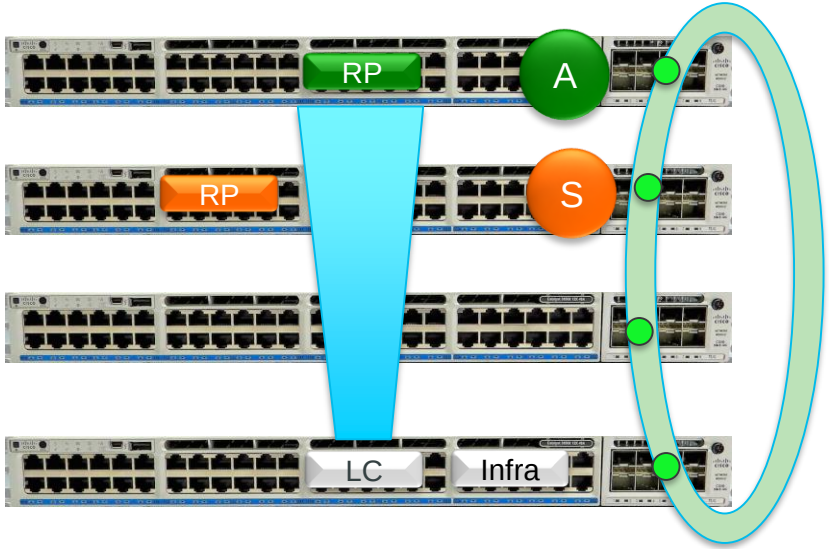

Stack Member Addition – Software Upgrade

- All stack members must have common IOS software version to pair in SSO redundancy state
- Stack member with version mis-match with ACTIVE switch will fail to RPR mode
- Enable “software auto-upgrade enable” command to automate upgrade process
- System must boot in install mode (default and recommended). Auto Upgrade not supported in Bundle Mode

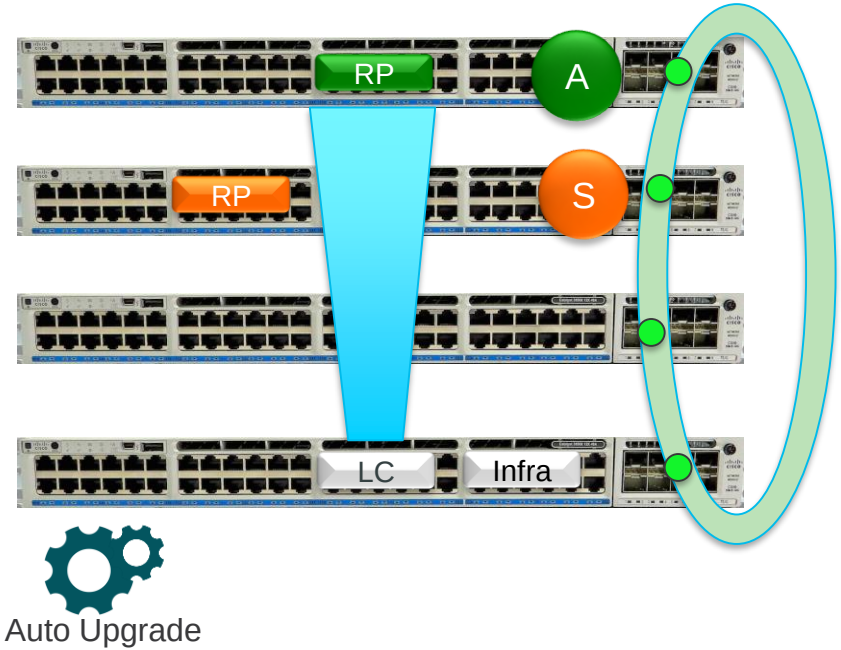

Stack Member Deletion

- Stack discovery initiated and completed
- Active detects member removal – and Clean up process is initiated
- Clean-up involves removing TCAM entries referencing removed member, MAC addresses, CDP tables – more like all ports on the member are shutdown
- Configuration is moved to Pre-Provisioned state

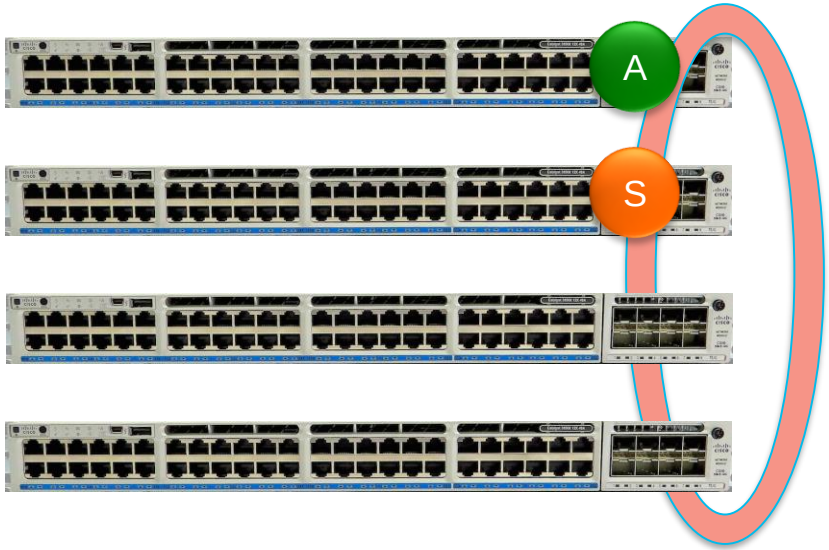

Show switch with SSO

```
Switch# show switch
Switch/Stack Mac Address : 2037.06cf.0e80
                                                    H/W   Current
Switch#   Role     Mac Address     Priority Version  State
------------------------------------------------------------------
*1        Active   2037.06cf.0e80     10      V01      Ready
 2        Standby  2037.06cf.3380     8       V00      Ready
 3        Member   2037.06cf.1400     6       V00      Ready
 4        Member   2037.06cf.3000     4       V00      Ready
```

Stack Mac follows Active initially

Active

Standby

Member

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

Show switch detail output

```
Switch# show switch detail
Switch/Stack Mac Address : 2037.06cf.0e80
                                                    H/W   Current
Switch#   Role    Mac Address     Priority Version  State
*1       Active   2037.06cf.0e80     10     V01     Ready
 2       Standby  2037.06cf.3380     8      V00     Ready
 3       Member   2037.06cf.1400     6      V00     Ready
 4       Member   2037.06cf.3000     4      V00     Ready


          Stack Port Status             Neighbors
Switch#   Port 1     Port 2           Port 1   Port 2
-----------------------------------------------------
- 1          OK          OK             2        4
  2          OK          OK             3        1
  3          OK          OK             4        2
  4          OK          OK             1        3
```

Stackport Information

Show redundancy states

```
Switch# show redundancy states
         my state = 13 -ACTIVE
       peer state =  8 -STANDBY HOT
             Mode = Duplex
          Unit ID = 1

Redundancy Mode (Operational) = SSO
 Redundancy Mode (Configured) = SSO
             Redundancy State = SSO
                    Manual Swact = enabled

 Communications = Up

            client count = 76
client_notification_TMR = 360000 milliseconds
          keep_alive TMR = 9000 milliseconds
         keep_alive count = 0
     keep_alive threshold = 9
            RF debug mask = 0
```

Terminal state for Active Unit.

Terminal state for Standby Unit for SSO.

Slot Number of Active Unit

Communication Channel Status between the Active/ Standby RP units

Show Redundancy Command Output…

```
Switch#sh redundancy
Redundant System Information :
------------------------------
       Available system uptime = 29 weeks, 2 days, 11 hours, 47 minutes
Switchovers system experienced = 2
              Standby failures = 0
        Last switchover reason = user_forced

                 Hardware Mode = Duplex
    Configured Redundancy Mode = SSO
     Operating Redundancy Mode = SSO
              Maintenance Mode = Disabled
                Communications = Up

Current Processor Information :
------------------------------
               Active Location = slot 1
        Current Software state = ACTIVE
        Uptime in current state = 1 week, 4 days, 22 hours, 38 minutes
        Image Version = Cisco IOS Software, IOS-XE Software, Catalyst L3 Switch Software (CAT3K_CAA-UNIVERSALK9-M),
Version 03.03.03E RELEASE SOFTWARE (fc1)

Peer Processor Information :
------------------------------
              Standby Location = slot 2
        Current Software state = STANDBY HOT
        Uptime in current state = 1 week, 4 days, 22 hours, 34 minutes
        Image Version = Cisco IOS Software, IOS-XE Software, Catalyst L3 Switch Software (CAT3K_CAA-UNIVERSALK9-M),
Version 03.03.03E RELEASE SOFTWARE (fc1)
```

Show Redundancy Command Output…

```
Switch#sh redundancy
Redundant System Information :
------------------------------
       Available system uptime = 29 weeks, 2 days, 11 hours, 47 minutes
Switchovers system experienced = 2
             Standby failures = 0
        Last switchover reason = user_forced

                 Hardware Mode = Duplex
    Configured Redundancy Mode = SSO
     Operating Redundancy Mode = SSO
              Maintenance Mode = Disabled
                Communications = Up

Current Processor Information :
------------------------------
               Active Location = slot 1
        Current Software state = ACTIVE
        Uptime in current state = 1 week, 4 days, 22 hours, 38 minutes
        Image Version = Cisco IOS Software, IOS-XE Software, Catalyst L3 Switch Software (CAT3K_CAA-UNIVERSALK9-M),
Version 03.03.03E RELEASE SOFTWARE (fc1)

Peer Processor Information :
------------------------------
             Standby Location = slot 2
        Current Software state = STANDBY HOT
        Uptime in current state = 1 week, 4 days, 22 hours, 34 minutes
        Image Version = Cisco IOS Software, IOS-XE Software, Catalyst L3 Switch Software (CAT3K_CAA-UNIVERSALK9-M),
Version 03.03.03E RELEASE SOFTWARE (fc1)
```

System uptime

Image version of current unit

StackPower - Overview

- Provides RPS functionality with Zero RPS Footprint
- Pay-as-you-grow architecture – similar to the Data Stack
- 1+N Redundancy with Inline Power
- Up to 4 Switches in a StackPower
- Multiple StackPower Possible within one Data Stack
- Flexible Installation, Better Efficiency

How StackPower Works?

- Pools Power from All PS
- All Switches in StackPower share the available Power in Pool
- Each Switch is given their Minimum Power Budget

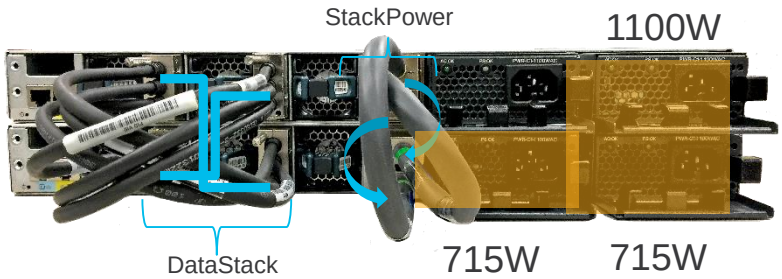

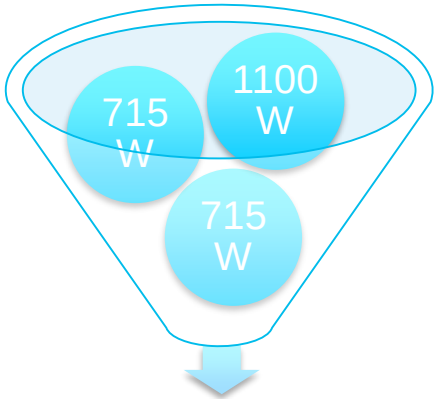

Total Input Power 2530W

Power Budget Modes

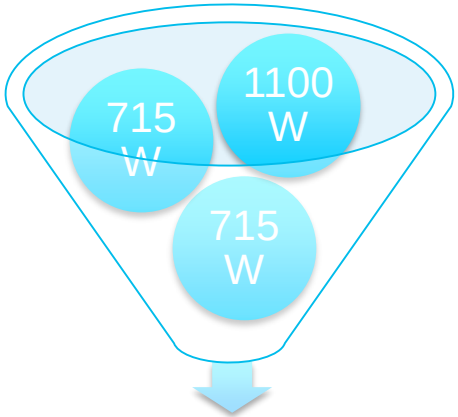

2530W – 30W

Power Sharing Mode

- The Default Mode
- Sum of All PS – 30~60W

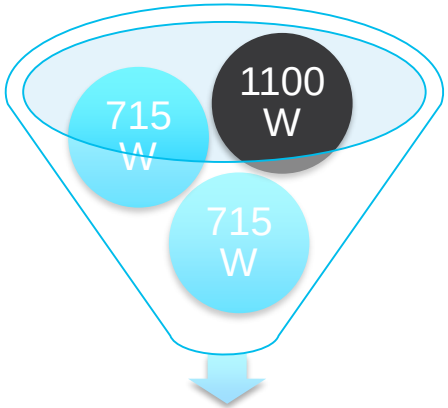

1430W – 30W

Redundant Mode

- User Configurable
- Sum of All PS – Largest PS - 30~60W

Global StackPower Reserve = 30W

Enforcement Modes

Strict & Loose Modes Control The Behavior of Load Shed

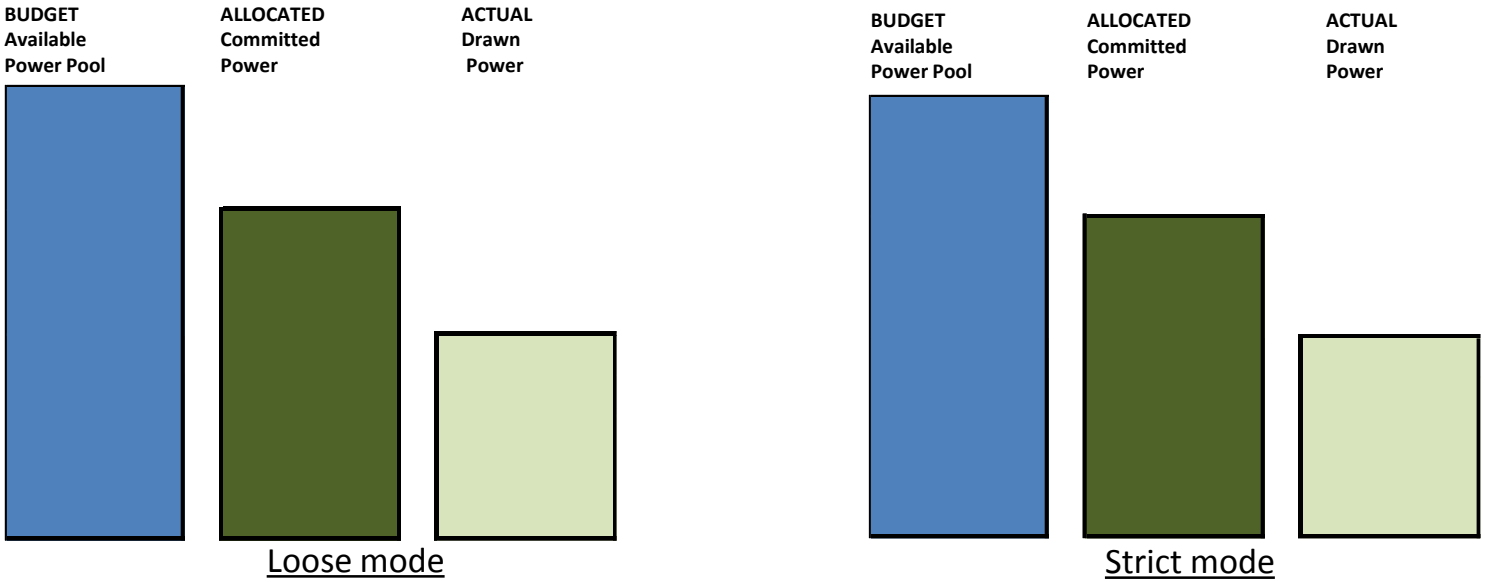

- Loose mode allows for a negative power budget
- Strict mode sheds load as soon as the power budget goes below the Allocated power level

Best Practices for Stackables

Define Stack Roles with minimal Downtime

Simplify Network Operations

- Power up the first Switch that you want to make it as Active
- Configure Priority of the switch
- (1-15) – 1 by default – the higher the better
- Power up the second member that you want to make as Standby
- Configure Priority less than the Active
- Power up the rest of the members
- Configure Priorities on those units

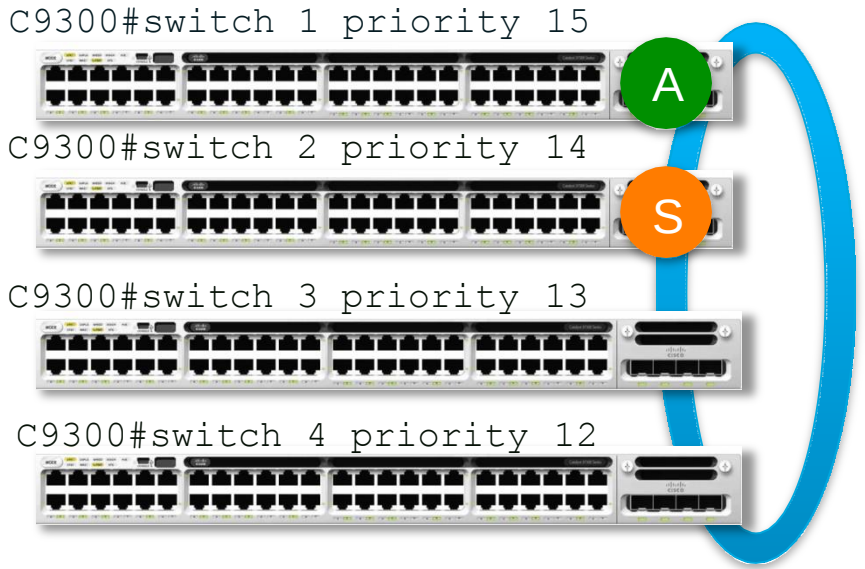

Stacking Convergence

Multi-Layer Access

- Active unit with uplink failure introduces two failures
  - Active control plane
  - Uplink interface
- When the Active fails,
  - the Standby will take over.
- Upstream, HSRP / GLBP will detect link down, and
  - D2 will start answering to the virtual MAC 0000.0c07.ac00
- Downstream traffic is
  - re-routed to D2 via L3 link

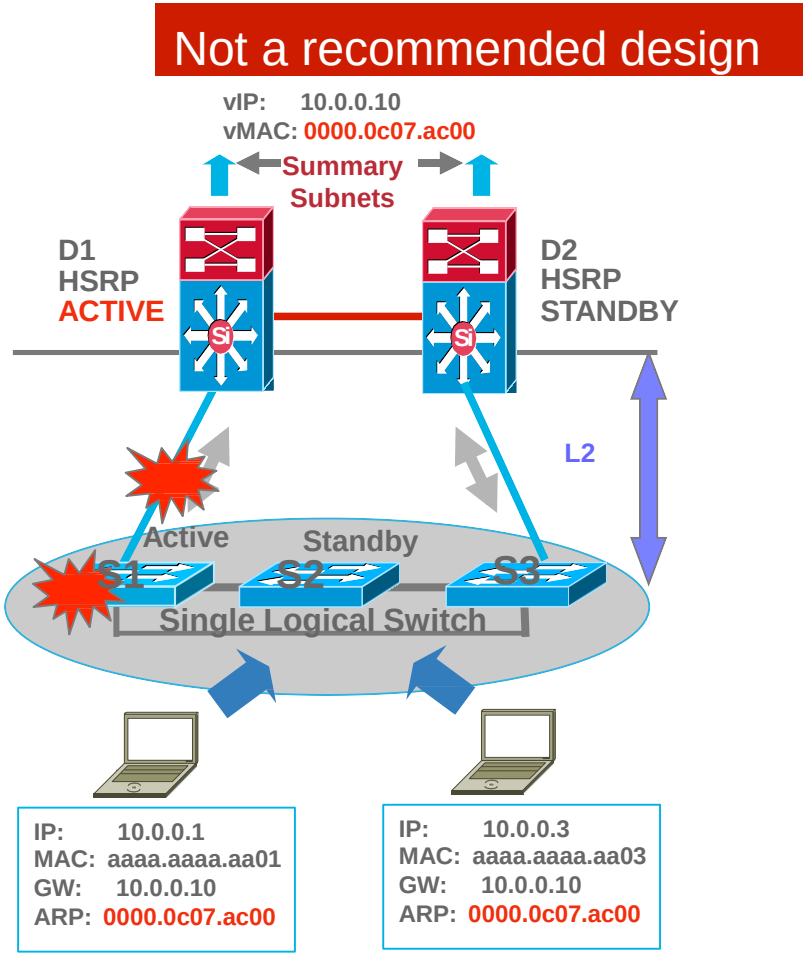

Stacking Convergence

Multi-Layer Access

- Active unit Failure (without uplink)
- When the Active fails,
  - the Standby will take over
  - No HSRP/GLBP failover, while the new Active being elected, MAC address of HSRP/GLPB still used by the rest of the stack for data forwarding
- No downstream
  - re-route convergence

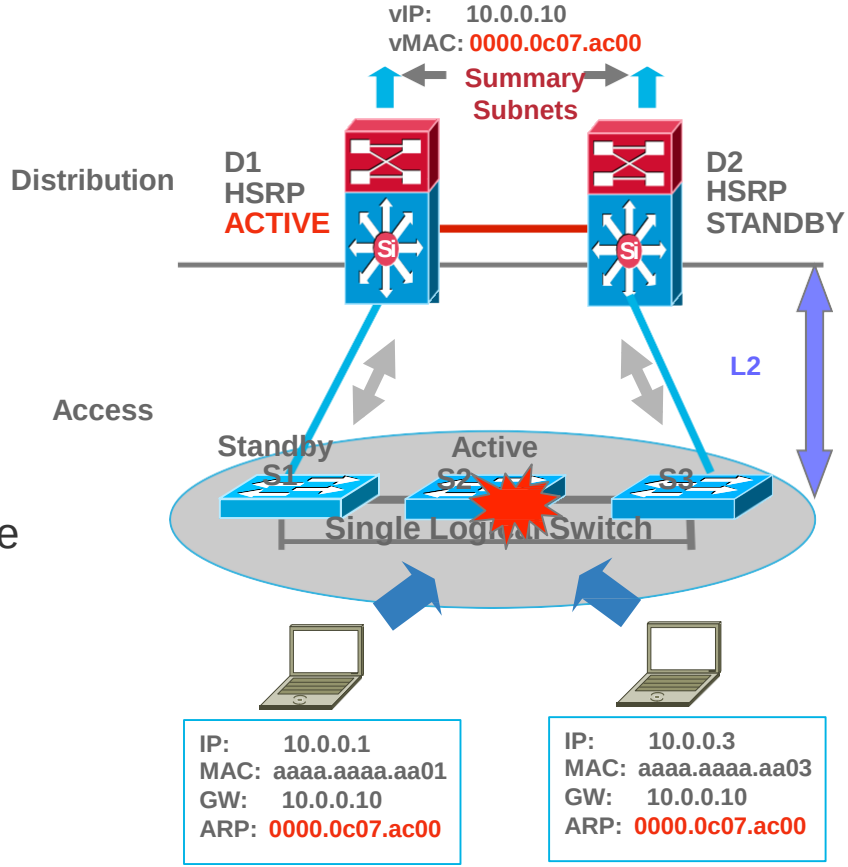

Catalyst 9300 StackWise

Routed Access

- CLI "stack-mac persistent timer 0" enables MAC consistency –
  - This is the default value for 9300
  - This is a change from the existing stacking model
  - New Active inherits the MAC address of the previous Active
  - No MAC changes for end hosts and adjacent routers, significantly improves upstream recovery
- **Caution –**
  - Do not re-introduce the 3x50 **elsewhere** in order to avoid duplicate MAC in your network

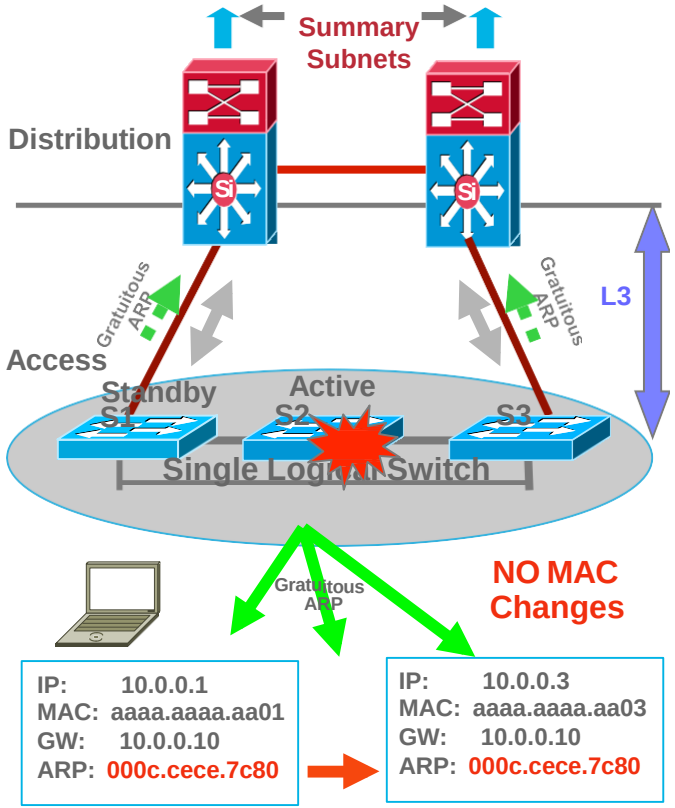

Changing Stack Mac on 9300 Stack

- By default the timer value is set to indefinite (0)
  - System continues to keep selected stack mac after switchover
  - Avoids Protocol flapping
- How to change it
  - A new command introduced

```
9300-1#stack-mac update force
```

```
9300-1#show switch
```

```
9300-1#show switch
Switch/Stack Mac Address : 2037.06cf.3380
Mac persistency wait time: Indefinite

                                                     H/W   Current
Switch#   Role    Mac Address     Priority Version  State
------------------------------------------------------------------
*1        Member  0000.0000.0000     10     V01      Removed
 2        Active  2037.06cf.3380     8      V00      Ready
 3        Member  2037.06cf.1400     6      V00      Ready
 4        Member  2037.06cf.3000     4      V00      Ready
```

Key Recommendations for Stacking

- Run the stack in full ring mode to get full bandwidth
- Configure the Active switch priority and Standby switch priority
  - Predetermine which switch is the Active and Standby which will become the Active should the Active fail
  - Simplifies operations
- Configure Active and Standby unit without uplinks if possible
  - If deploying a stack of 4 or more switches keep the Active and Standby switches without uplinks, this will simplify the convergence and reduce the outage time
- Do Not change the stack-mac timer value
  - By default the value is 0 (indefinite)
  - Avoids protocol flapping
  - There is a command to change the stack-mac when needed

Fast Software Upgrade

Achieving High Availability on Catalyst Access Switch

Fast Software Upgrade

- FSU provides a mechanism to upgrade and downgrade the software image with minimal impact to the Data Traffic
- During Software upgrade, Control plane is updated first and only then Data Plane is reset impacting the data traffic
- Traffic impact is almost reduced by half compared to regular upgrades

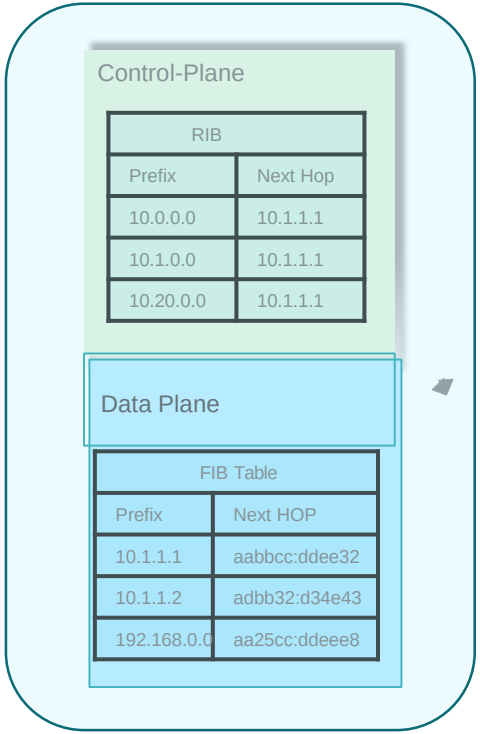

Fast Software Upgrade

Normal Upgrade Impact

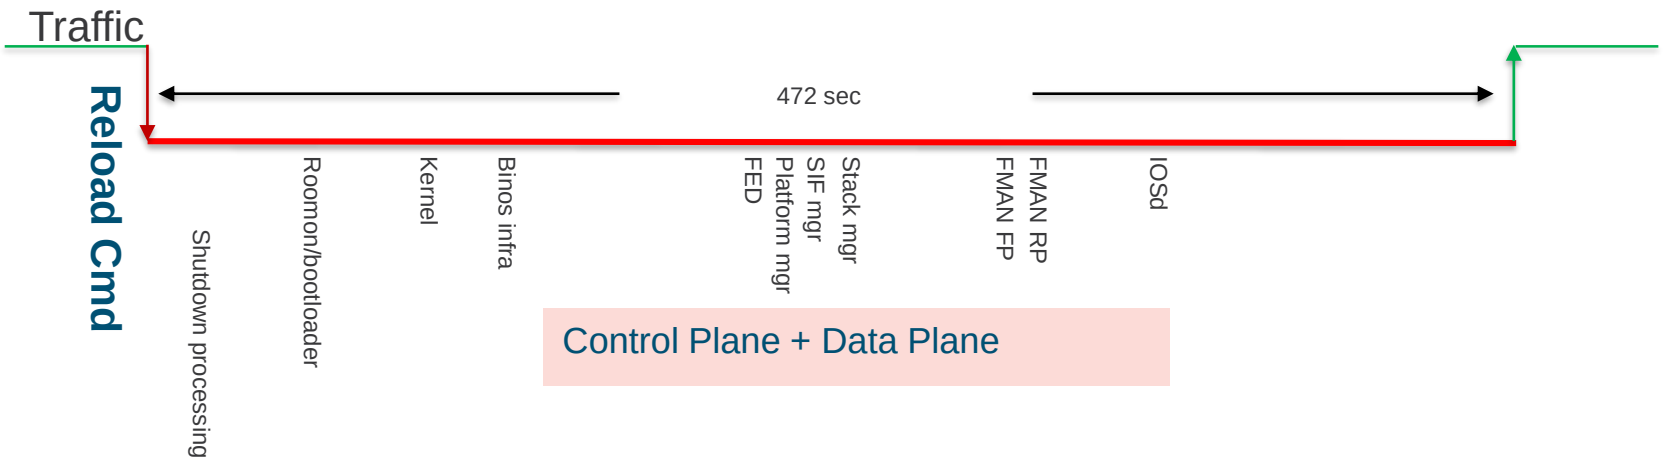

Fast Software Upgrade

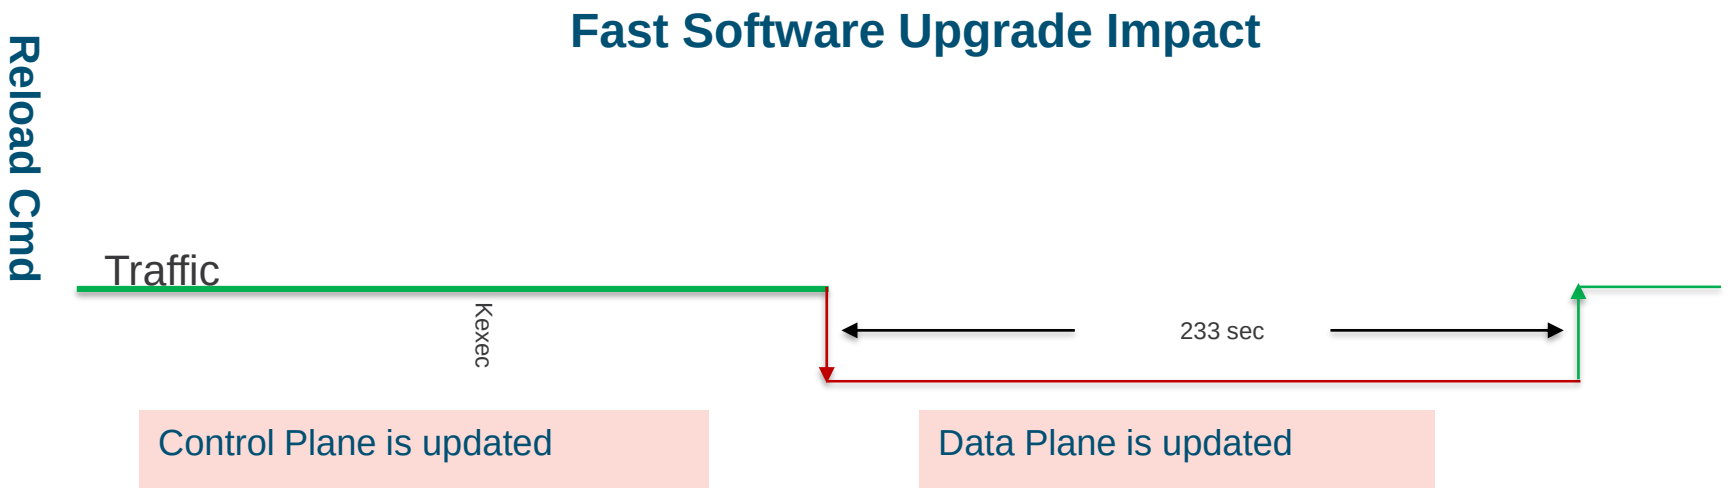

Fast Software Upgrade

FSU Workflow

The process starts by copying the images to the switch

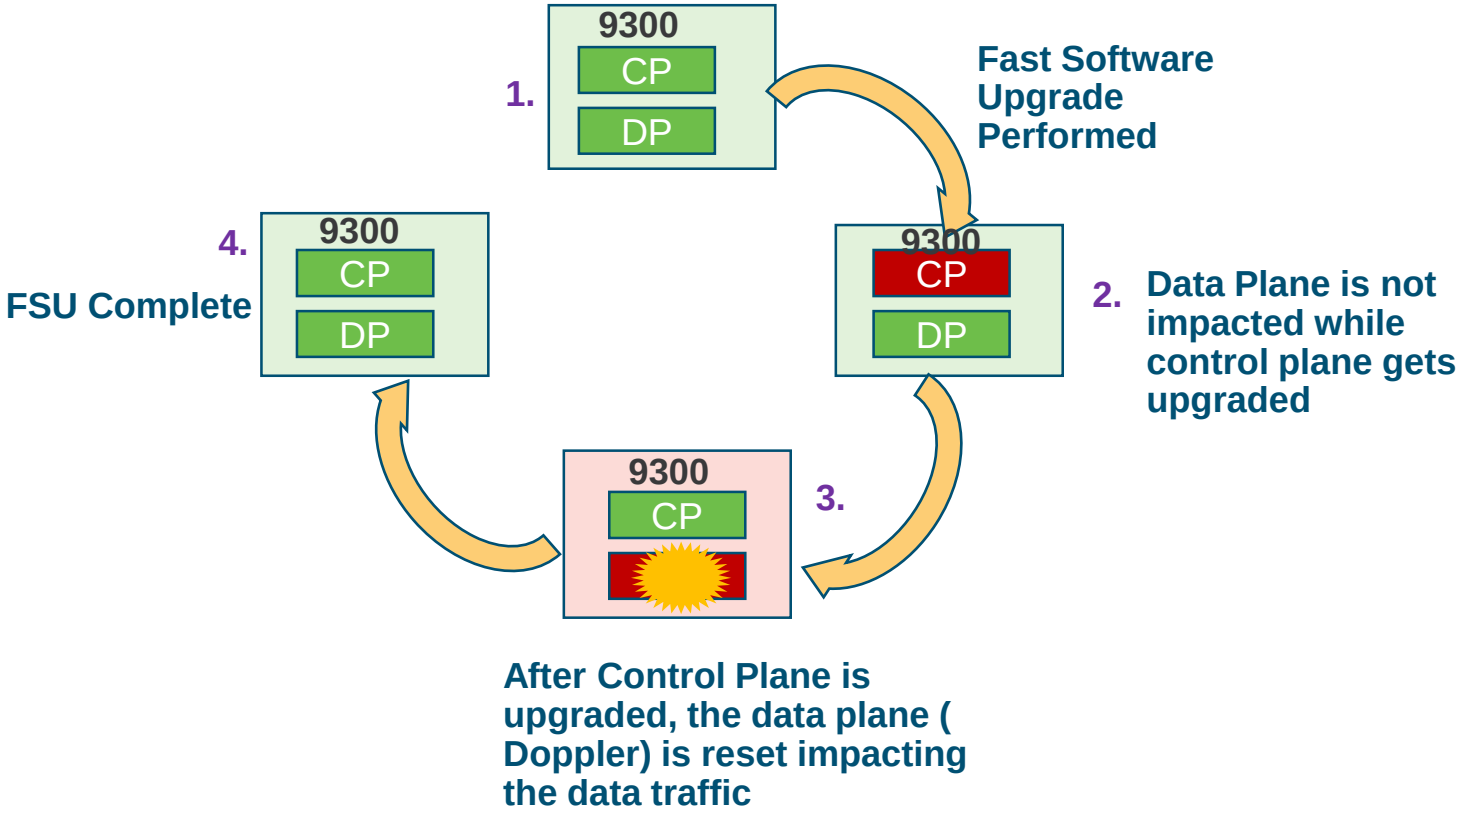

Fast Software Upgrade
Restrictions

- FSU is not supported in Routed Access Topology
- Switch cannot have more than one forwarding port and hence is only suitable for Access layer
- FSU is not supported on Stackwise Virtual

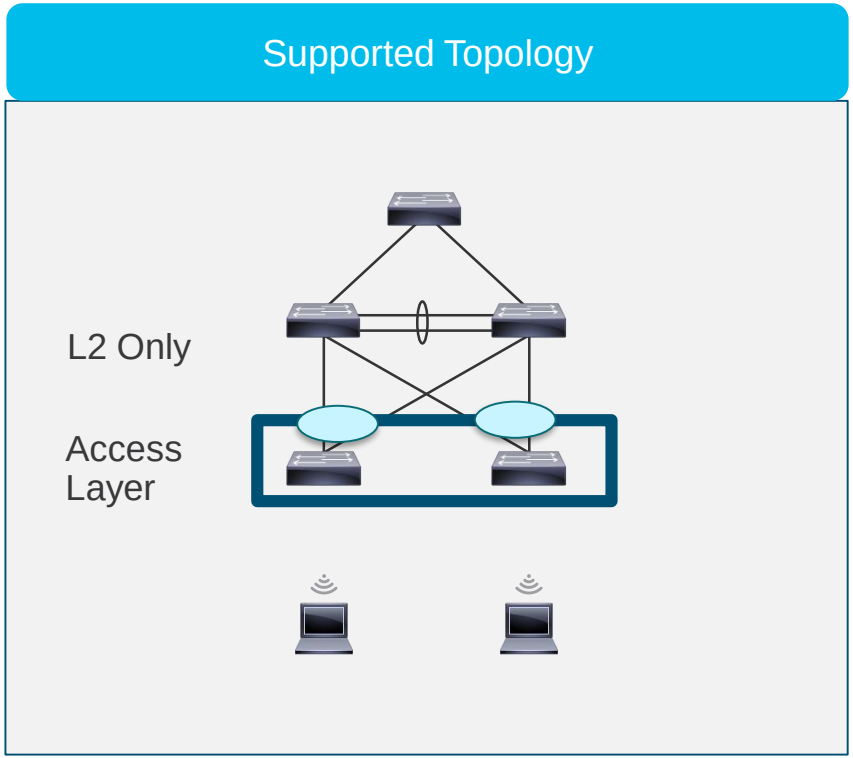

Fast Software Upgrade

CLI Commands

- FSU is supported only in install mode
- One step command which activates the fast software upgrade and commits it

```
install add file flash:cat9k_iosxe.BLD_V168 activate reloadfast commit
```

- Fast Reload without Software upgrade

```
Reload Fast
```

Agenda

- High Availability Overview and Evolution
- High Availability Solution on the Campus Access
  - Stackable High Availability Solution
  - Modular High Availability Solution
- High Availability Solution on the Campus Distribution
- High Availability Solution in IOS
- Summary/Q&A

Catalyst 9400 High Availability (Modular)

- There are two dedicated supervisors
- Switch Boots Up
- Reads registers on backplane to determine the inserted card types
- Active Election begins after Discovery exits
- Active Supervisor asserts mastership in the HW
- Other Supervisor will become Standby

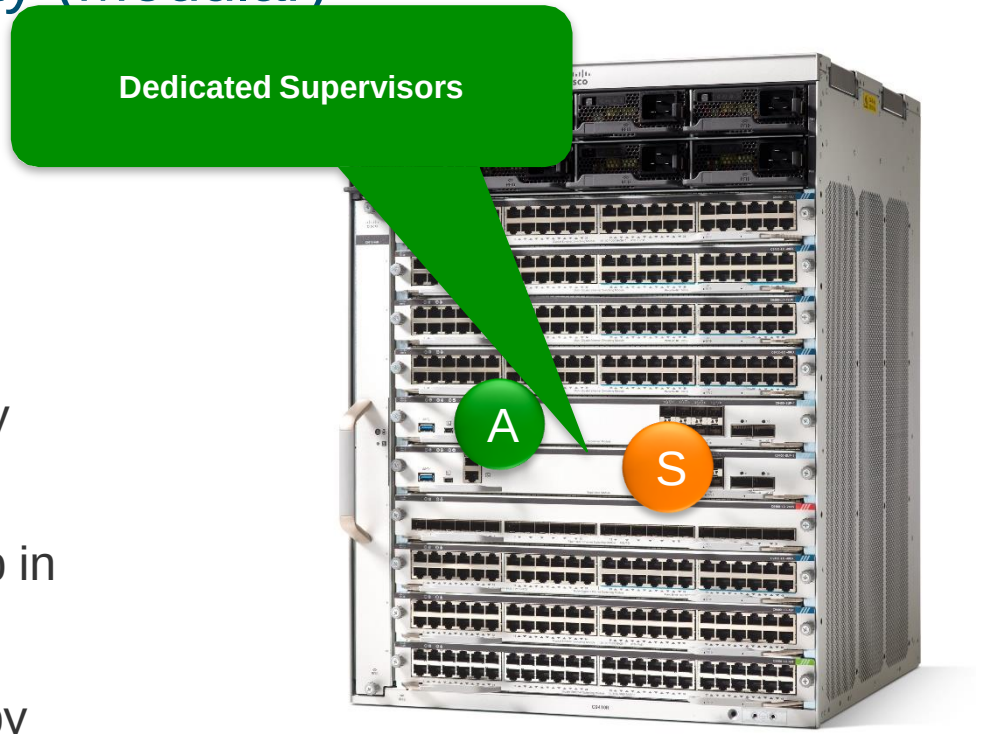

Catalyst 9400 High Availability State Machine

- Active starts its software processes
- Standby starts its software processes
- **Active Programs the local Sup HW**
- Standby Start Bulk sync with Active RP
- Standby Reaches “Standby Hot”
- **Standby Programs the local Sup HW**

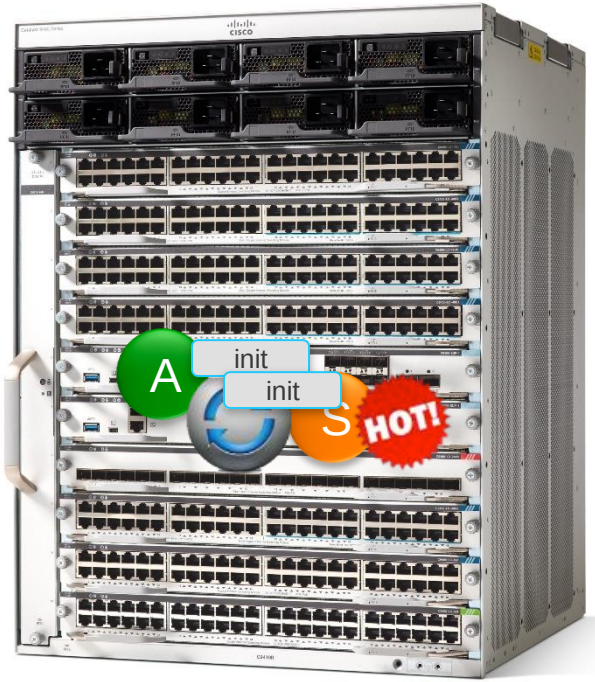

Best Practices for Modular

Uplinks on the Catalyst 9400

- Two redundant 40 Gig uplinks active between the supervisors and actively forwarding at the same time
- The Active Sup controls all the uplinks on both itself and the Standby unit
- Two wire speed 40 Gig uplinks are active all the time even if one of the units fail.
- **Recommendation** – Connect uplinks on different Supervisors

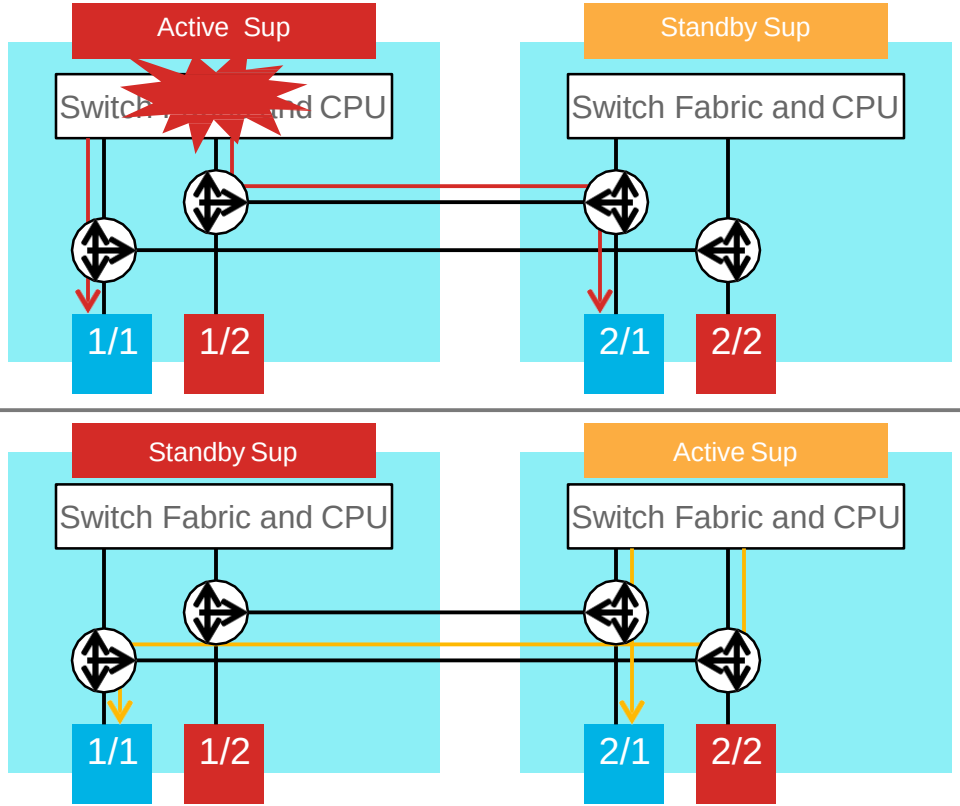

Power Supply Redundancy on the Catalyst 9400

- Has Eight power supply bays that can be run in redundant or combined mode.

- Power Supply configuration modes
- **Combined Mode:** supervisor engines manages the combined power budget of both units
  - Only used for powering POE devices
  - that require more power than the 1 supply can provide

- **Redundant Mode**: One or N Power supplies are standby and remaining are active

Power Supply Redundancy Modes

	Normal	PS failure
Combined (Default)	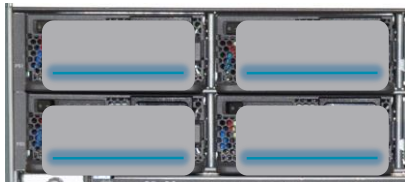 Load sharing on all PSs	 Load sharing on functional PSs
Redundant	 Load sharing on active PSs Standby PS in output disabled	 Standby PS becomes active System enters alarm state

Failed PS

ACTIVE

STANDBY

Power Redundancy: N+1 and N+N

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

- Default active is PS1-7 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

Key Recommendations for Modular

- Redundant Supervisors for better Availability
- Split the Uplinks between the Active and Standby units in a redundant system
  - All uplinks are Actively forwarding traffic
  - Active Supervisor will control all uplinks even if the other unit is failed
- Power Redundancy
  - Default is redundant Power mode
  - Choose the combined mode for running POE devices requiring more power than 1 supply can provide

Agenda

- High Availability Overview and Evolution
- High Availability Solution on Campus Access
  - Stackable High Availability Solution
  - Modular High Availability Solution
- High Availability Solution on Campus Distribution
- High Availability Solution in IOS
- Summary/Q&A

Importance of High Availability for Access switches

Feature and Device rich layer

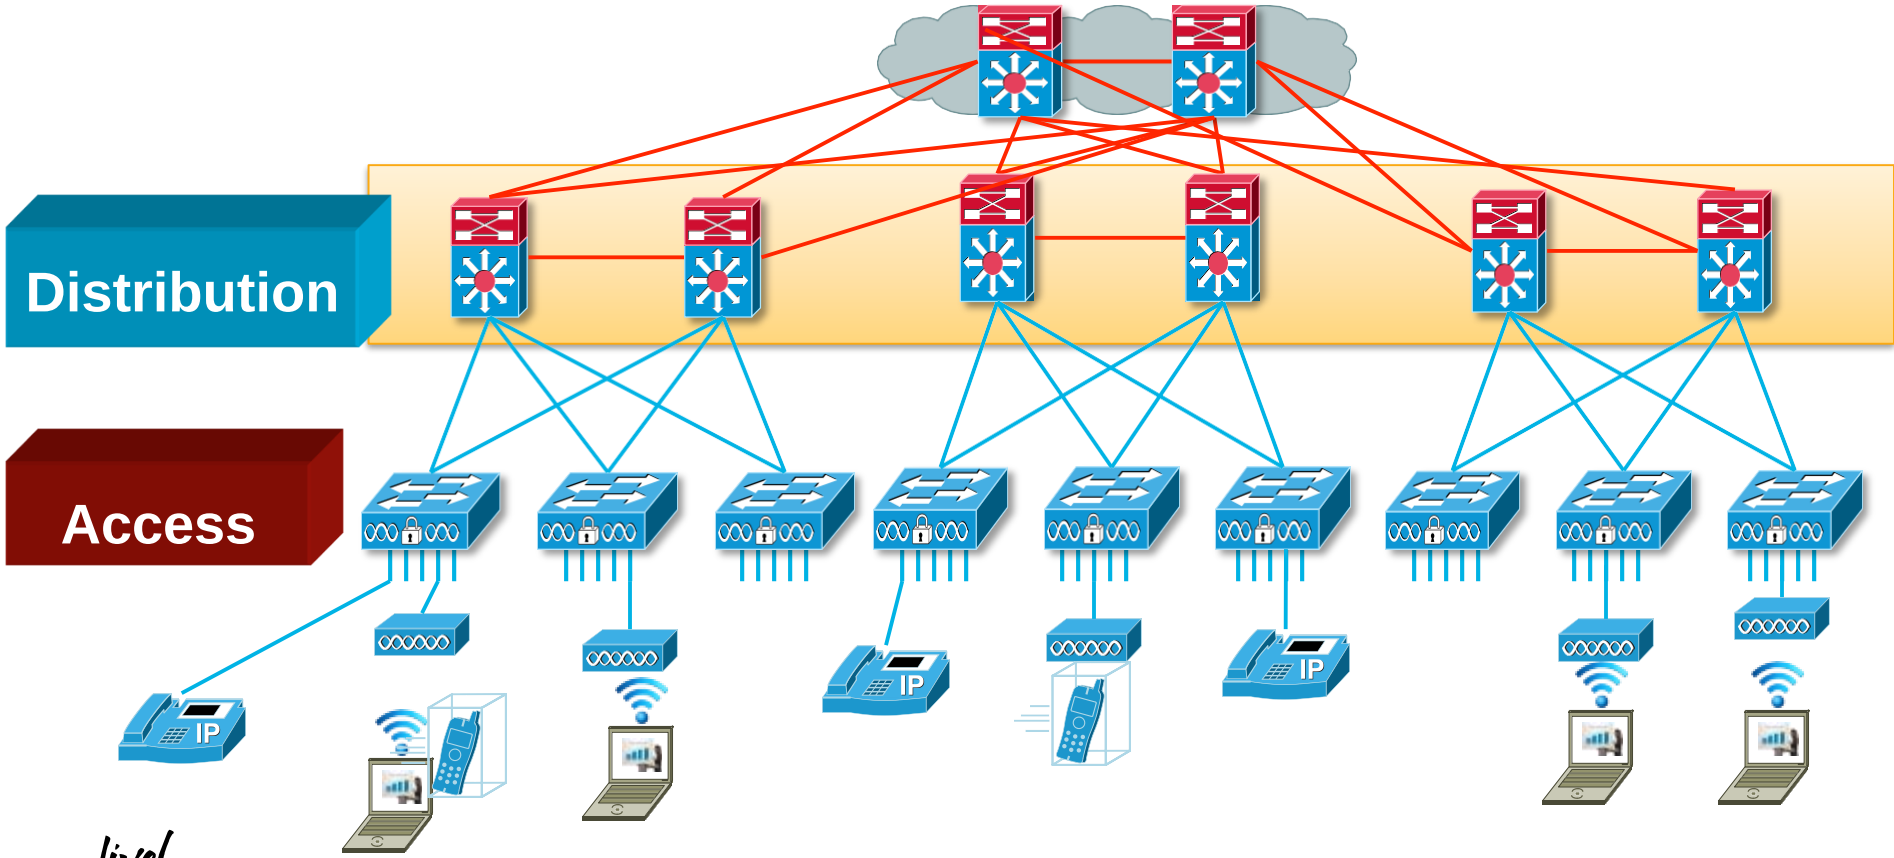

Catalyst 9500: Optimized for Enterprise Deployments

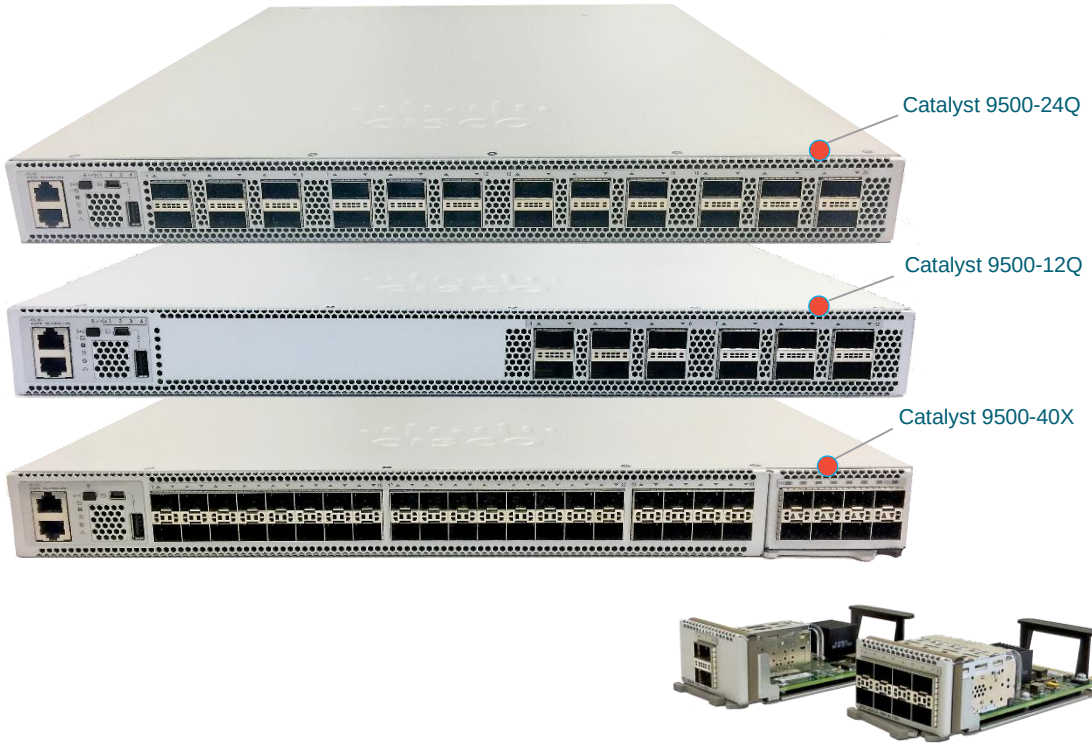

Stackwise Virtual

Achieving High Availability in Distribution Layer

Providing HW redundancy with Stackwise Virtual

- **Unified Control Plane**
  - Manage, Configure and troubleshoot two switches as a single switch

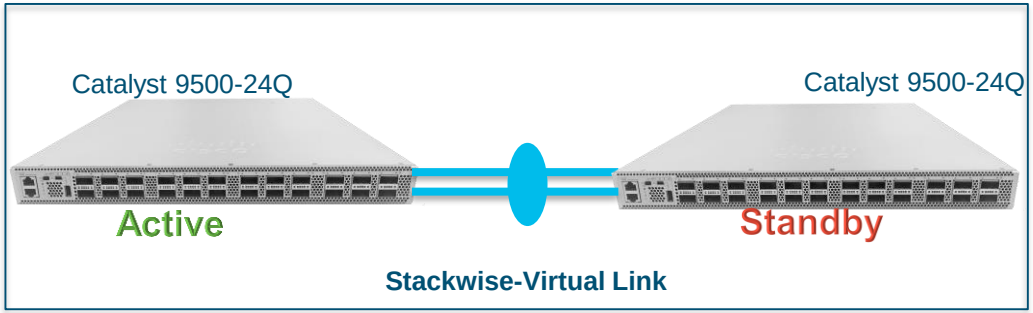

Achieving High Availability in Distribution Layer

Providing HW redundancy with Stackwise Virtual

- **Active/Active Data Plane**
  - Both the switches are capable of forwarding the traffic locally

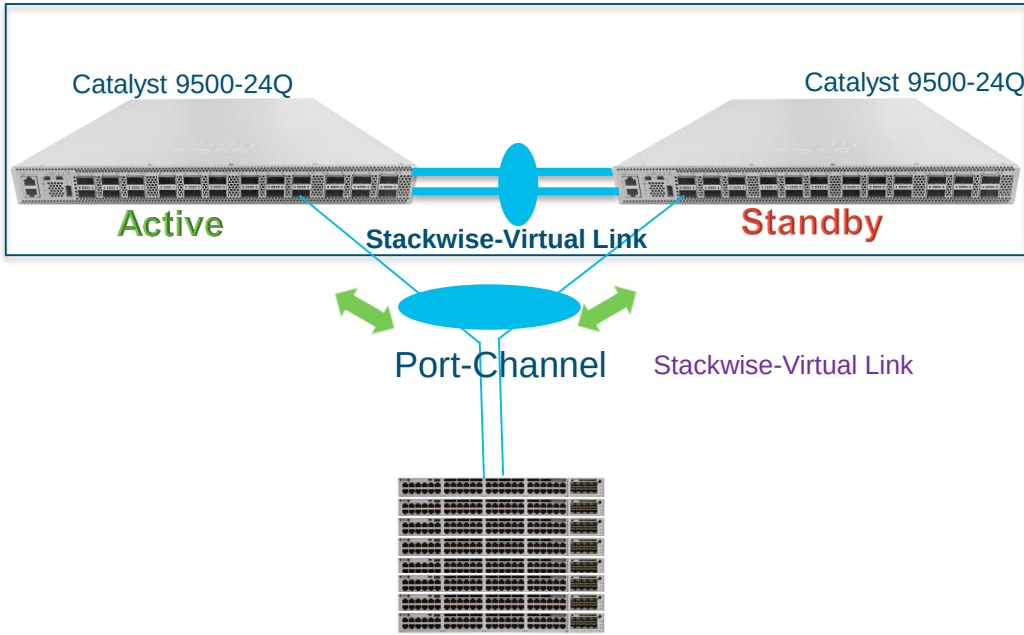

Achieving High Availability in Distribution Layer

Stackwise Virtual Components

- **Stackwise Virtual Link**
  - Dedicated Stacking Link facilitating communication between the switches
- **Dual Active Detection Link**
  - Dedicated Connection to check and avoid dual-active scenario
- **Multi-Chassis Ether-channel**
  - Port-Channel Spanning across Stackwise virtual switches
  - L2 and L3 Port-channels

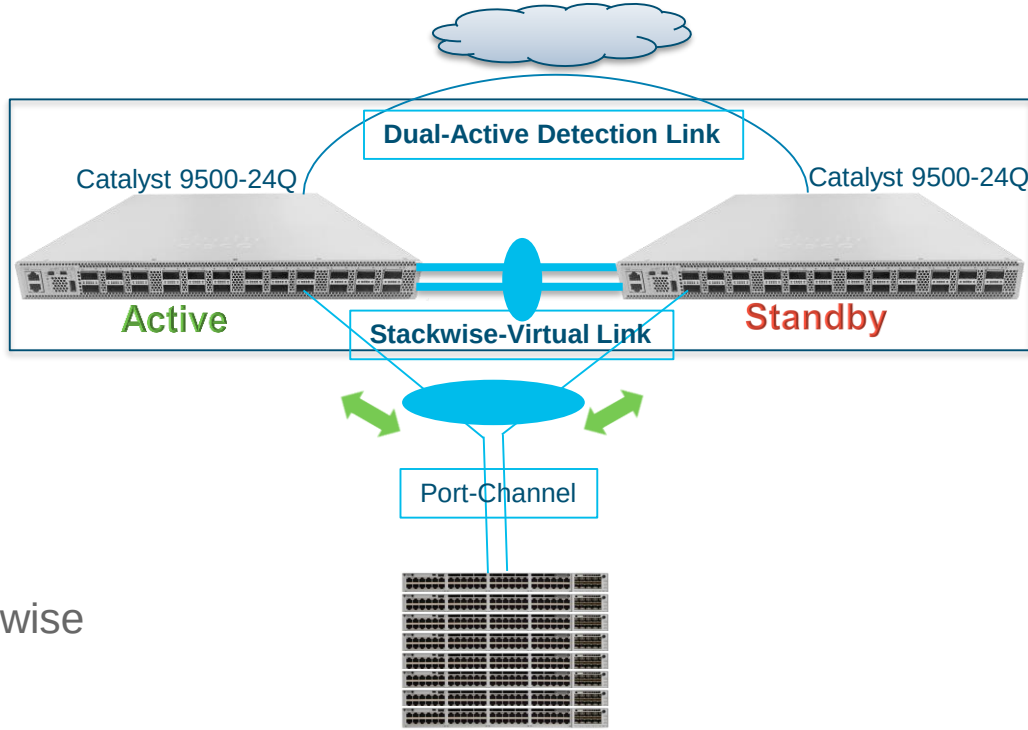

Achieving High Availability in Distribution Layer

Stackwise Virtual Link

- **Inter-Chassis System Link**
  - No network protocol operations
  - Invisible in network topology
  - Transparent to network level troubleshooting
- **SVL Control Link**
  - Carries all system internal control traffic
  - Single member-link and dynamic election during bootup
  - Shared interface for network/data traffic
- **Payload Overhead**
  - Every single packet encapsulated with 64B of StackWise Virtual Header (SVH)
  - Non-bridgeable and Non-routeable.
  - SVL must be directly connected between two stack-member switch systems

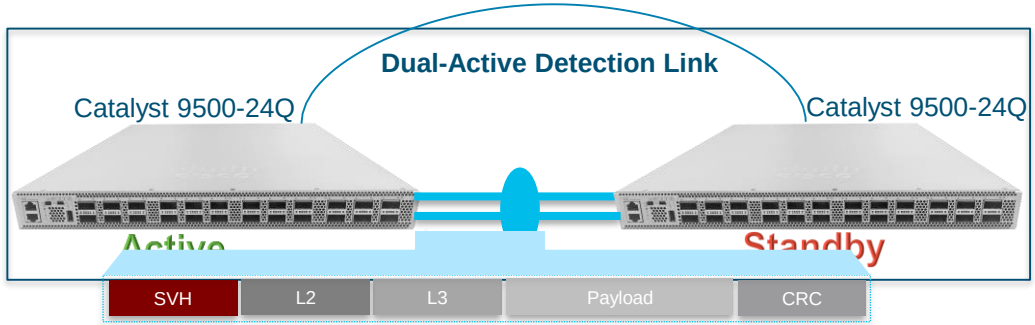


```
Dist-1# show stackwise-virtual link

Stackwise Virtual Link(SVL) Information:
-----------------------------------------------------
Flags:
--------
Link Status
---------------
U-Up D-Down
Protocol Status
--------------------
S-Suspended P-Pending E-Error T-Timeout R-Ready
-----------------------------------------------------------------------
Switch      SVL         Ports                          Link-Status
            Protocol-Status
------      ---         -----                          -----------
            ---------------
1           1           FortyGigabitEthernet1/1/1  U       R
                        FortyGigabitEthernet1/1/2  U       R
2           1           FortyGigabitEthernet2/1/1  U       R
                        FortyGigabitEthernet2/1/2  U       R
```

Stackwise Virtual Roles

- Primary election is **based on role priority**
- Rules of election:
  - Lower priority wins
  - Lower system mac wins
- Role is **non-preemptive**
- Operational role may be different from the priorities configured under the domain

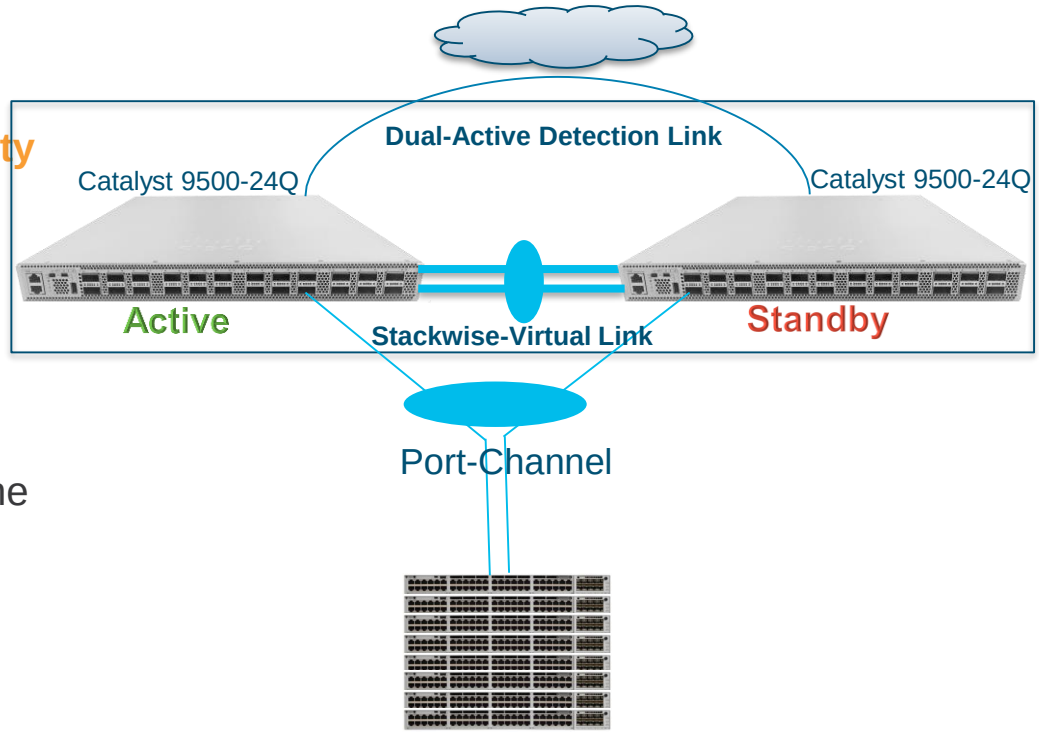

High Availability
Dual-Active Protocols

Fast Hello

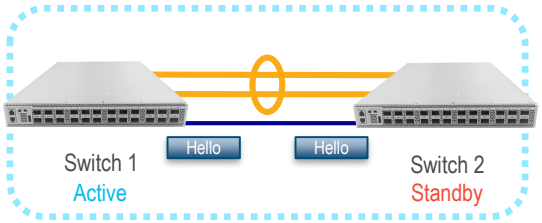

❖ **Direct L2 Point-to-Point Connection**

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

Enhanced PAGP

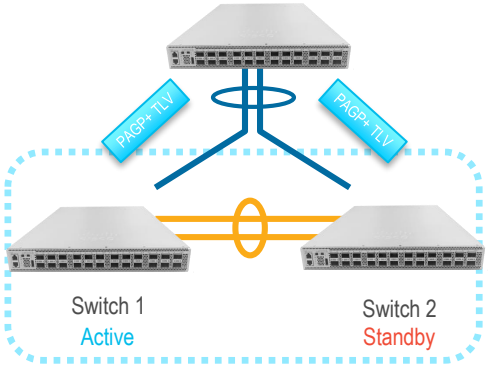

❖ **Requires ePAGP capable neighbor:**

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

StackWise Virtual – VPN Technologies

GRE Network Design

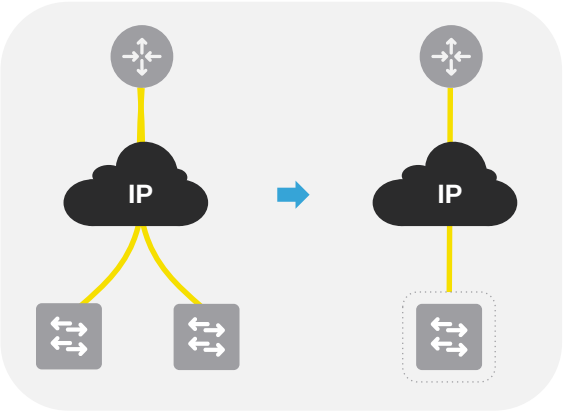

Multi-VRF Network Design

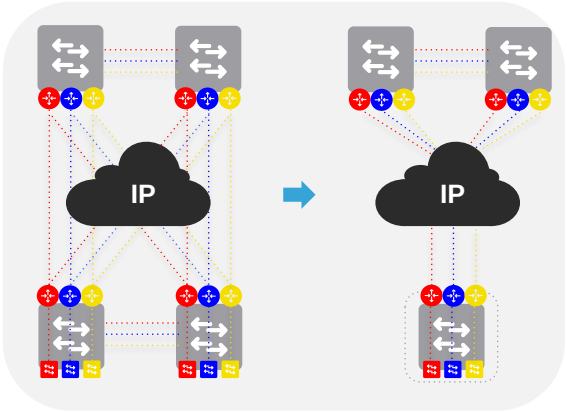

MPLS VPN Network Design

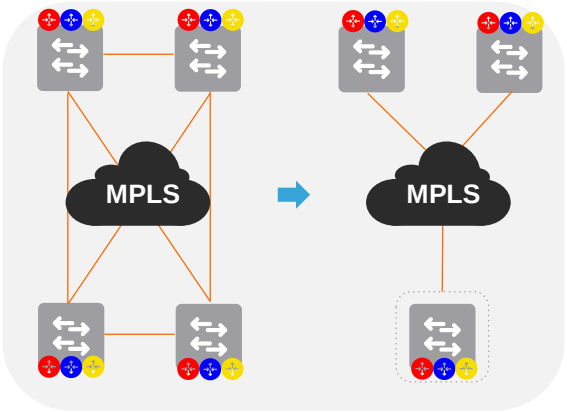

- Network Virtualization key benefits with StackWise Virtual :
  - Simplified : Underlay and Overlay networks topologies gets simplified
  - Scalable : Simplified topologies creates opportunity for larger network scale
  - Resilient : Hardware-based fault detection and recovery, instead per protocol instance
- StackWise Virtual supports GRE, Multi-VRF and MPLS VPN from initial IOS software release.

Deploying StackWise Virtual

Preparing Deployment

StackWise Virtual Deployment Pre-Requisite

	SW-1	SW-2
Model	Catalyst 9k	Catalyst 9k
Software Version	16.6.1	16.6.1
Software License	Essentials | Advantage	Essentials | Advantage
Inter-Chassis Link	40G | 10G	40G | 10G
Max Inter-Chassis Link	8	8
Dual-Active Detection Link	4	4

- Cisco StackWise Virtual Catalyst 9k series systems must have consistent IOS Version
- Both StackWise Virtual members must have same IOS Software License type to pair in same domain.

StackWise Virtual – Configuration at a glance

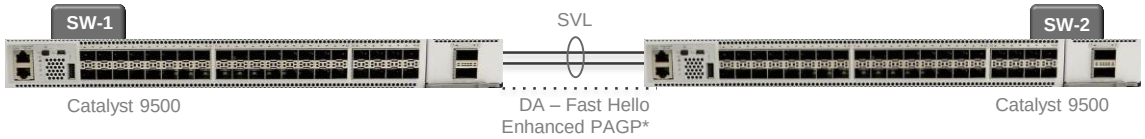

Step 1

SW-1	SW-2
9500-Dist-1(config)# stackwise-virtual	9500-Dist-2(config)# stackwise-virtual
9500-Dist-1(config)# domain <1-255>	9500-Dist-2(config)# domain <1-255>

Step-1 : StackWise Virtual Domain

- Enable StackWise Virtual
- Configure StackWise Domain–ID. Range 1-255. Default=1. Optional.

SW-1	SW-2
9500-Dist-1(config)# interface range FortyG x/y/z	9500-Dist-2(config)# interface range FortyG x/y/z
9500-Dist-1(config-if)# stackwise-virtual link <1 | 255>	9500-Dist-2(config-if)# stackwise-virtual link <1 | 255>

Step-2: StackWise Virtual Link

- Configure StackWise Virtual Links
- 10G or 40G SVL interface. Auto EtherChannel
- Max 8 SVL member-links

SW-1	SW-2
9500-Dist-1(config)# interface range TenG x/y/z	9500-Dist-2(config)# interface range TenG x/y/z
9500-Dist-1(config-if)# stackwise-virtual dual-active-detection	9500-Dist-2(config-if)# stackwise-virtual dual-active-detection

Step-3: Dual-Active Detection

- Configure StackWise Virtual Links
- 10G or 40G Dual-Active Support
- Max 8 Dual Active Fast Hello links

SW-1	SW-2
9500-Dist-1# copy run start	9500-Dist-2# copy run start
9500-Dist-1# reload	9500-Dist-2# reload

Step-4: Save and Reload to Convert

- Copy configuration on startup and storage
- Reload both system to enable StackWise-Virtu

StackWise Virtual – Software Upgrade

Auto Software Upgrade

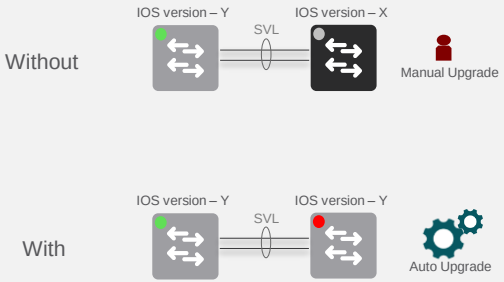

- StackWise Virtual members must have common IOS software version to pair in SSO redundancy state
- Stack member with version mis-match with ACTIVE switch will fail to RPR mode.
- Enable "**software auto-upgrade enable**" command to automate upgrade process.
- System must boot in Install mode (Default and Recommended). Auto Upgrade not supported in Bundle mode.

Cisco Prime Infra SWIM Upgrade

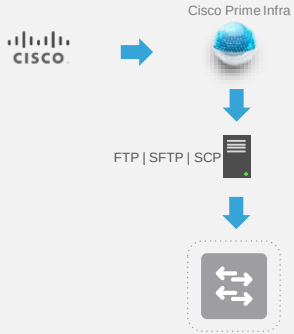

- Cisco IOS software upgrade from centralized Cisco Prime Infrastructure Software Image Management (SWIM)
- Supports internal or external file distribution server with – FTP, SFTP and SCP protocols
- Upgrade single or multiple StackWise Virtual domains based on automated schedule or on-demand.

In-Service Software Upgrade (ISSU)

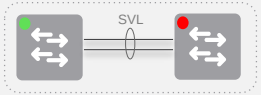

- Cat 9500 series systems deployed in StackWise Virtual mode will support ISSU*
- Plan for network downtime during software upgrade on both StackWise Virtual systems

Stackwise Virtual ISSU

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 IOS upgrade (or downgrade)
- Key technology is the ISSU Infrastructure
  - Allows SSO between different versions

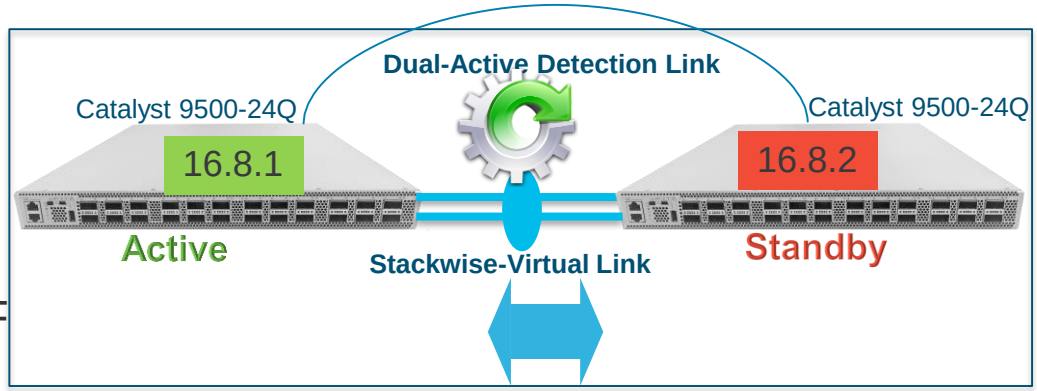

Stackwise Virtual ISSU

ISSU Workflow

The process starts by copying the images to both stackwise virtual nodes

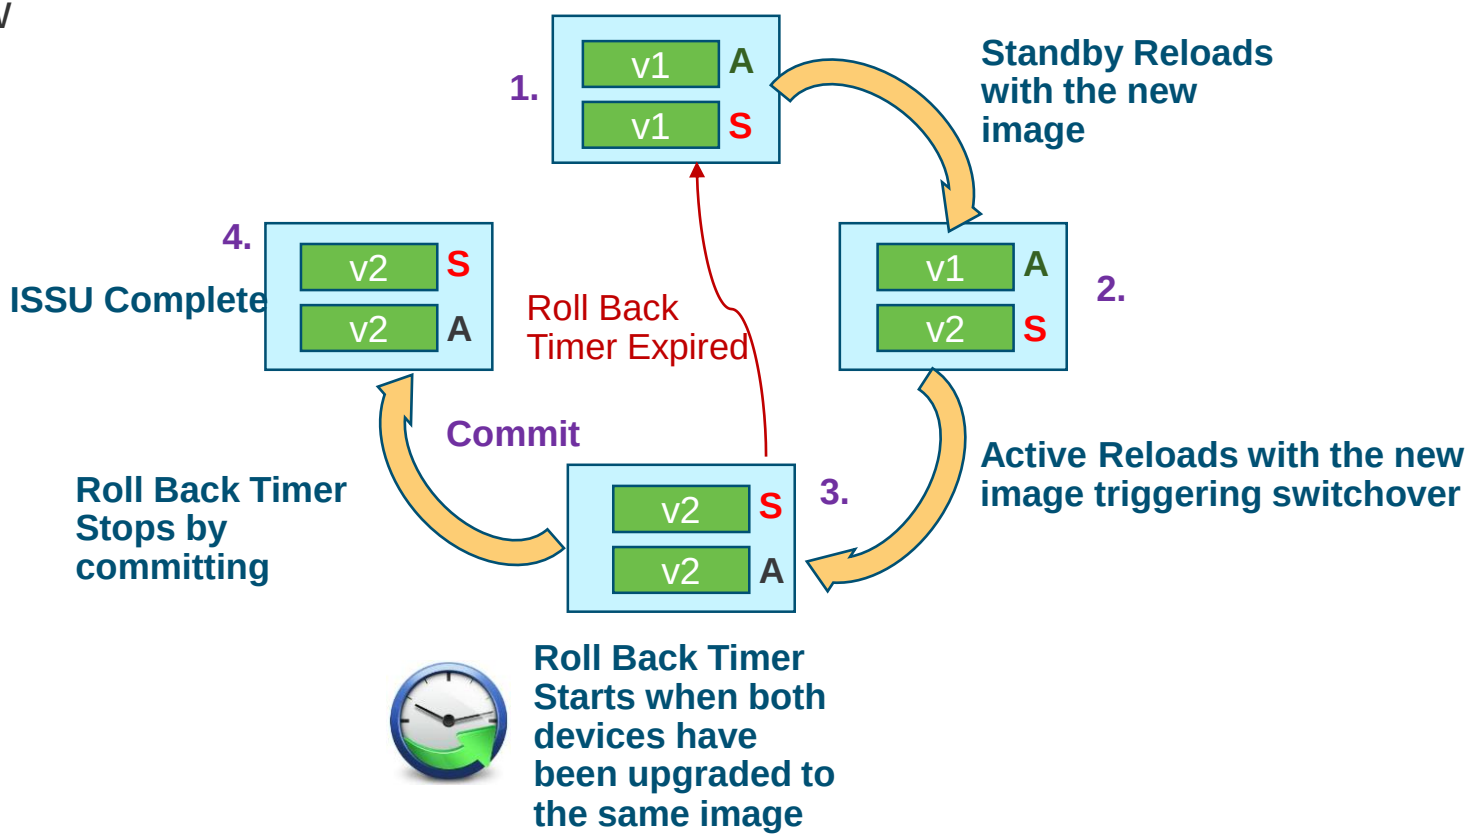

Stackwise Virtual ISSU

ISSU Approach

- New image is provisioned on the standby switch and is reloaded to come up with new image
- Active switch is provisioned with the new image and is reloaded triggering switchover
- ISSU completes with both switches running the same image

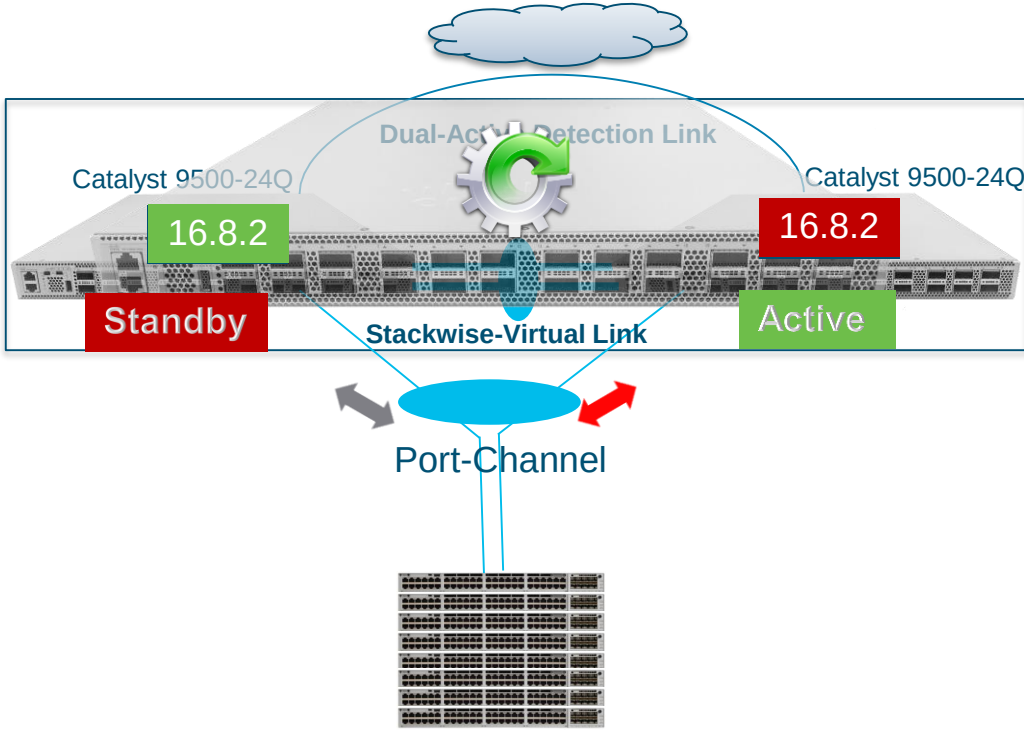

Stackwise Virtual ISSU

ISSU Steps

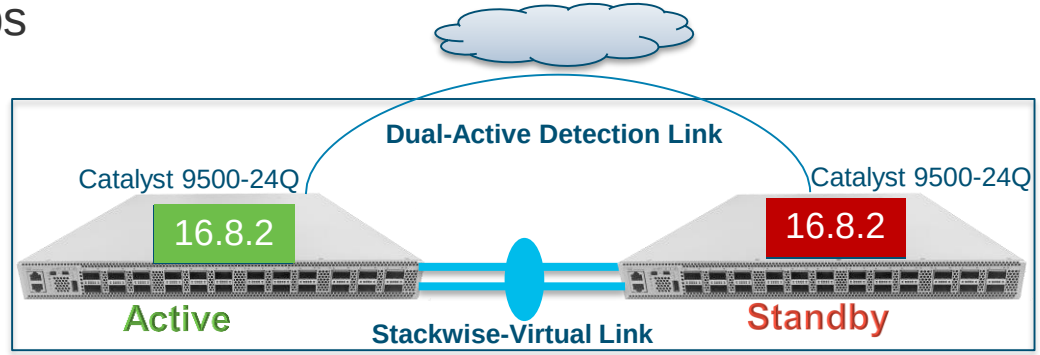

3 Step Process

- Install add file <tftp/ftp/flash/disk:*.bin>
- Install activate ISSU
- Install commit

1 Step Process

- Install add file <tftp/ftp/flash/disk:*.bin>activate ISSU commit

Stackwise Virtual ISSU

ISSU Caveats

- ISSU is only supported in Install Mode
- ISSU is only supported if both switches are in SSO Active/Standby
- ISSU will not work if there is insufficient memory in the flash to hold the existing and new image

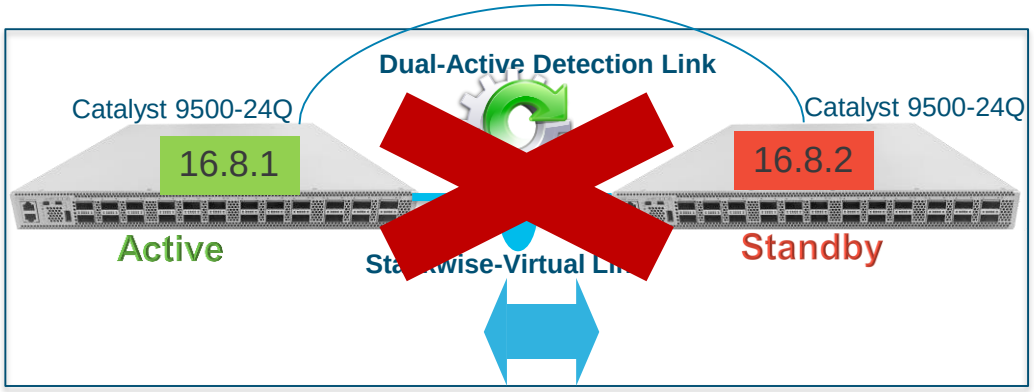

Graceful Insertion and Removal

Graceful Insertion and Removal

Upgrades with no or Minimal Traffic Loss

Comprehensive Node Isolation Framework

Easy Execution with a single command

Highly Customizable workflow

Simple
Customizable
Non-Traffic
Impacting

Graceful Insertion and Removal on Catalyst 9000

Isolation of Switch from network Gracefully

- Isolate a switch from the network in order to perform debugging or an upgrade.
- Shutdown Vs. Isolate Mode
  - **Shutdown**: All protocols are gracefully brought down and all physical ports are shut down. (7.2.1)
  - **Isolate**: All protocols are gracefully brought down but is not shutdown. (7.3.0)

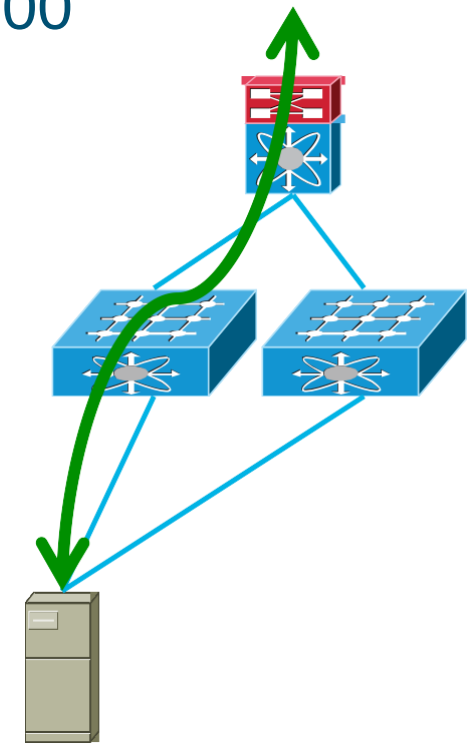

L2 and L3 Topology with GIR Isolation

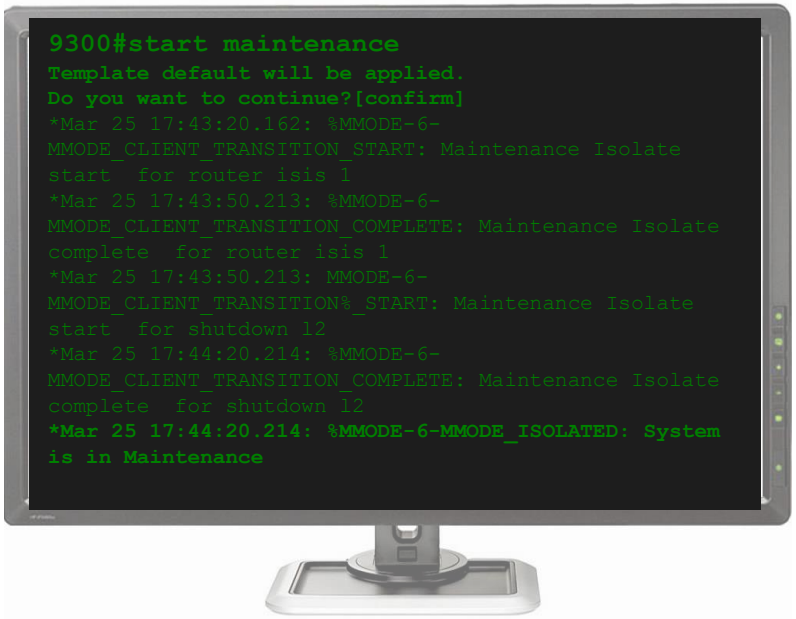

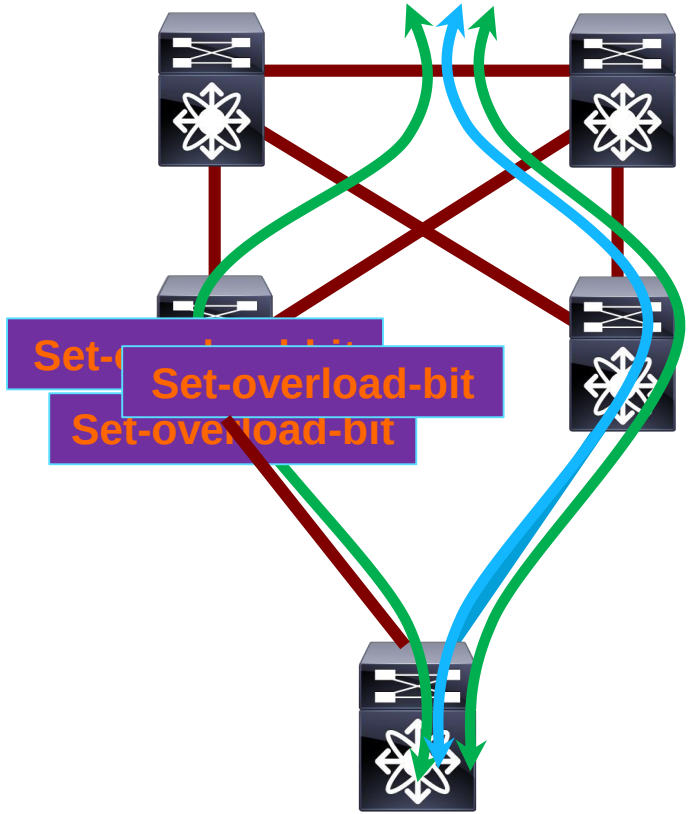

Order for Maintenance:

EGP -> IGPs 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
```

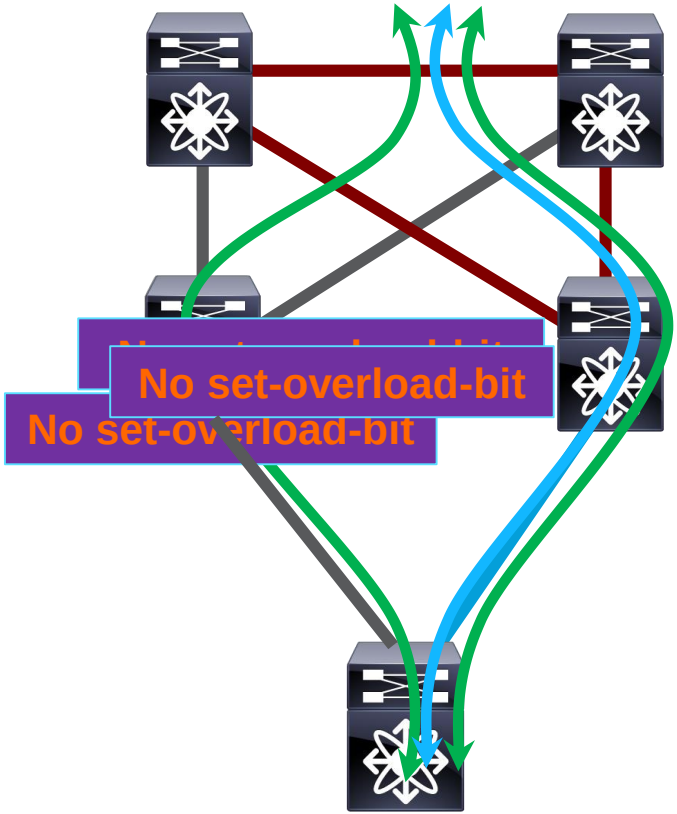

Order for Maintenance:

L2 → IGPs in parallel (ISIS) -> EGP

Graceful Insertion and Removal

Default and Customizable Templates

- **Default Template**
  - **System Generated Profile based on the switch configuration**
- **Customized Template**
  - **User Configured Profile based on specific configuration or use case**

```
9300L#show system mode maintenance template default
System Mode: Normal
default maintenance-template details:
router isis 1
shutdown l2
9300L#show system mode maintenance template test
System Mode: Normal
Maintenance Template test details:
shutdown l2
```

Graceful Insertion and Removal

Snapshots

- **Automatic Snapshots**
  - **Snapshots are automatically generated when entering and exiting maintenance mode**
  - **Captures operational data from the running system like Vlan's, Routes etc.**
- **User Configured Snapshots**
  - **Snapshots can be collected manually for comparing and troubleshooting**

```
Switch#show system snapshots compare before_maintenance
after_maintenance
================================================================================
Feature              Tag              .before_maintenance .after_maintenance
================================================================================
[interface]

--------------------------------------------------------------------------------
      [Name:Vlan1]
                     packetsinput                  181587          **181589**
      [Name:GigabitEthernet1/0/3]
                     packetsinput                  101531          **101550**
                     broadcasts                    80893           **80910**
                     packetsoutput                 211568          **211594**
      [Name:GigabitEthernet1/0/8]
                     output                        00:00:00,       **00:00:04,**
                     packetsinput                  6915            **6918**
                     packetsoutput                 57677           **57706**
      [Name:GigabitEthernet1/0/17]
                     packetsinput                  101528          **101550**
                     broadcasts                    80891           **80910**
                     packetsoutput                 211570          **211600**
```

Graceful Insertion and Removal

Maintenance Profile Options

- **On-Reload**
  - **If the switch is reloaded in maintenance mode, the switch will come back in maintenance mode**
- **Failsafe**
  - **Timeout for Client Acknowledgement**
- **Duration**
  - **The Switch will come out of maintenance after the configured duration**

```
9300(config)#system mode maintenance
9300(config-maintenance)#?
maintenance mode submode configuration commands:
default     Set a command to its defaults
exit        Exit from maintenance configuration mode
failsafe    Client ack timeout
no          Negate a command or set its defaults
on-reload   On reload maintenance mode configuration
template    use maintenance-template
timeout     maintenance duration
```

Configuration Profiles

- Maintenance-mode profile is applied when entering GIR mode,
- Normal-mode profile is applied when GIR mode is exited.

Automatic Profiles	Custom Profiles
• Generated by default • GIR is applied to all protocols running on the system • GIR state machine uses Registry mechanism to interface with client protocols • Use: Maintenance Windows	• User created profile for maintenance-mode and normal-mode using “templates” • Flexible selection of protocols for isolation • Use: maintenance windows and isolation during troubleshooting using preconfigured templates

Best Practices for Distribution

Stackwise Control Plane Recommendations

Stackwise Virtual Link & Dual Active Detection

- For the SVL Link, it is good to use two 10/40 GbE ports on separate ASIC's for resilient connectivity.
- SVL links should be **point-to-point** connection.
- It is recommended that you use a dedicated link for Dual Active Detection.

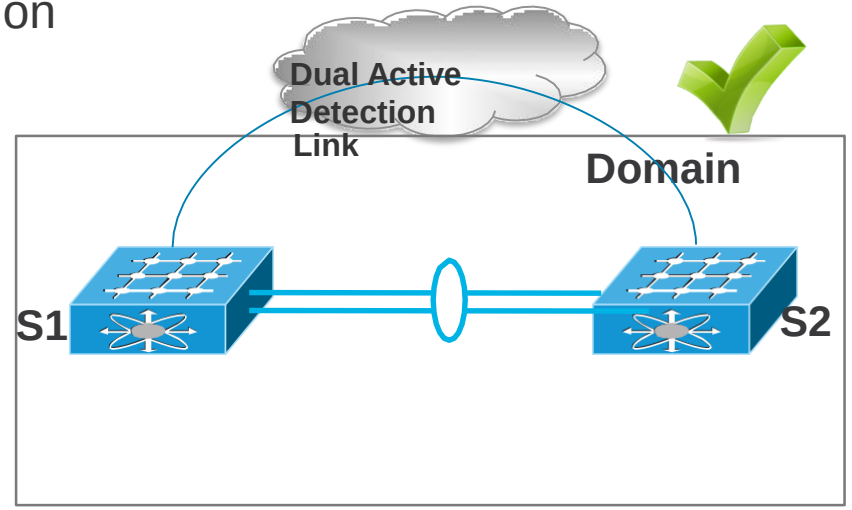

Member connectivity Recommendations

Dual Home Devices

- Make sure to dual attach all devices into the SVL Domain.
- It is good to use LACP for the Port Channels of MEC Member ports.

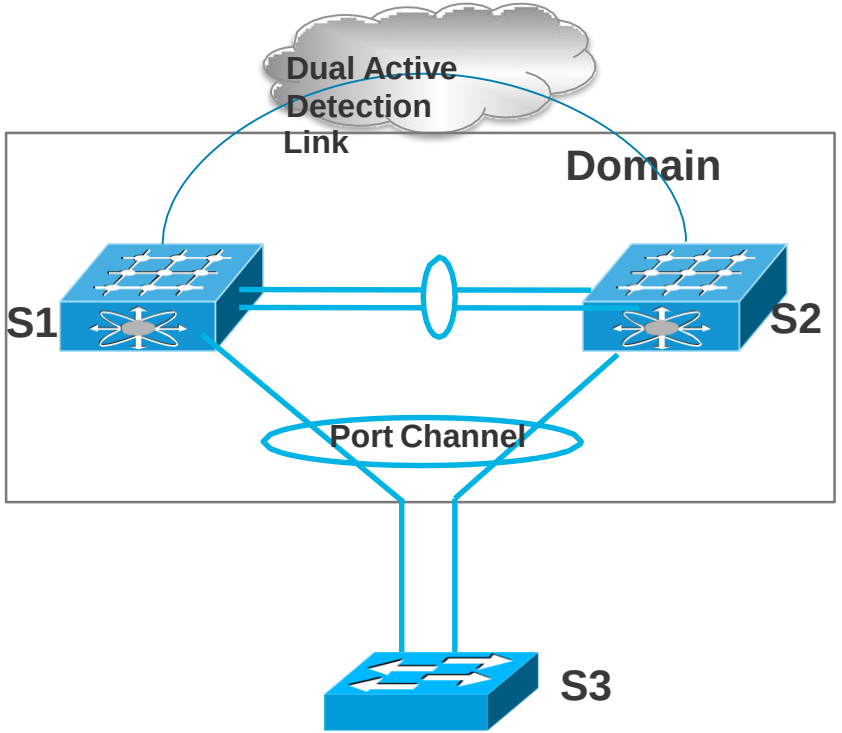

Graceful Insertion and Removal Recommendations

Dual Home devices

- Similar to SVL, it is recommended that you have an alternate path for your devices to take
- Traffic will be lost during the maintenance mode operation from that unit for units that are singly attached

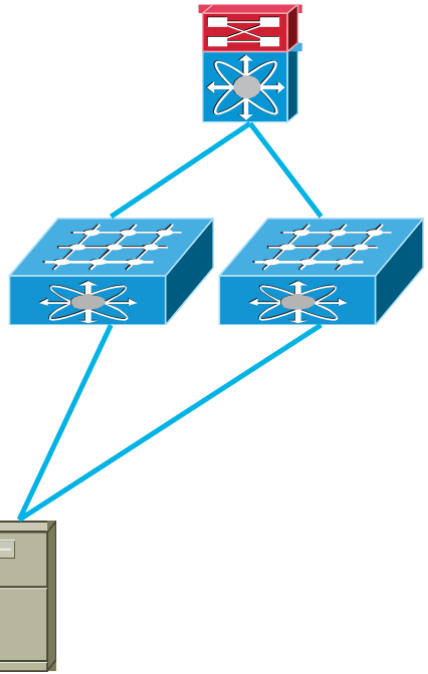

Key Recommendations

- SVL usage for better Availability
- SVL Control Plane Recommendation
  - Use Different ASIC's for resiliency
  - SVL links should be point-to-point
  - DAD link should be used as a dedicated link
- MEC member connectivity recommendations
  - Always dual attach devices to SVL domain
  - Best to use LACP as the protocol for the SVL members
- Graceful Insertion and Removal recommendations
  - Set maintenance timer to a reasonable value (60 recommended)
  - Always dual attach devices to have an alternate traffic path

Agenda

- High Availability Overview and Evolution
- High Availability Solution on the Campus Access
  - Stackable High Availability Solution
  - Modular High Availability Solution
- High Availability Solution on the Campus Distribution
- High Availability Solution in IOS
- Summary/Q&A

Open IOS-XE

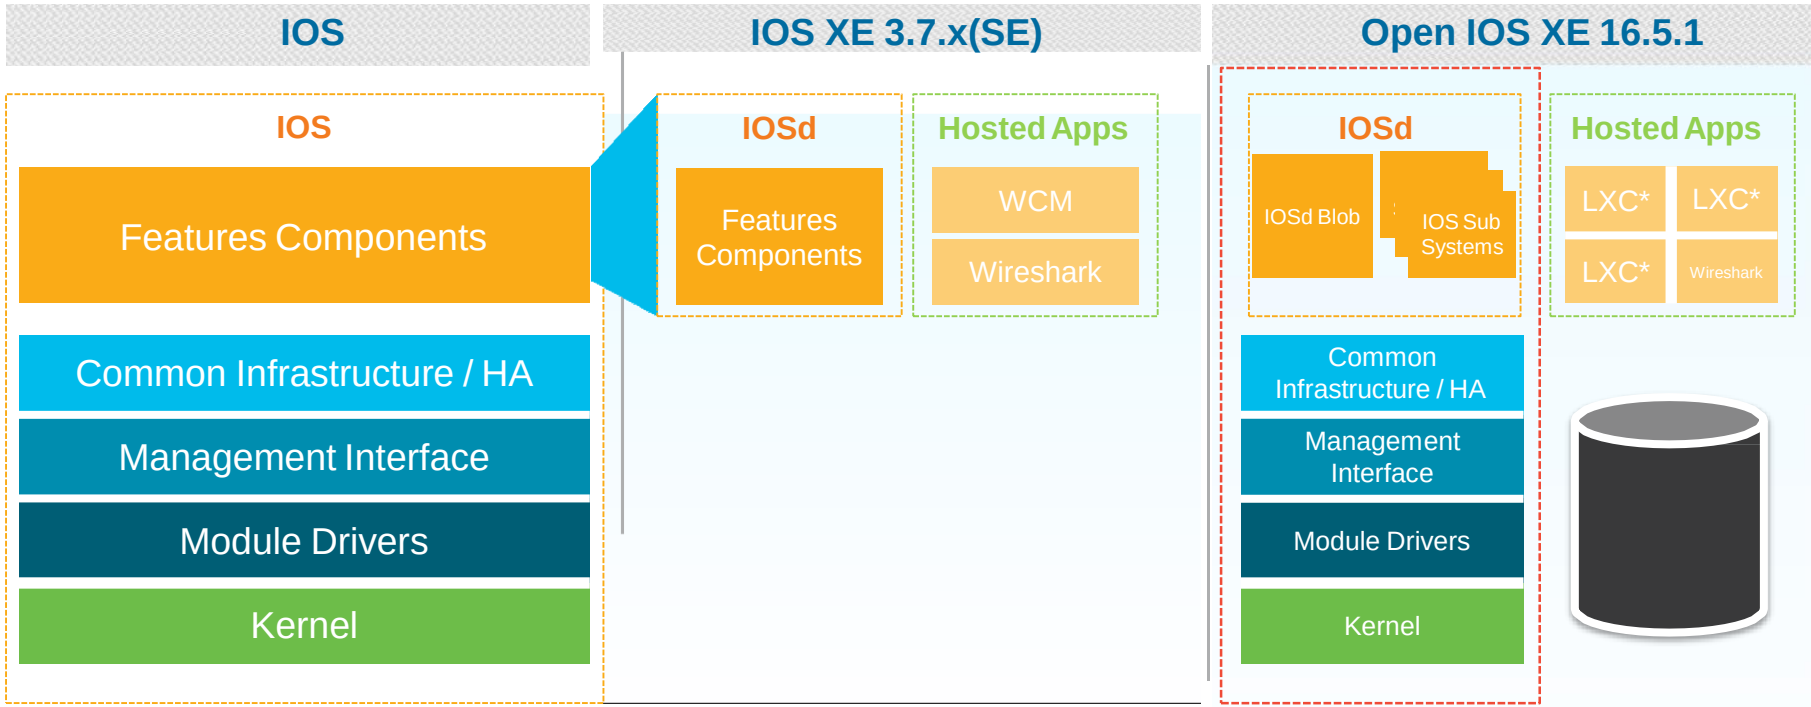

Same Binary Image Across all Catalyst 9K Family

IOS XE – IOS-XE 16.x (Illustrative View)

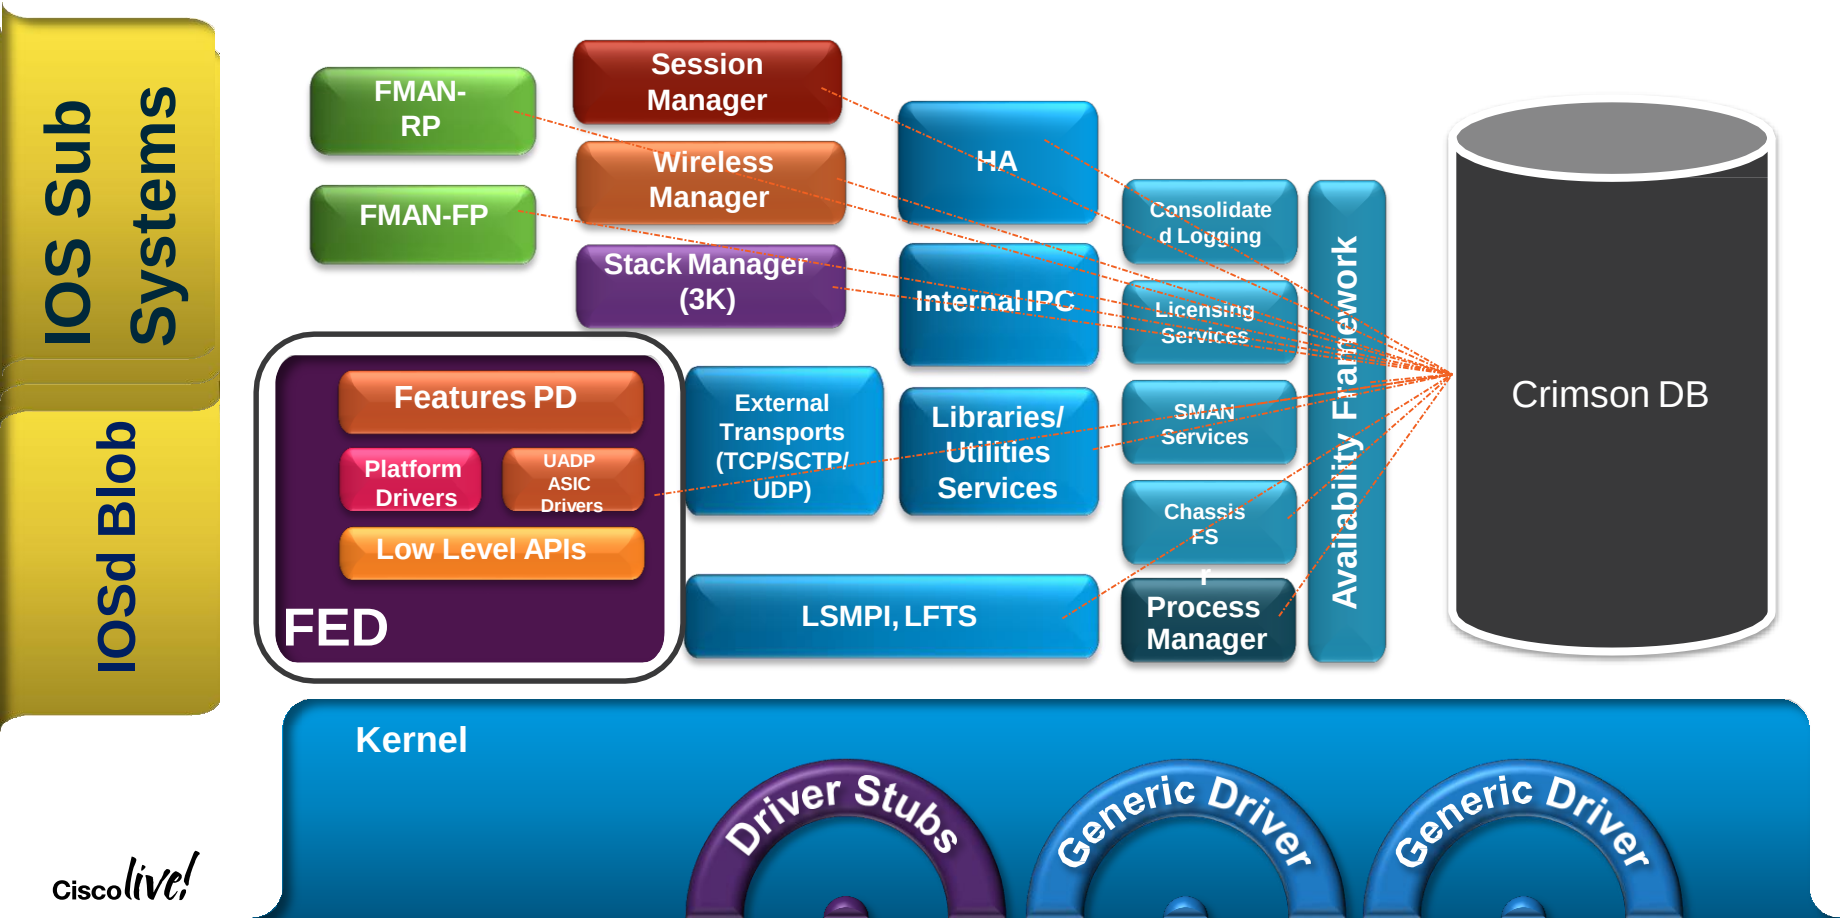

Stateful Switchover Mode – IOS

SSO-Aware and SSO-Compliant IOS Applications

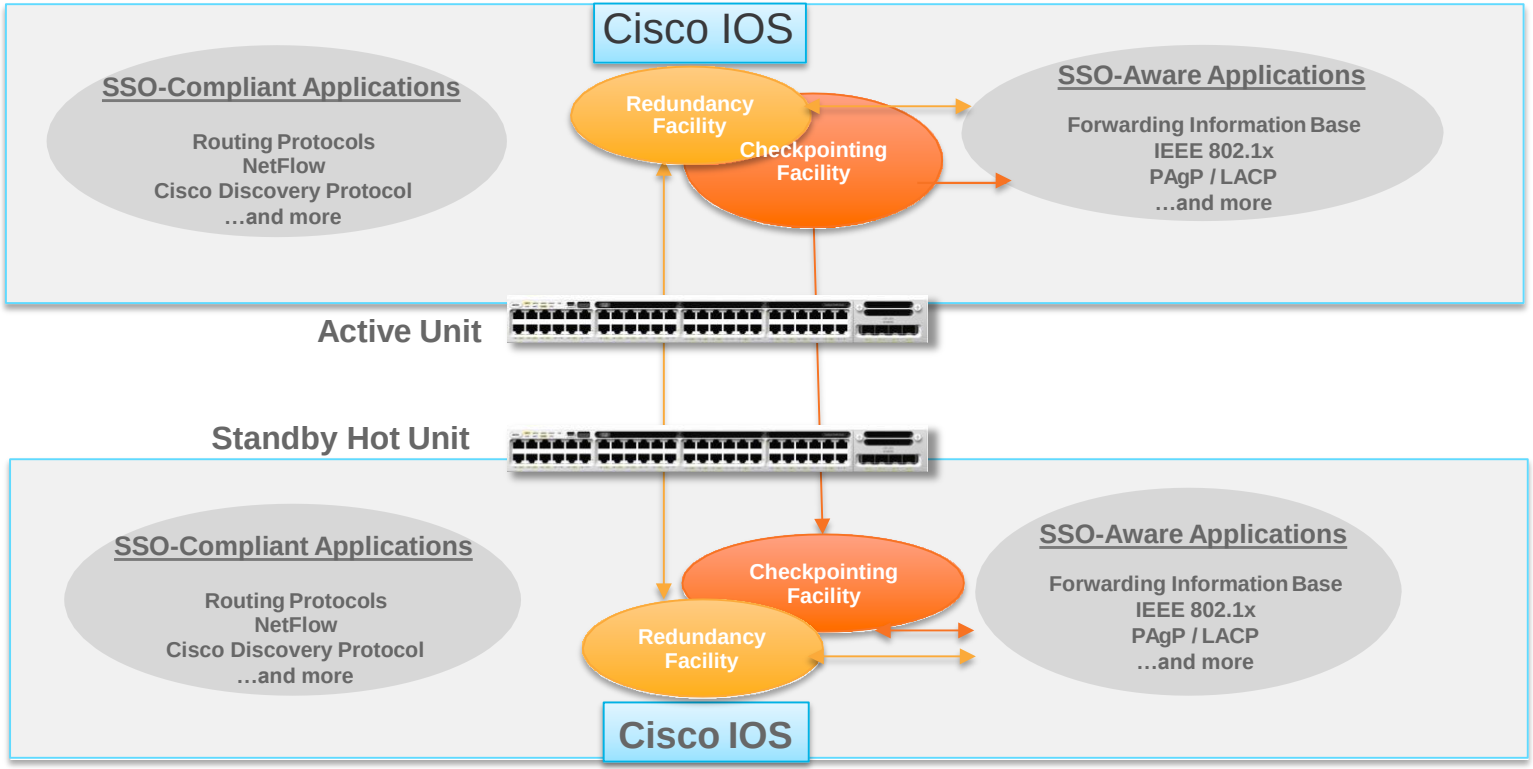

SSO Compliant Redundancy Clients – examples

```
Switch#sh redundancy clients

Group ID =  1
    clientID = 20002        clientSeq =   4   EICORE HA Client
    clientID = 24001        clientSeq =  28   Table Manager Client
    clientID = 20010        clientSeq =  31   SNMP SA HA Client
    clientID = 20007        clientSeq =  34   Installer HA Client
    clientID =  29          clientSeq =  60   Redundancy Mode RF
    clientID =  139         clientSeq =  61   IfIndex
    clientID =  3300        clientSeq =  62   Persistent Variable
    clientID =  25          clientSeq =  68 CHKPT RF
    clientID = 20005        clientSeq =  74  IIF-shim
```

Security	Other Applications
AAA	Bonjour
DHCP	Net flow
CTyptoRSA	Wireless
ity	
ARP	

Some processes have more than one HA client

```
Switch#sh redundancy slaves

Group ID =    1
Slave/Process ID =   6175    Slave Name = [installer]
Slave/Process ID =   6177    Slave Name = [eicored]
Slave/Process ID =   6198    Slave Name = [snmp_subagent]
Slave/Process ID =  12981    Slave Name = [wcm]
Slave/Process ID =  12982    Slave Name = [table_mgr]
Slave/Process ID =  12985    Slave Name = [iosd]
```

SSO by itself Does Not Provide Redundancy for the Routing Protocols

Graceful Restart, Non-Stop Forwarding and Non- Stop Routing

- Non-Stop Forwarding was developed by Cisco to maintain traffic forwarding by a router experiencing a control plane switchover event. The router will essentially synchronize its Forwarding Information Base between an Active and Standby Route Processor as well as signal to its routing neighbors to continue forwarding traffic while routing topology information is exchanged
- The IETF developed standards based implementations similar to Cisco NSF
- The IETF implementations use different terminology including the terms “Graceful Restart” to describe the signaling used between the routers
- Graceful Restart(GR) and Non-Stop Forwarding (NSF) are terms often used interchangeably
- Graceful Restart/Non-Stop Forwarding as well as Non-Stop Routing (NSR) all allow for the forwarding of data packets to continue along known routes while the routing protocol information is being restored (in the case of Graceful Restart) or refreshed (in the case of Non Stop Routing) following a processor switchover.
- Each routing protocol has its own unique implementation and signaling mechanisms

Routing Protocol Redundancy With NSF

Active Supervisor Engine Slot 1

EIGRP RIB	
Prefix	Next Hop
10.0.0.0	10.1.1.1
10.1.0.0	10.1.1.1
10.20.0.0	10.1.1.1

OSPF RIB	
Prefix	Next Hop
192.168.0	192.168.0.1
192.168.55..0	192.168.55.1
192.168.32.0	192.168.32.1

ARP Table	
IP	MAC
10.1.1.1	aabbcc:ddee32
10.1.1.2	adbb32:d34e43
10.20.1.1	aa25cc:ddeee8

Standby Supervisor Engine Slot 2

EIGRP RIB	
Prefix	Next Hop
-	-
-	-
-	-

OSPF RIB	
Prefix	Next Hop
-	-
-	-
-	-

ARP Table	
IP	MAC
-	-
-	-
-	-

FIB Table	
Prefix	Next HOP
10.1.1.1	aabbcc:ddee32
10.1.1.2	adbb32:d34e43
192.168.0.0	aa25cc:ddeee8

SSO

Redundancy Facility

Checkpoint Facility

FIB Table	
Prefix	Next HOP
10.1.1.1	aabbcc:ddee32
10.1.1.2	adbb32:d34e43
192.168.0.0	aa25cc:ddeee8

Routing Protocol Redundancy With NSF

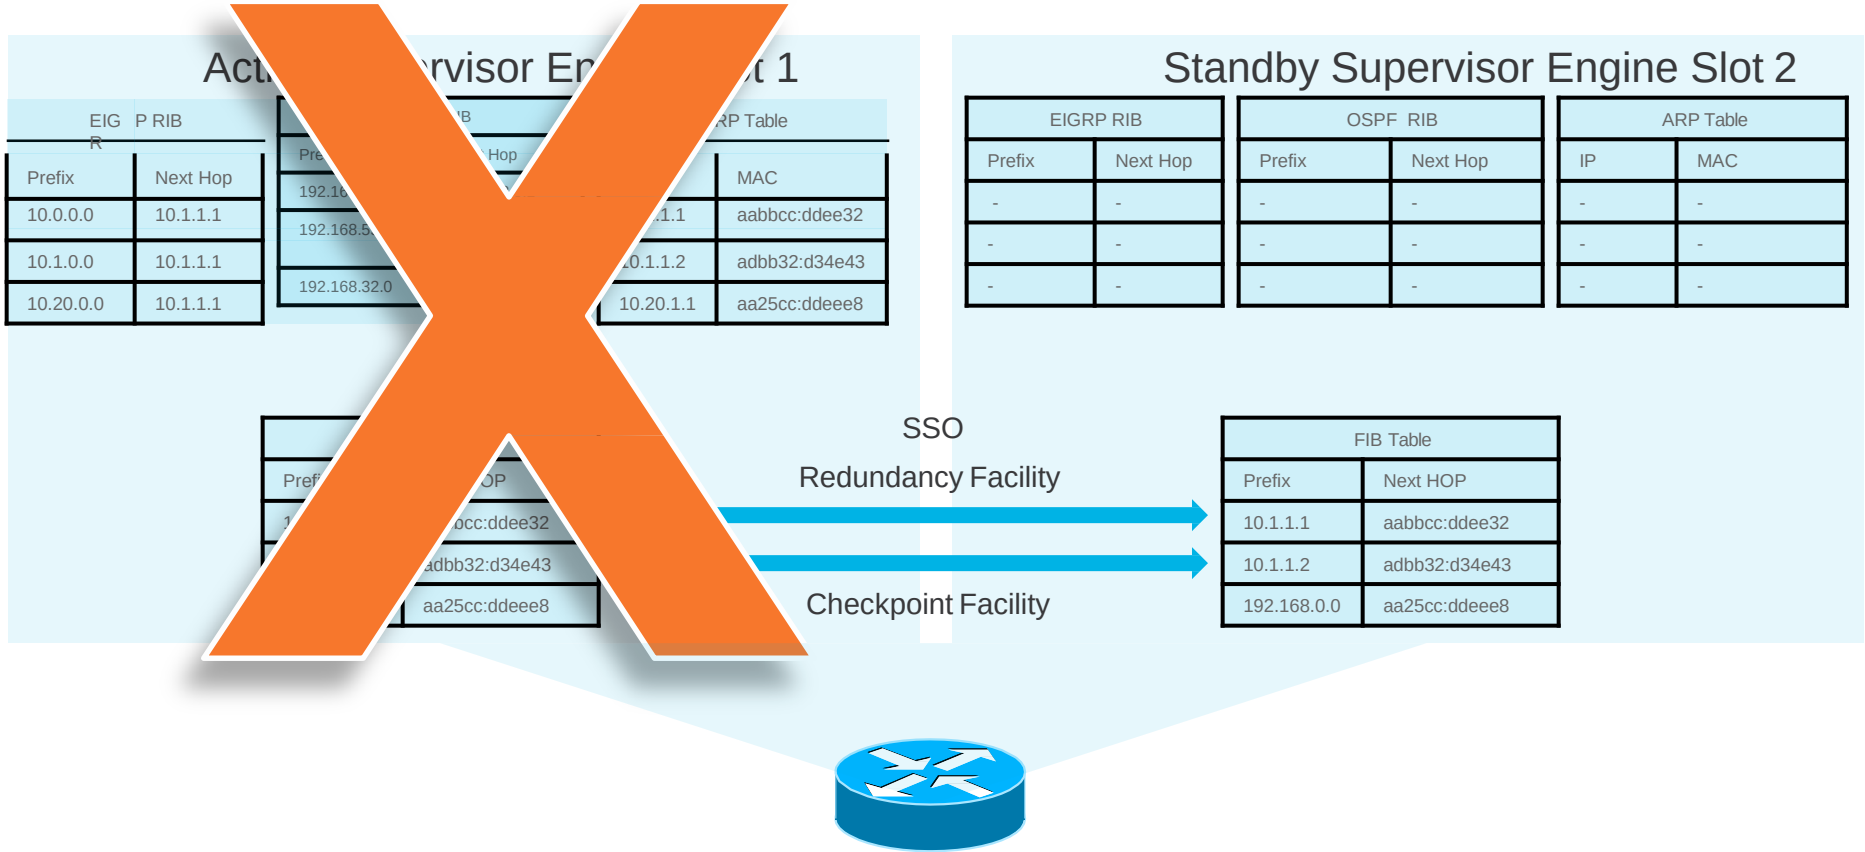

Routing Protocol Redundancy With NSF

Standby Supervisor Engine Slot 2

EIGRP RIB	
Prefix	Next Hop
1-0.0.0.0	-10.1.1.1
-10.1.0.0	-10.1.1.1
-10.20.0.0	10.1.1.1

OSPF RIB	
Prefix	Next Hop
192.168.0 -	192.168.0.1 -
192.168.55..0 -	192.168.55.1 -
-192.168.32.0	-192.168.32.1

ARP Table	
IP	MAC
-10.1.1.1	aabbcc:ddee32
-10.1.1.2	-adbb32:d34e43
-10.20.1.1	-aa25cc:ddeee8

FIB Table	
Prefix	Next HOP
10.1.1.1	aabbcc:ddee32
10.1.1.2	adbb32:d34e43
192.168.0.0	aa25cc:ddeee8

GR/NSF Signaling per protocol

Synchronization per protocol

Non Stop Forwarding Router Roles

- Non-Stop Forwarding, NSF, allows a router to continue forwarding data along routes that are already known, while the routing protocol information is being restored

➡NSF Aware router or
NSF Helper router*

- A router running NSF-compatible software, capable of assisting a neighbor router perform an NSF restart

➡NSF Capable router

- A router configured to perform an NSF restart, therefore able to rebuild routing information from neighbor NSF-aware or NSF capable router

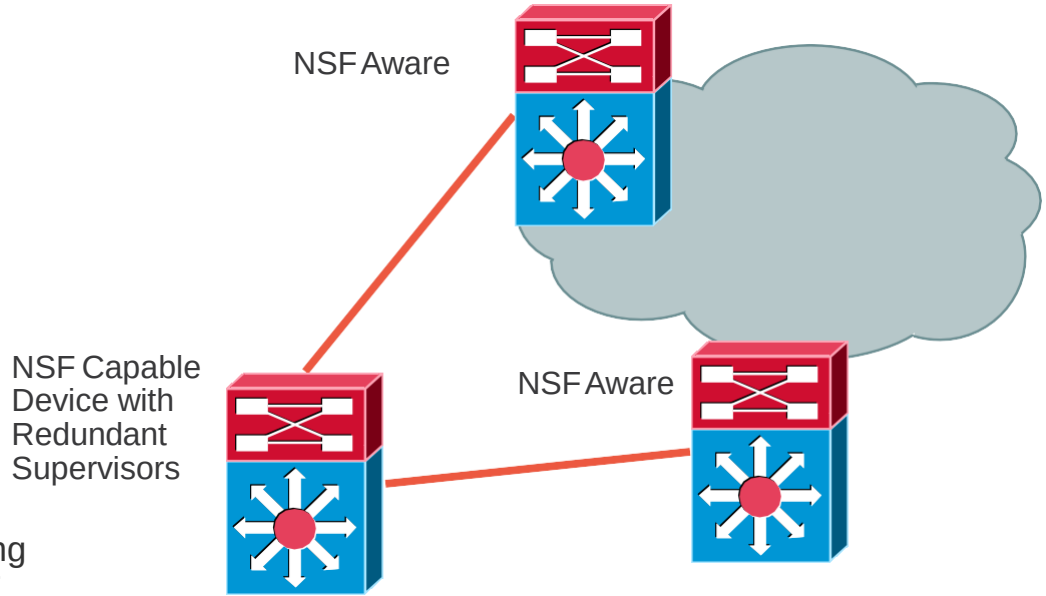

* NSF Helper - This term is used in IETF terminology

Stateful Switchover and Non-Stop Forwarding

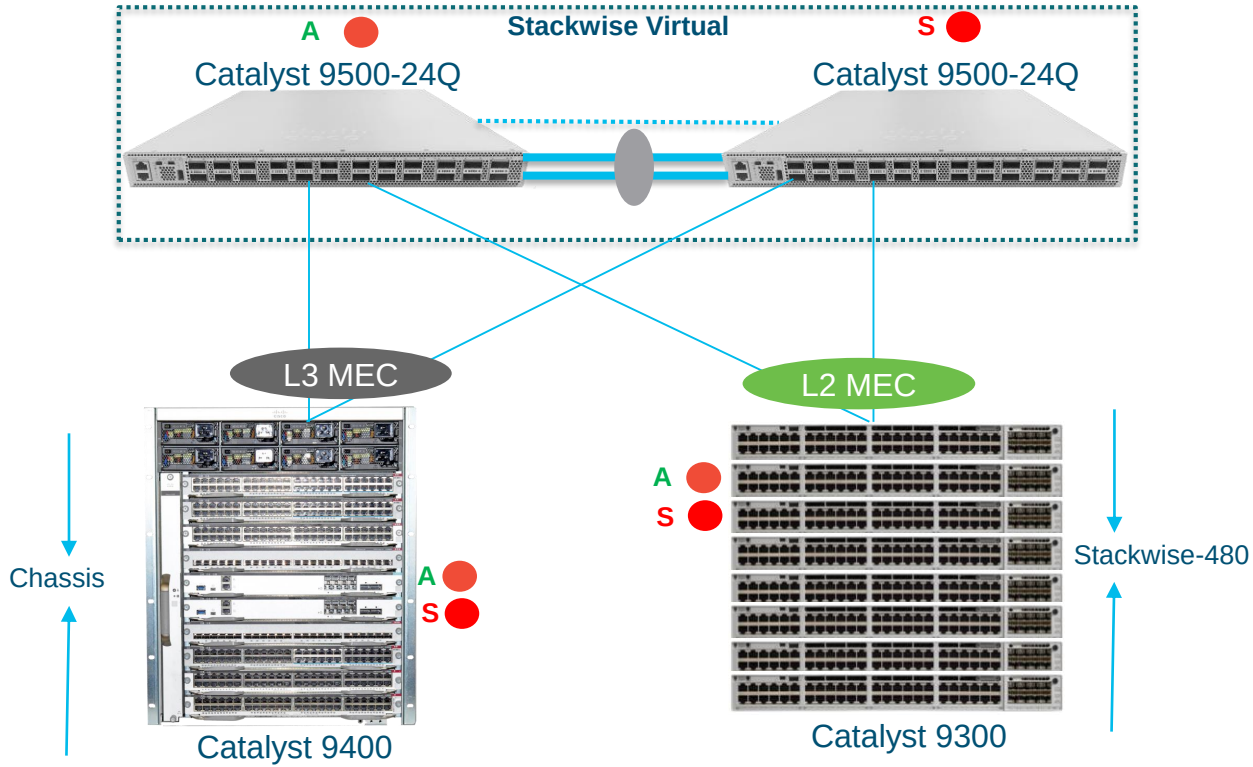

Stack Switchover Example

```
9300-1# show redundancy states

         my state = 13 -ACTIVE
       peer state =  8 -STANDBY HOT
            Mode = Duplex
         Unit ID = 1
```

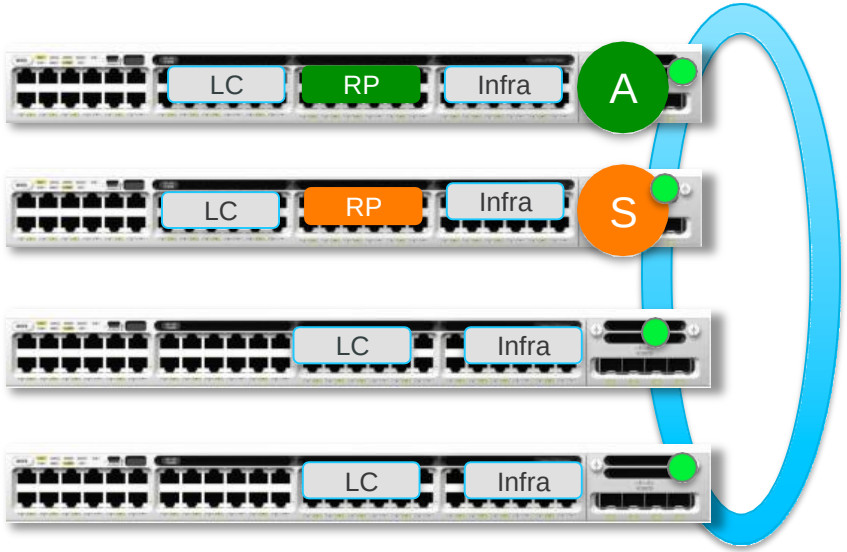

Stack Switchover Example

```
9300-1# show redundancy states
```

```
9300-1# show switch detail
     Switch/Stack Mac Address : 2037.06cf.0e80
                                                 H/W   Current
Switch#   Role     Mac Address     Priority Version  State
-----------------------------------------------------------
*1        Active   2037.06cf.0e80     10      V01     Ready
 2        Standby  2037.06cf.3380      8      V00     Ready
 3        Member   2037.06cf.1400      6      V00     Ready
 4        Member   2037.06cf.3000      4      V00     Ready

          Stack Port Status             Neighbors
Switch#   Port 1     Port 2           Port 1   Port 2
-----------------------------------------------------------
- 1         OK         OK               2        4
  2         OK         OK               3        1
  3         OK         OK               4        2
  4         OK         OK               1        3
```

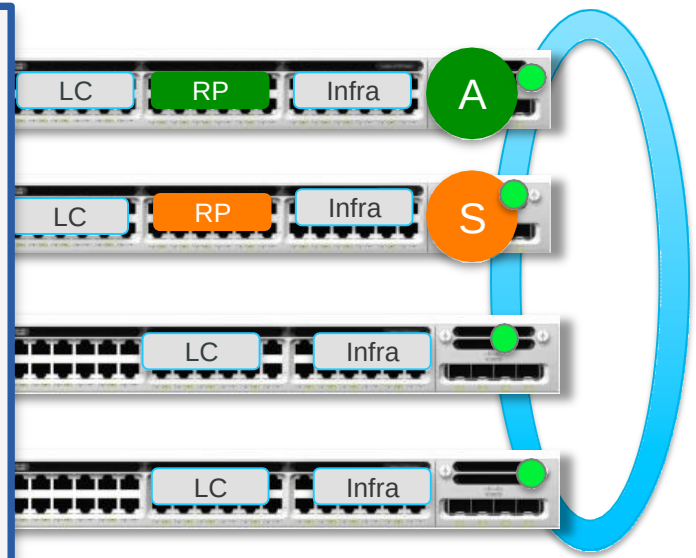

Stack Switchover Example

1. Active Unit Goes Down

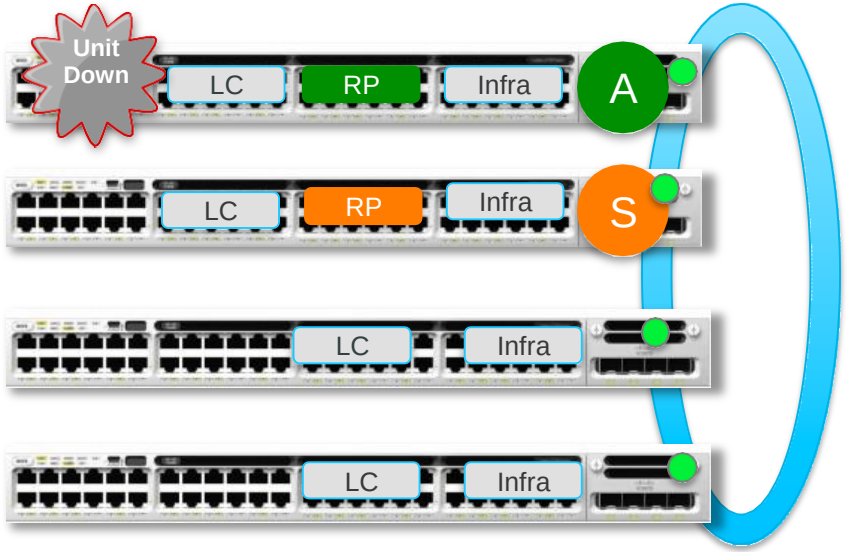

Stack Switchover Example

1. Active Unit Goes Down
2. Standby Unit Becomes the new Active

```
9300-2# show redundancy states

        my state = 13 -ACTIVE
      peer state =  1 - DISABLED
           Mode = Simplex
        Unit ID = 2
```

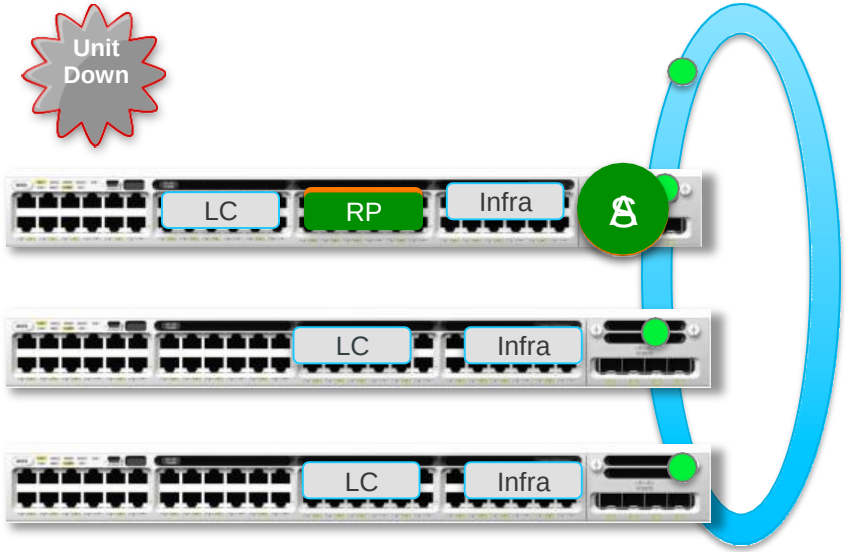

Stack Switchover Example

1. Active Unit Goes Down
2. Standby Unit Becomes the new Active

```
9300-1# show switch detail
    Switch/Stack Mac Address : 2037.06cf.0e80
9300-2# show redundancy states                                  H/W   Current
Switch#    Role    Mac Address     Priority Version  State
-------------------my state = 13 -ACTIVE----------------------------
*1       Member   0000.0000.0000     10       V01     Removed
 2       Active   2037.06cf.3380     8        V00     Ready
 3       Member   2037.06cf.1400     6        V00     Ready
 4       Member   2037.06cf.3000     4        V00     Ready

          Stack Port Status             Neighbors
Switch#   Port 1     Port 2           Port 1   Port 2
-----------------------------------------------------------
  2         OK         DOWN             3        None
  3         OK         OK               4        2
  4         DOWN       OK               None     3
```

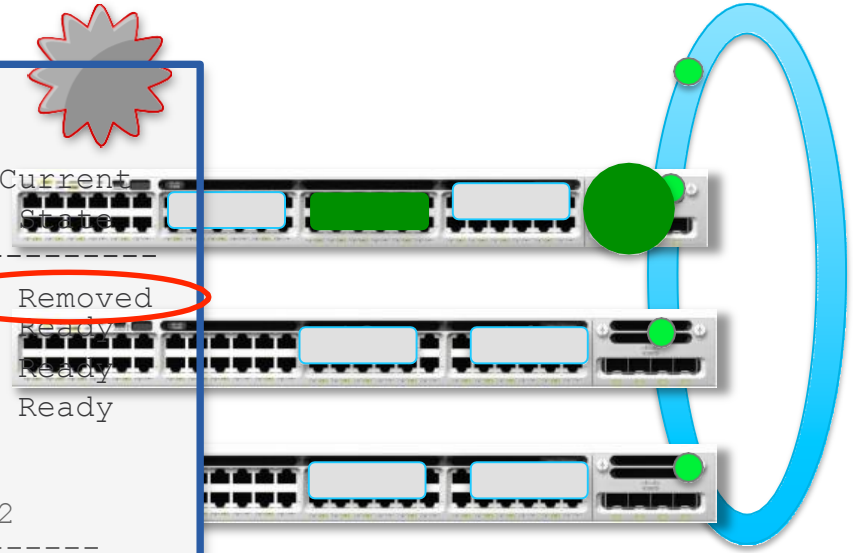

Stack Switchover Example

- Active Unit Goes Down
- Standby Unit Becomes the new Active
- Starts 2min Timer to elect Standby
- Active elects Standby
- Standby starts RP Domain locally
- Starts Bulk Sync with Active RP
- Standby reaches “Standby Hot”

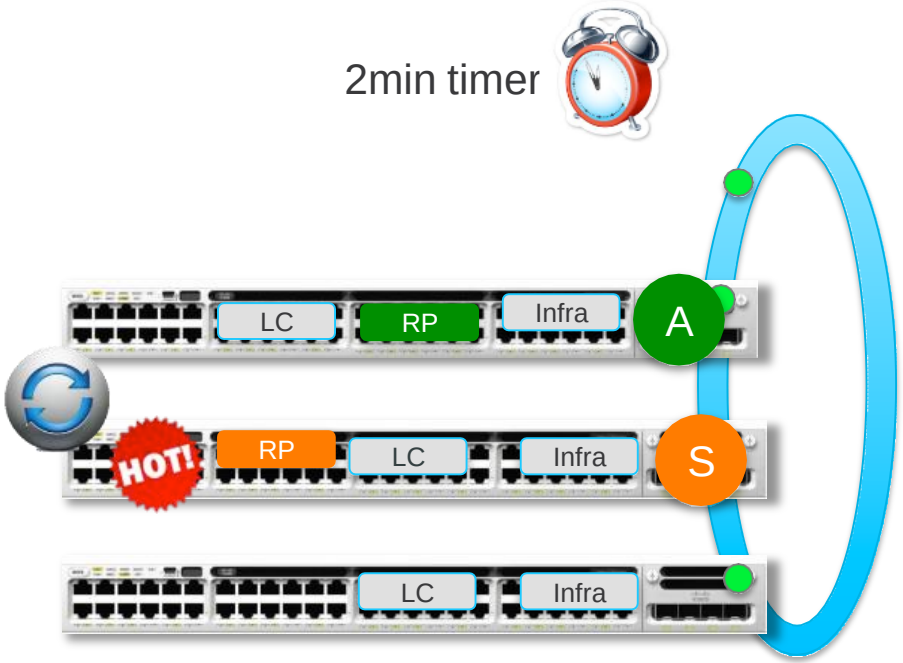

Stack Switchover Example

- Active Unit Goes Down
- Standby Unit Becomes the new Active
- Starts 2min Timer to elect Standby
- Active elects Standby
- Standby starts RP Domain locally
- Starts Bulk Sync with Active RP 7.

```
9300-2# show redundancy states

        my state = 13 -ACTIVE
      peer state =  8 - STANDBY HOT
           Mode = Duplex
        Unit ID = 2
```

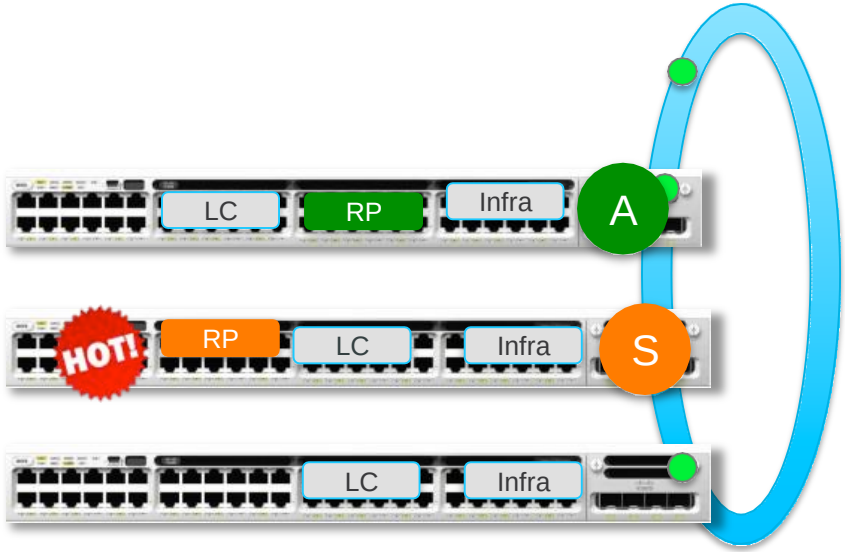

Stack Switchover Example

- Active Unit Goes Down
- Standby Unit Becomes the new Active
- Starts 2min Timer to elect Standby
- Active elects Standby
- Standby starts RP Domain locally
- Starts Bulk Sync with Active RP
- Standby reaches “Standby Hot”
- Member rejoins the Stack
- Active Programs hardware on all LCs
- Traffic resumes

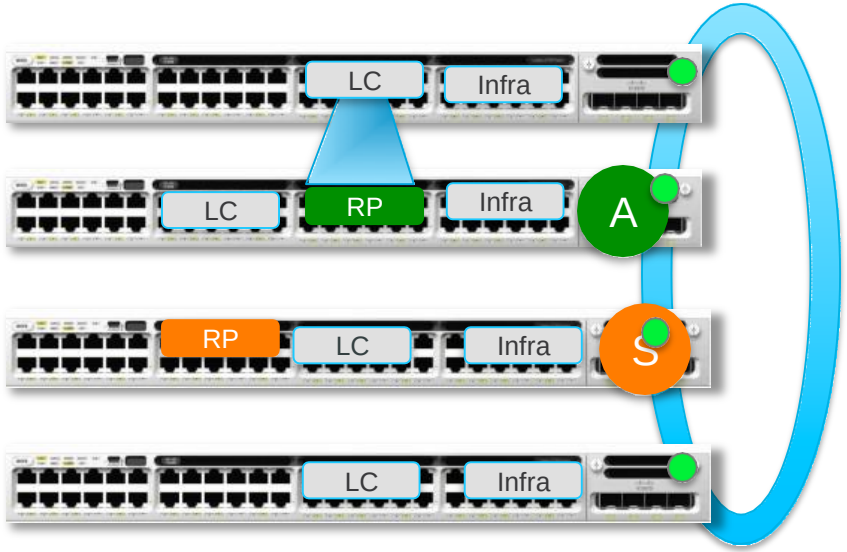

Stack Switchover Example

1. Active Unit Goes Down
2. Standby Unit Becomes the new Active
3.
4.
5.
6.
7.
8.
9.
10.

```
9300-1# show switch detail
      Switch/Stack Mac Address : 2037.06cf.0e80

                                                 H/W   Current
Switch#   Role     Mac Address     Priority Version  State
-----------------------------------------------------------
*1        Member   2037.06cf.0e80     10      V01     Ready
 2        Active   2037.06cf.3380     8       V00     Ready
 3        Standby  2037.06cf.1400     6       V00     Ready
 4        Member   2037.06cf.3000     4       V00     Ready

          Stack Port Status             Neighbors
Switch#   Port 1     Port 2           Port 1   Port 2
-----------------------------------------------------------
- 1         OK         OK               2        4
  2         OK         OK               3        1
  3         OK         OK               4        2
  4         OK         OK               1        3
```

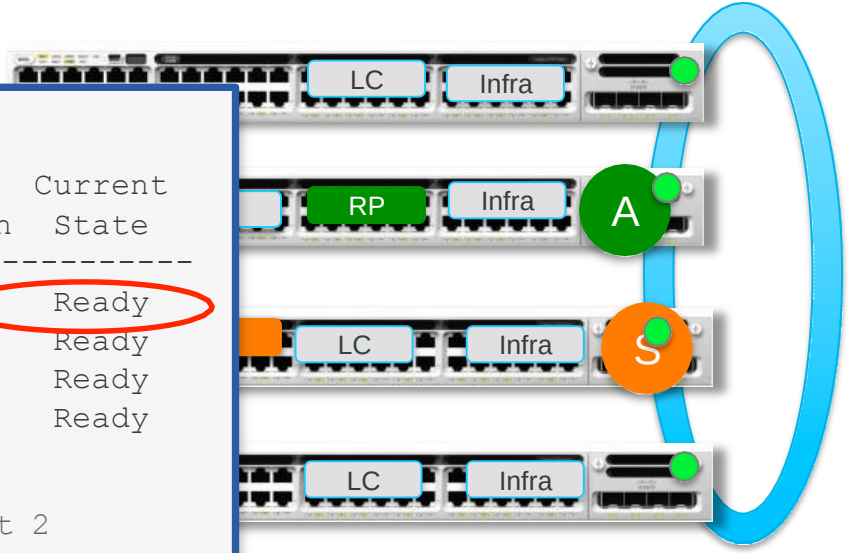

Catalyst 9400 Switchover Example

```
9400# show redundancy states

          my state = 13 -ACTIVE
        peer state =  8 -STANDBY HOT
              Mode = Duplex
           Unit ID = 7
```

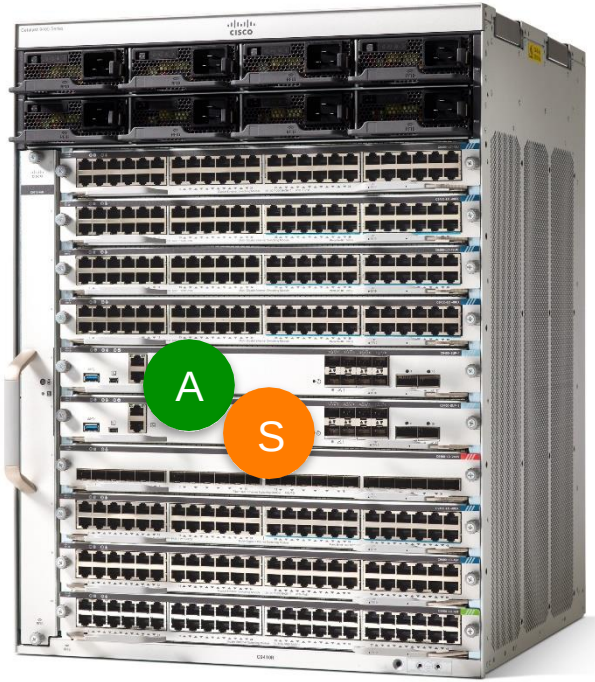

Catalyst 9400 Switchover Example

- Active Unit Goes Down
- Standby Takes over as Active

```
9400# show redundancy states

        my state = 13 -ACTIVE
      peer state =  1 -DISABLED
            Mode = Simplex
         Unit ID = 8
```

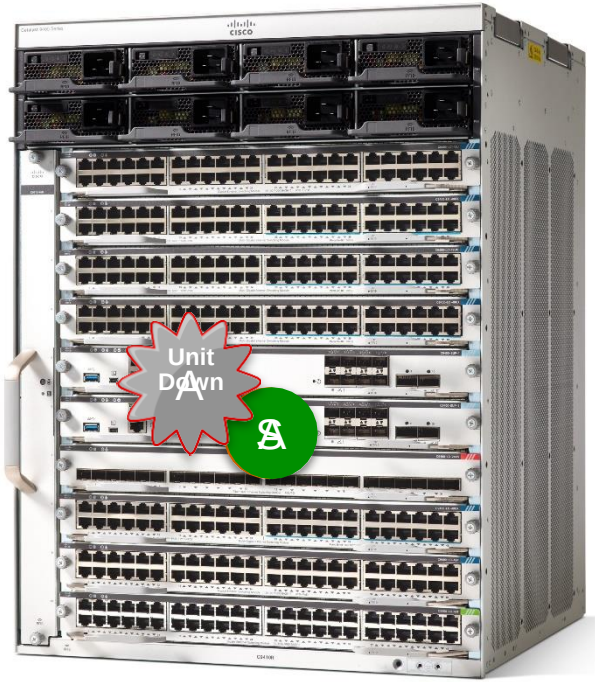

Catalyst 9400 Switchover Example

- Active Unit Goes Down
- Standby Takes over as Active
- Standby Unit Boot up
- Standby Start Bulk sync with Active RP
- Standby Reaches “Standby Hot”

```
9400# show redundancy states

         my state = 13 -ACTIVE
       peer state =  8 -STANDBY HOT
             Mode = Duplex
          Unit ID = 8
```

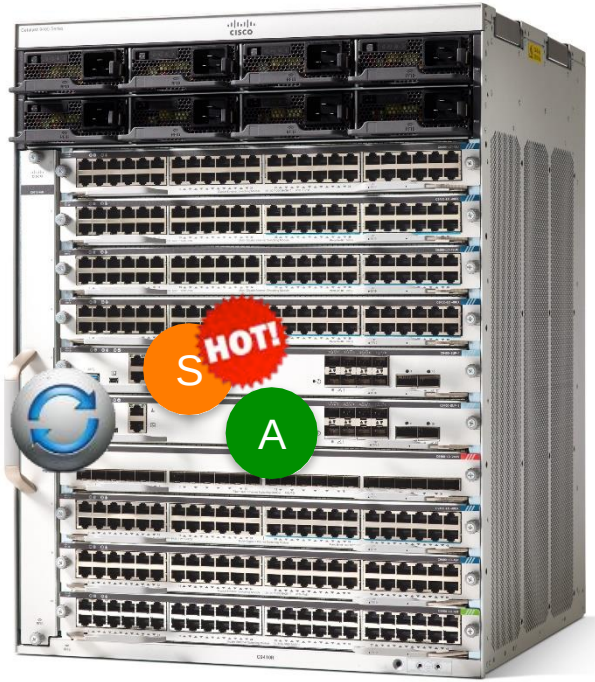

IOS XE Patchability/SMU

Software Maintenance Update (SMU)

- SMU (Software Maintenance Upgrade) is an emergency point fix positioned for expedited delivery to a customer in case of a network down or revenue affecting scenario. SMUs are:
  - Quick (able to deliver point fixes much faster than possible in IOS)
  - Effective (does not require a monolithic code upgrade)
  - Focused (target the specific area of code which has the issue)

- SMU is effectively like a **medication**:
  - It addresses the issue effectively.
  - In theory, there is no limit to the number you can take.
  - In practice, you want to be selective when SMU'ing

Why SMUs are needed?

Software Upgrades are Challenging

Cost

- Expensive Upgrades - Business Loss
- Each device upgrade causes Network outage

Time

- Reduced IT staff slows software roll out
- Physical presence required

Scope

- New Code requires bug analysis, certification

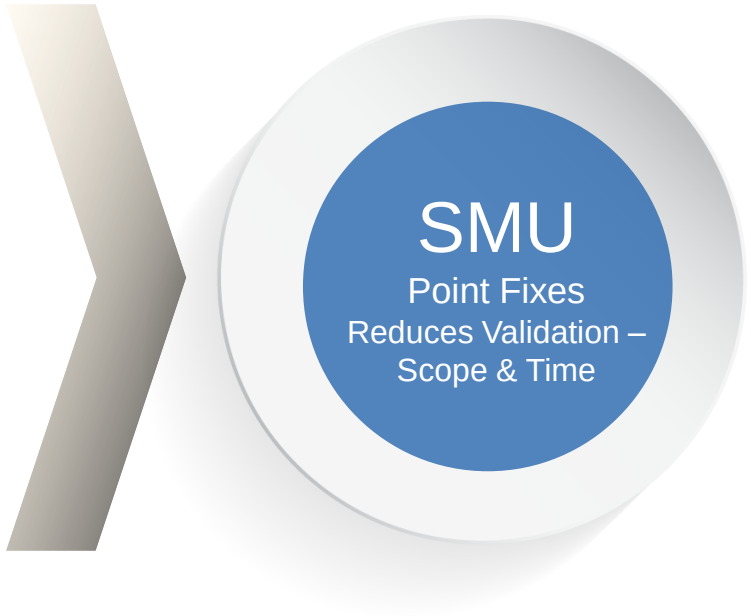

SMU Types

- Cold Patching (traffic-affecting)
  - Install of a SMU will require a system reload in the first release
- Hot Patching (non traffic-affecting)
  - Hot Restart of the patched process can be supported in the future
  - Install of a SMU will not require a system reload
- ISSU (non traffic-affecting)
  - Install the SMU using the ISSU mechanism
  - Will be available only on switches with Redundant CPU's

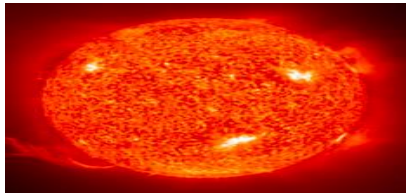

SMU Lifecycle – CLI

Switch# install remove …

SMU Committed
Memory: Process:

Switch# install commit …

SMU Removed
Memory: Process:

Switch# install deactivate …

SMU Committed
Memory: Process:

Switch# install commit …

SMU Applied
Memory: Process:

Switch# install activate …

Copy to Device
Memory: Process:

Switch# install add …

SMU Repository
SMU SMU

```
show install active
show install committed
show install inactive
show install packages
```


Software Update Creation – Work Flow

TAC/HTTS

DDTS

Fixed ?

NO

Investigate

Service Impacting ?

NO

Commit to next SW release

YES

Request Patch

SMU Council approval

Create SW Update

Documentation

Post on CCO

DE Engineer

DT Engineer

Build Engineer

Create SW update

...

Unit /Dev Test

Integration Test

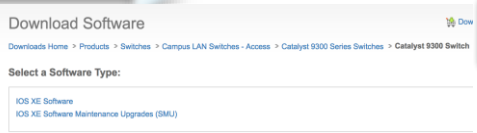

Patching Highlights

- SMUs are TAC supported.
- SMUs are synched to standby supervisor or device if auto-upgrade is enabled
- During Standby SUP or device replacement, patch will be synchronized.
- SMUs are not for feature implementation. A SMU cannot change the configuration.

SMU Management Options

Problem: SMU Life Cycle Mgmt. at Scale is a challenge with (1) Device types (2) Sw versions

There are three potential solutions

CLI	Controller (APIC-EM DNA-C)	Programmable APIs (3rd Party tools - Chef/Puppet/Ansible)
➢ Small Scale Deployments ➢ Per Device Access ➢ Full Control	➢ Mass Scale Deployments ➢ SMU Analysis ➢ SMU Life Cycle Mgmt ➢ SMU Alerts and Notification ➢ SMU Orchestration across Geo's	➢ Mass Scale Deployments ➢ Standard Programmatic Interfaces • Open Standards APIs • Consistent across multiple platforms ➢ Script Support (Shell, Perl, Python)

SMU Deployment Experience with DNA-C

- Download SMU to APIC-EM file server
- Analyze SMU impact
- Test SMU on Pilot setup
- Schedule SMU deployment

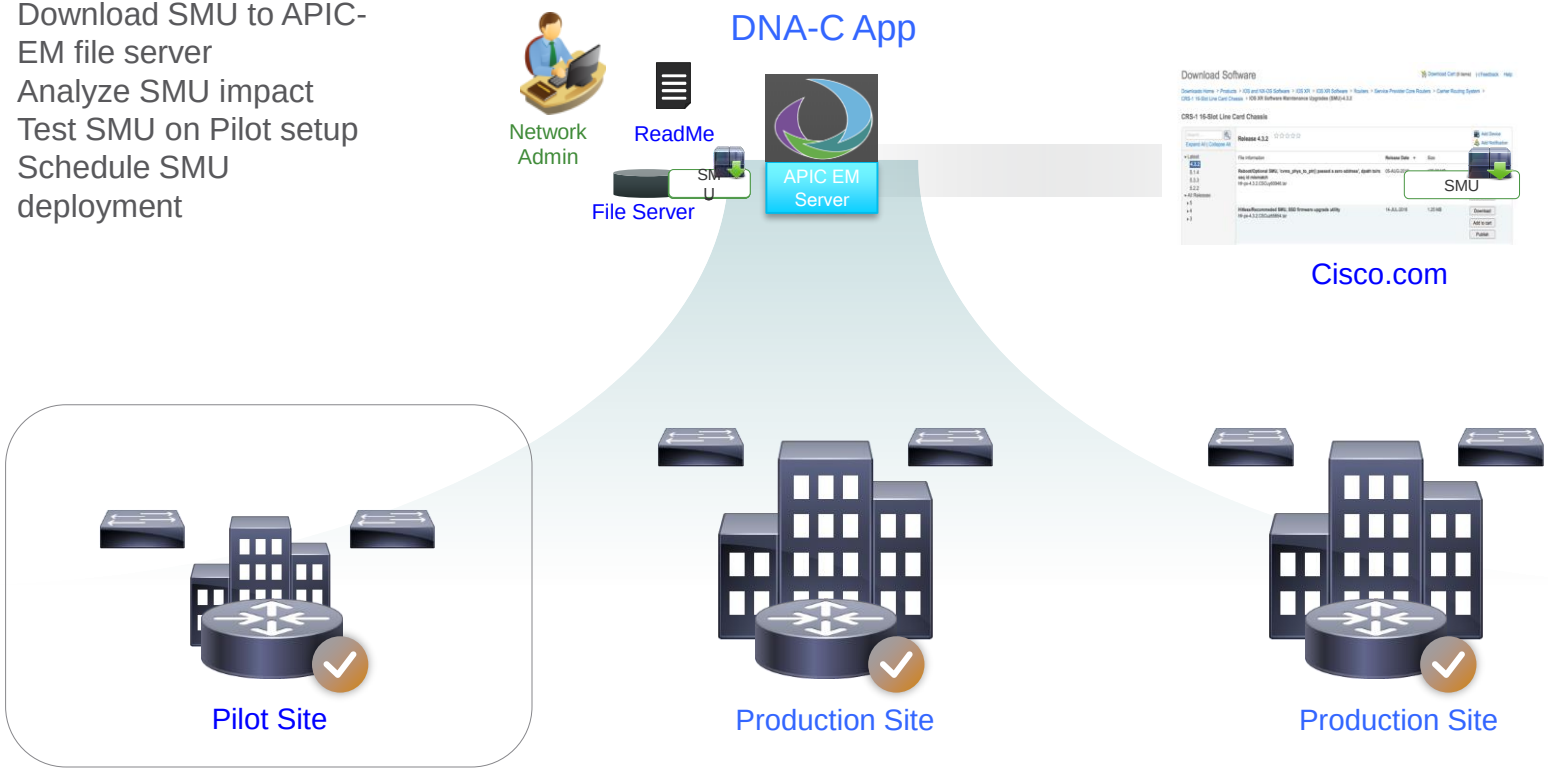

Cold Patch VS Hot Patch Components

Cold Patch Components – Reload will be Required (Will require SMU Council Approval)

Catalyst 9k All Models	
Multicast	IPV6
PIM	SSH
MVPN	ACL
ISIS	Device Sensor
BGP	CDP
RIB	SNMP
OSPF	QOS
DOT1X	DHCP
LLDP	AAA
FHS	RBAC
LISP/VXLAN	TrustSec

Hot Patch Components – Reload not Required

- From 16.6.1 Release onwards, All native yang models will be patchable without reloading the switch
- Update of existing yang models and adding new models will be supported based on case by case basis
- Operational/Config/IETF Yang models will only be supported in the future

Agenda

- High Availability Overview and Evolution
- High Availability Solution on the Campus Access
  - Stackable High Availability Solution
  - Modular High Availability Solution
- High Availability Solution on the Campus Distribution
- High Availability Solution in IOS
- Summary/Q&A

Stackwise Virtual Topology – L2 + L3

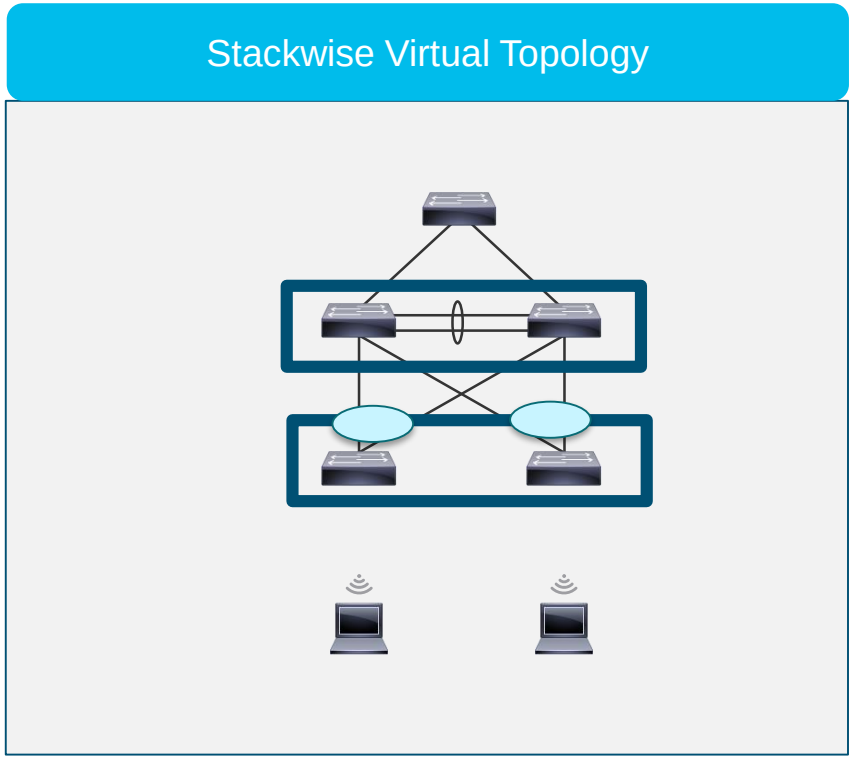

- Stackwise Virtual
- ISSU
- Stackwise-480
- Fast Software Upgrade
- Dual Supervisors
- SSO/NSF
- Patchability

Routed Access – L3

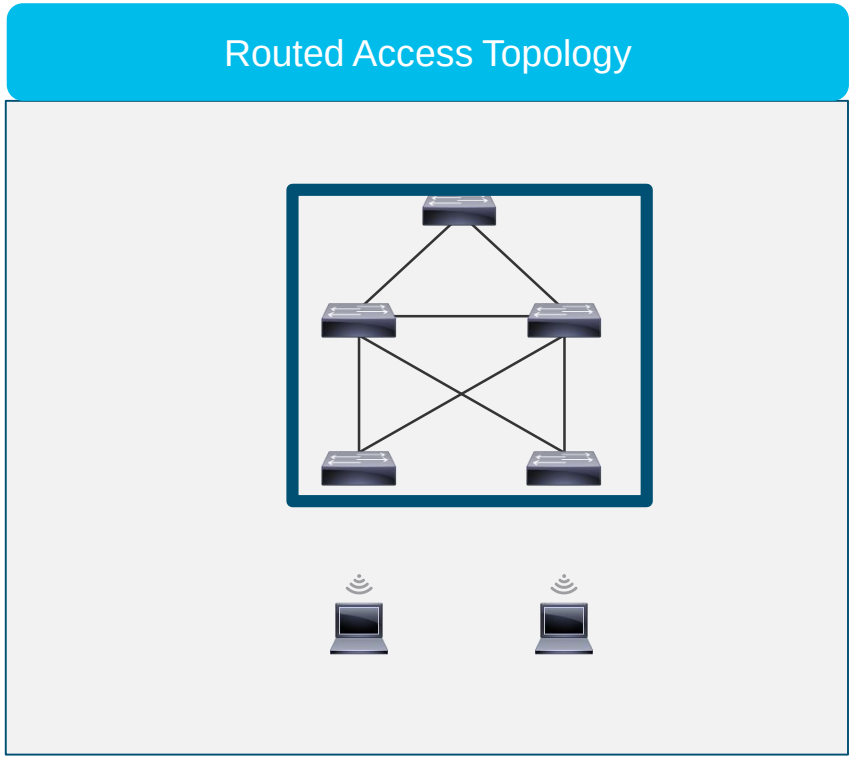

- GIR
- Stackwise-480
- Fast Software Upgrade
- Dual Supervisors
- SSO/NSF
- Patchability

Summary

- Importance of High Availability to the Access
- Feature rich and device rich deployments
- Single Points of Failures
- Best Practices on Stacking (Campus)
- Connect the switches in full ring
- Configure the Active switch priority and Standby switch priority
- Configure Active and Standby unit without uplinks where possible
- Do Not change the stack-mac timer value
- Best Practices on Modular (Campus)
- Run the system with Redundant Sups
- Split the Uplinks between the Active and Standby units in a redundant system
- Chose the right Power Redundancy Mode

Summary

- Best Practices in Enterprise Campus
- Use SVL for better availability in Distribution Layer
- Dual Connects all Devices to SVL domain
- Use multiple modules for SVL links in SVL domain
- Use separate connection for DAD Connection
- Maintenance Window timer should be configured to a reasonable value

Cisco Spark

Questions?

Use Cisco Spark to communicate with the speaker after the session

How

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

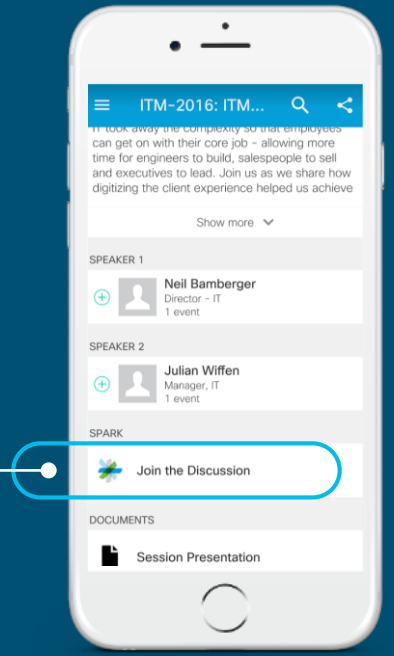

cs.co/ciscolivebot#BRKCRS-2650

Complete Your Online Session Evaluation

- Please complete your Online Session Evaluations after each session
- Complete 4 Session Evaluations & the Overall Conference Evaluation (available from Thursday) to receive your Cisco Live T-shirt
- All surveys can be completed via the Cisco Live Mobile App or the Communication Stations

Don’t forget: Cisco Live sessions will be available for viewing on-demand after the event at www.ciscolive.com/global/on-demand-library/.

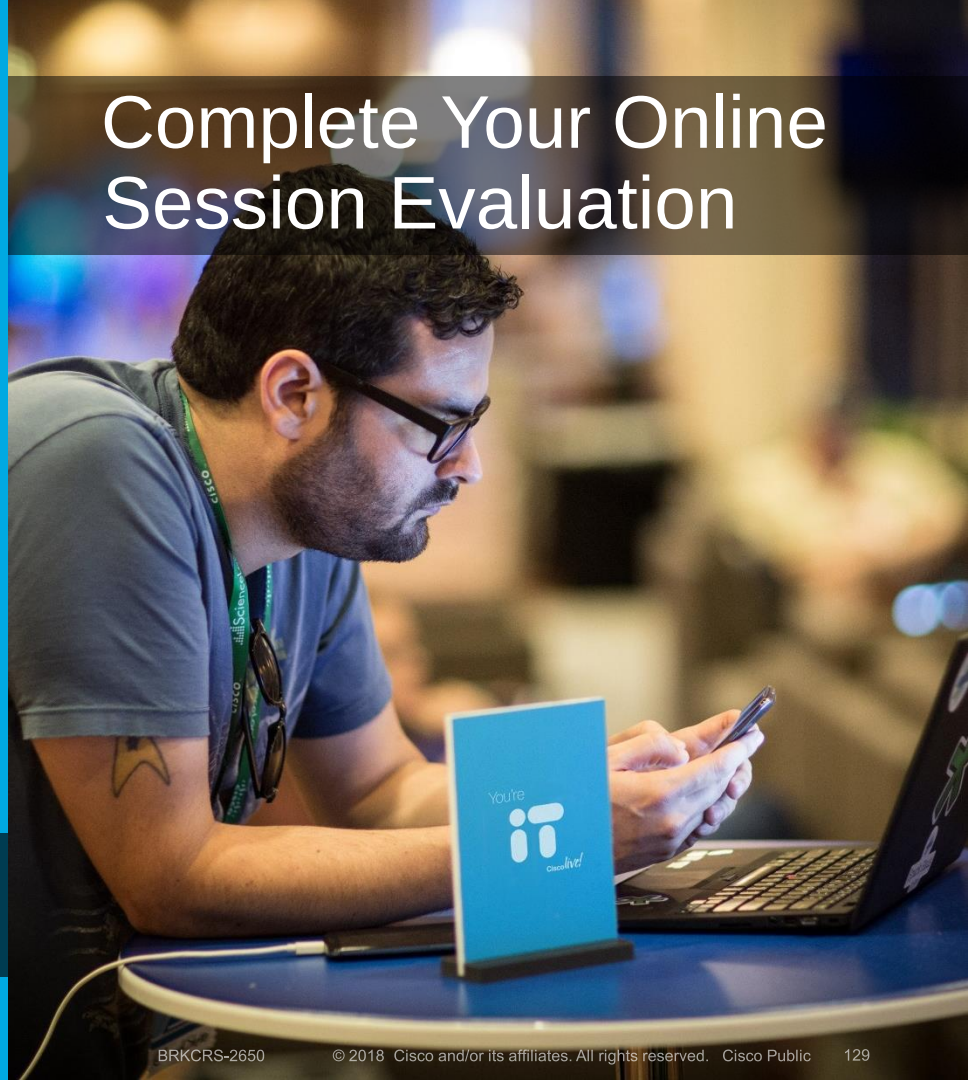

Continue Your Education

- Demos in the Cisco campus
- Walk-in Self-Paced Labs
- Tech Circle
- Meet the Engineer 1:1 meetings
- Related sessions

Thank you


You're
Cisco live!