Designing High Availability for your Network with Resiliency

Samer Theodossy
Distinguished Engineer
@SamerTheodossy

BRKDCN-2947

Cisco Webex App

Questions?

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

How

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

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

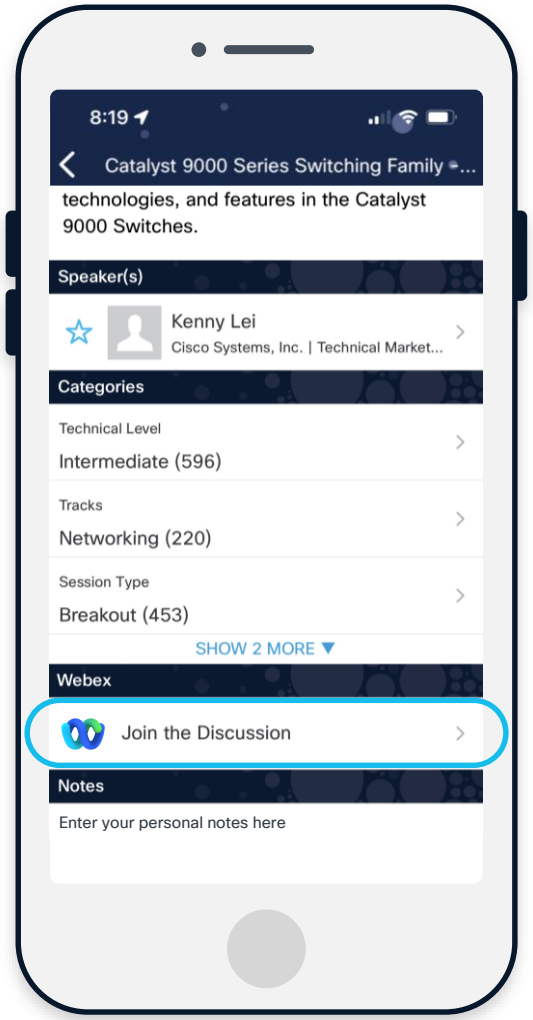

https://ciscolive.ciscoevents.com/
ciscolivebot/**#BRKDCN-2947**

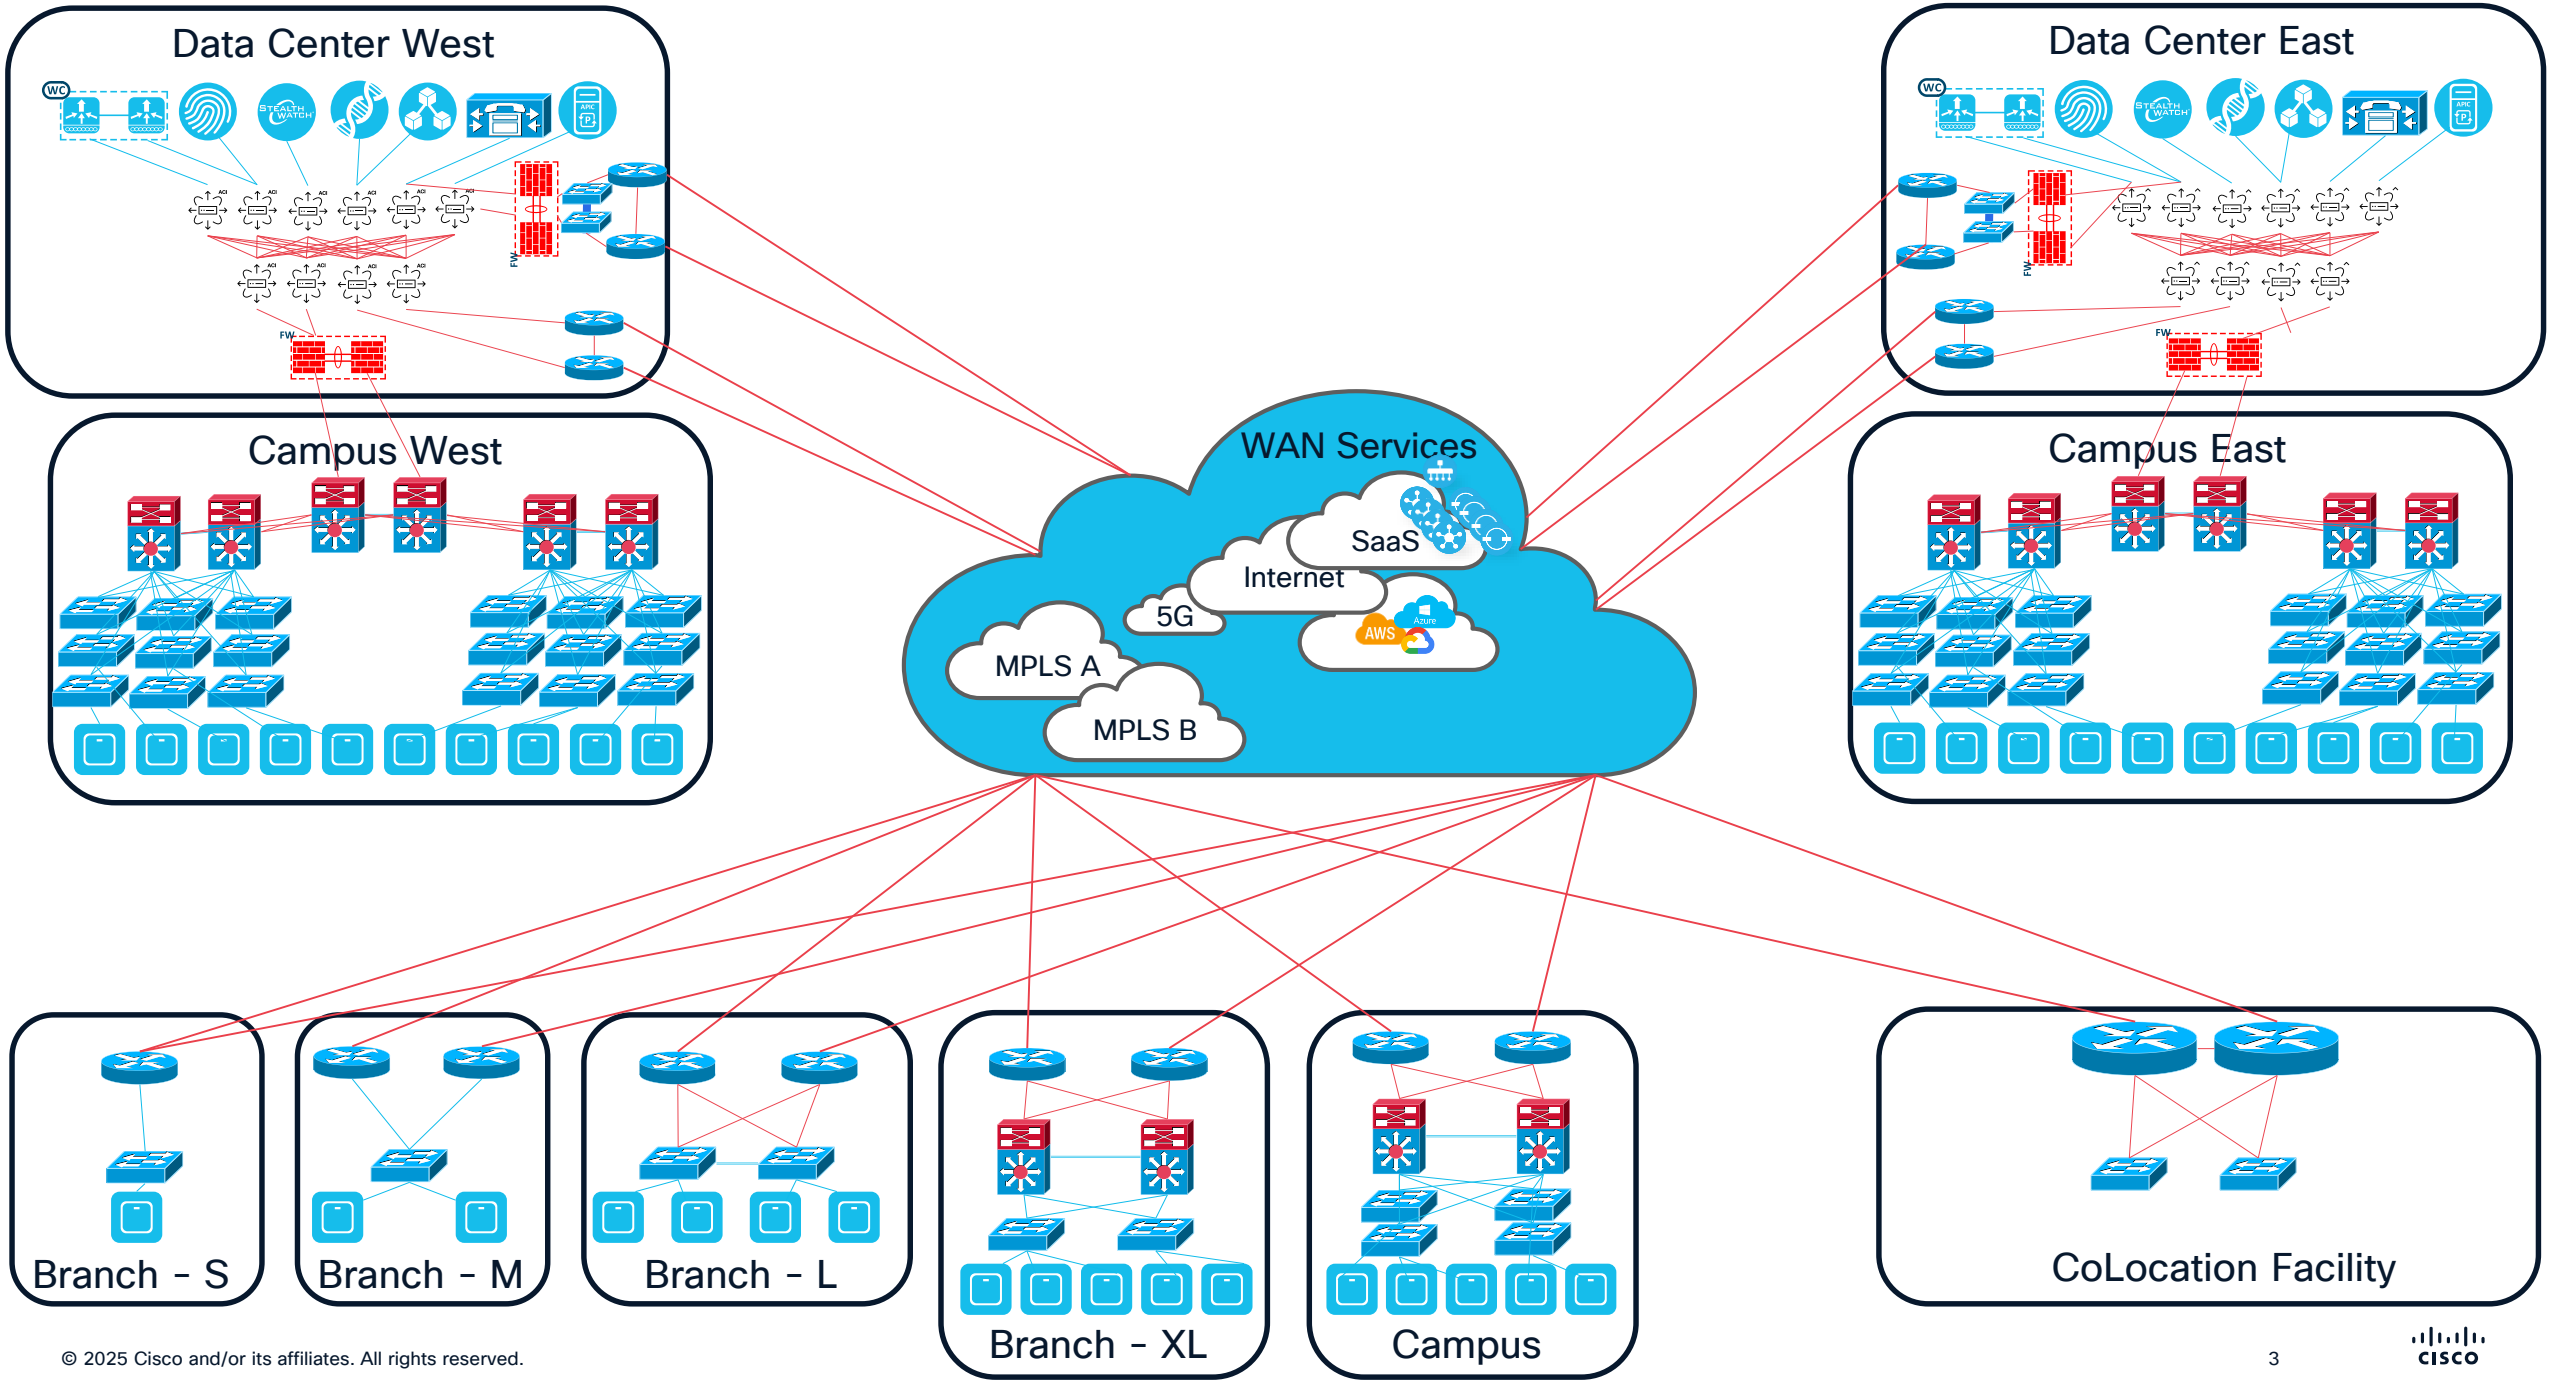

Agenda

01 Introduction

02 System Level High Availability

03 Service Level High Availability

04 Network Level Availability

05 Conclusion

Planned vs. Unplanned Outages

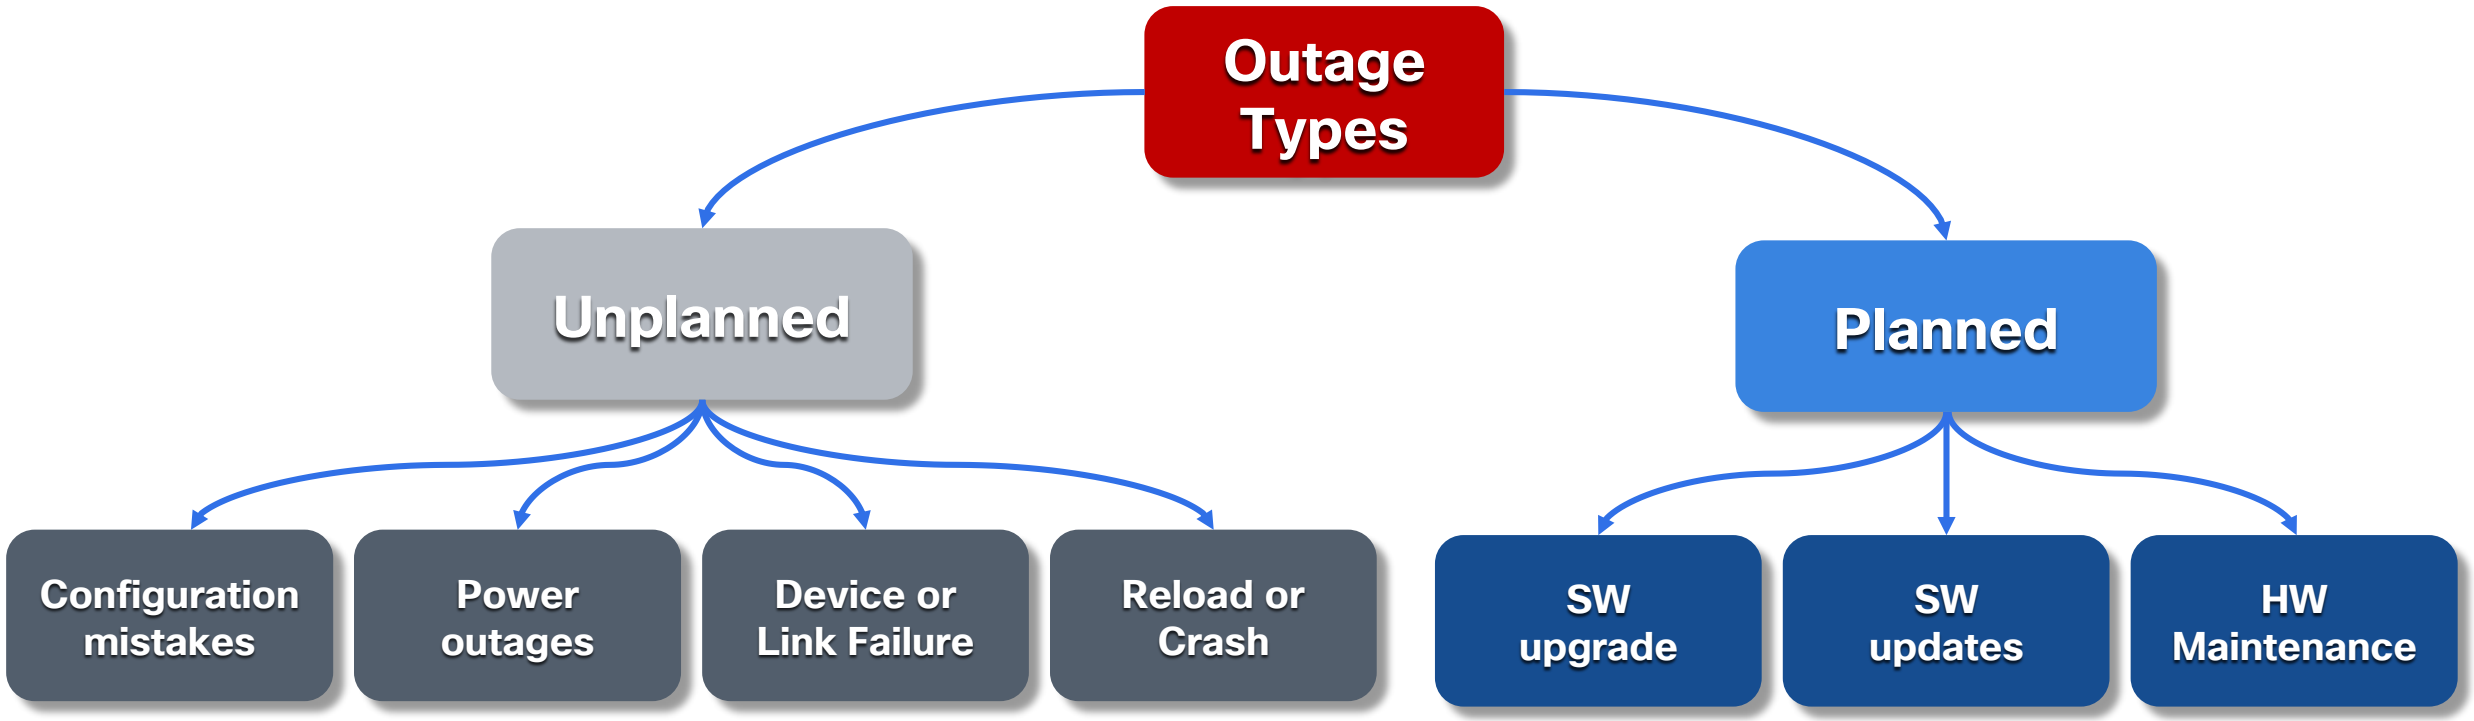

GOAL: Minimize the impact of outages on clients, network and applications

About Me

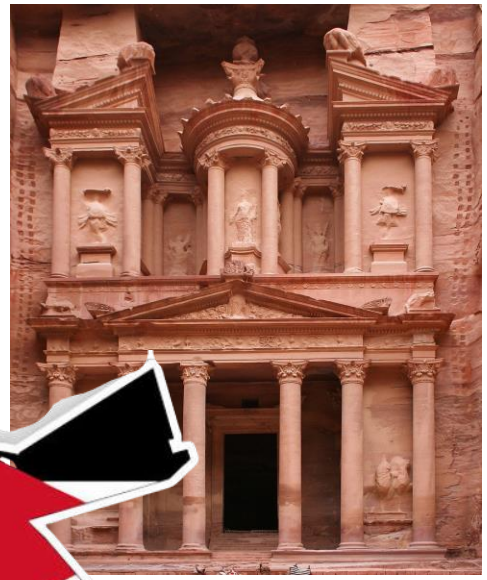

Jordan

Bullet

28+

About the speaker

Samer Theodossy – Distinguished Engineer Cloud Networking

I'm a **Distinguished Engineer** with the Cloud Networking Software team at Cisco. I've been with Cisco for 28+ years. I have been intimately involved with the design and implementation of the systems and High Availability aspect of a lot of the platforms. Well versed in the design of the SSO/NSF technology and its interworking and helped design the first ISSU solution at Cisco that has been commended with the Pioneer Award and best in class in the industry.
I was part of the team that delivered Cat3850 Architecture and integrated the High Availability Solution on this Access Product as well as on the Cisco Cat4500.
I'm here to help you learn more about to simplify your Upgrade experience and talk about modularity of the software.

I have worked on a wide spectrum of Products in Cisco (Nexus7k, Nexus3k/9k, BPX, MGX, 7600, 7500, 7300, 7200, Cat4500, C10K and the Cat9K) and Operating Systems (NXOS, ACI, Classic IOS, IOS-XE).

How can we measure a predicted availability?

- It’s function of:
- Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR)

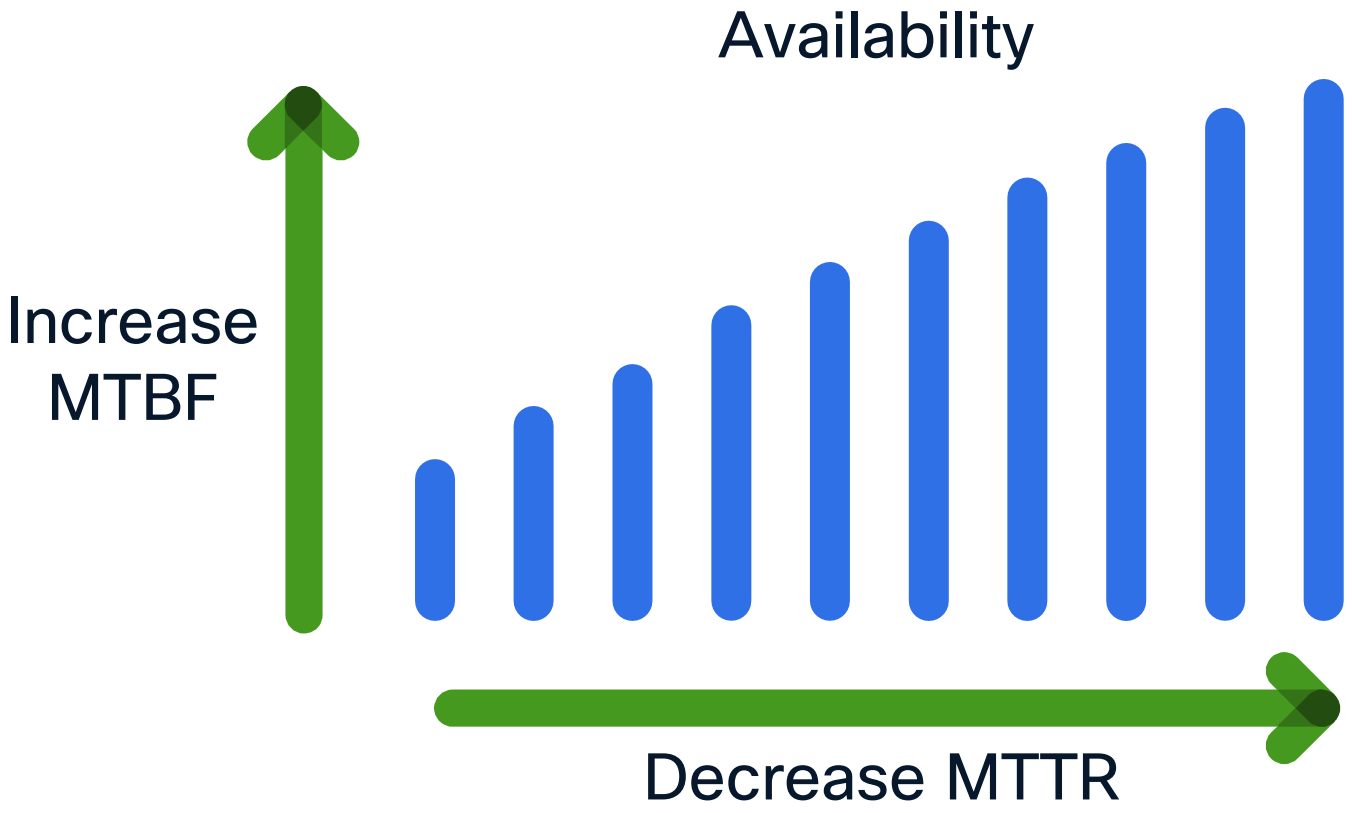

A basic predicted availability equation

Predicted Availability Equation

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

MTBF: Mean Time Between Failures

MTTR: Mean Time To Repair

Example predicted availability calculations

Component with MTBF=87,600 hours

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

24 hour depot replacement

2 hour 24 minutes predicted annual downtime

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

4 hour depot replacement

24 minutes predicted annual downtime

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

Warm spare

1 minute predicted annual downtime

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

The redundancy effect

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

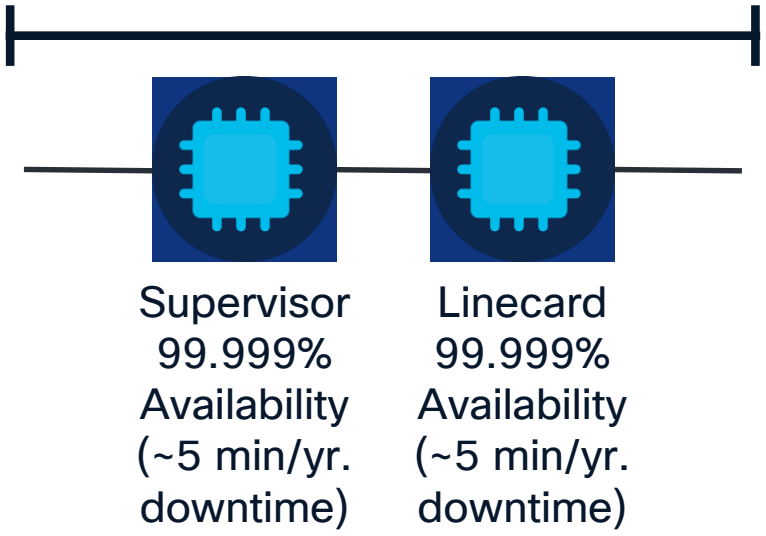

- Redundant components functioning in parallel
- System predicted availability: **99.999999%** (~½ second/year predicted downtime)

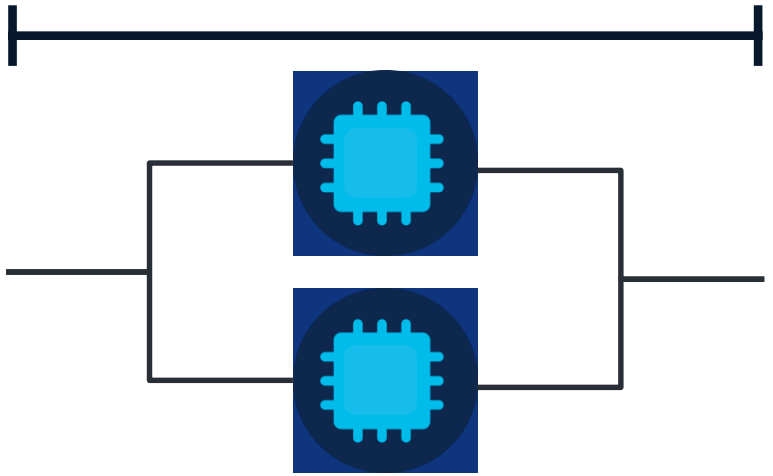

Redundancy and Resiliency

Network Level

- ❑ Redundant Links
- ❑ Redundant devices

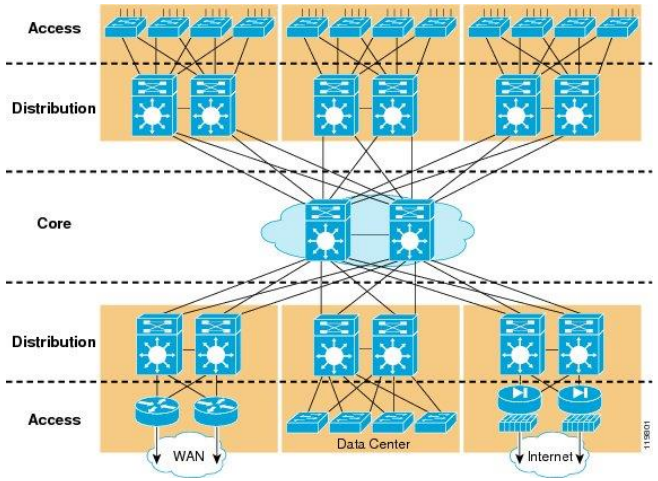

Device Level

- ❑ Integrated hardware resiliency
- ❑ Redundant power supplies

Software Level

- ❑ Recovery from Failures

Data Center
High Availability Design

Data Center High Availability

Reducing downtime for Upgrades and Unplanned Events

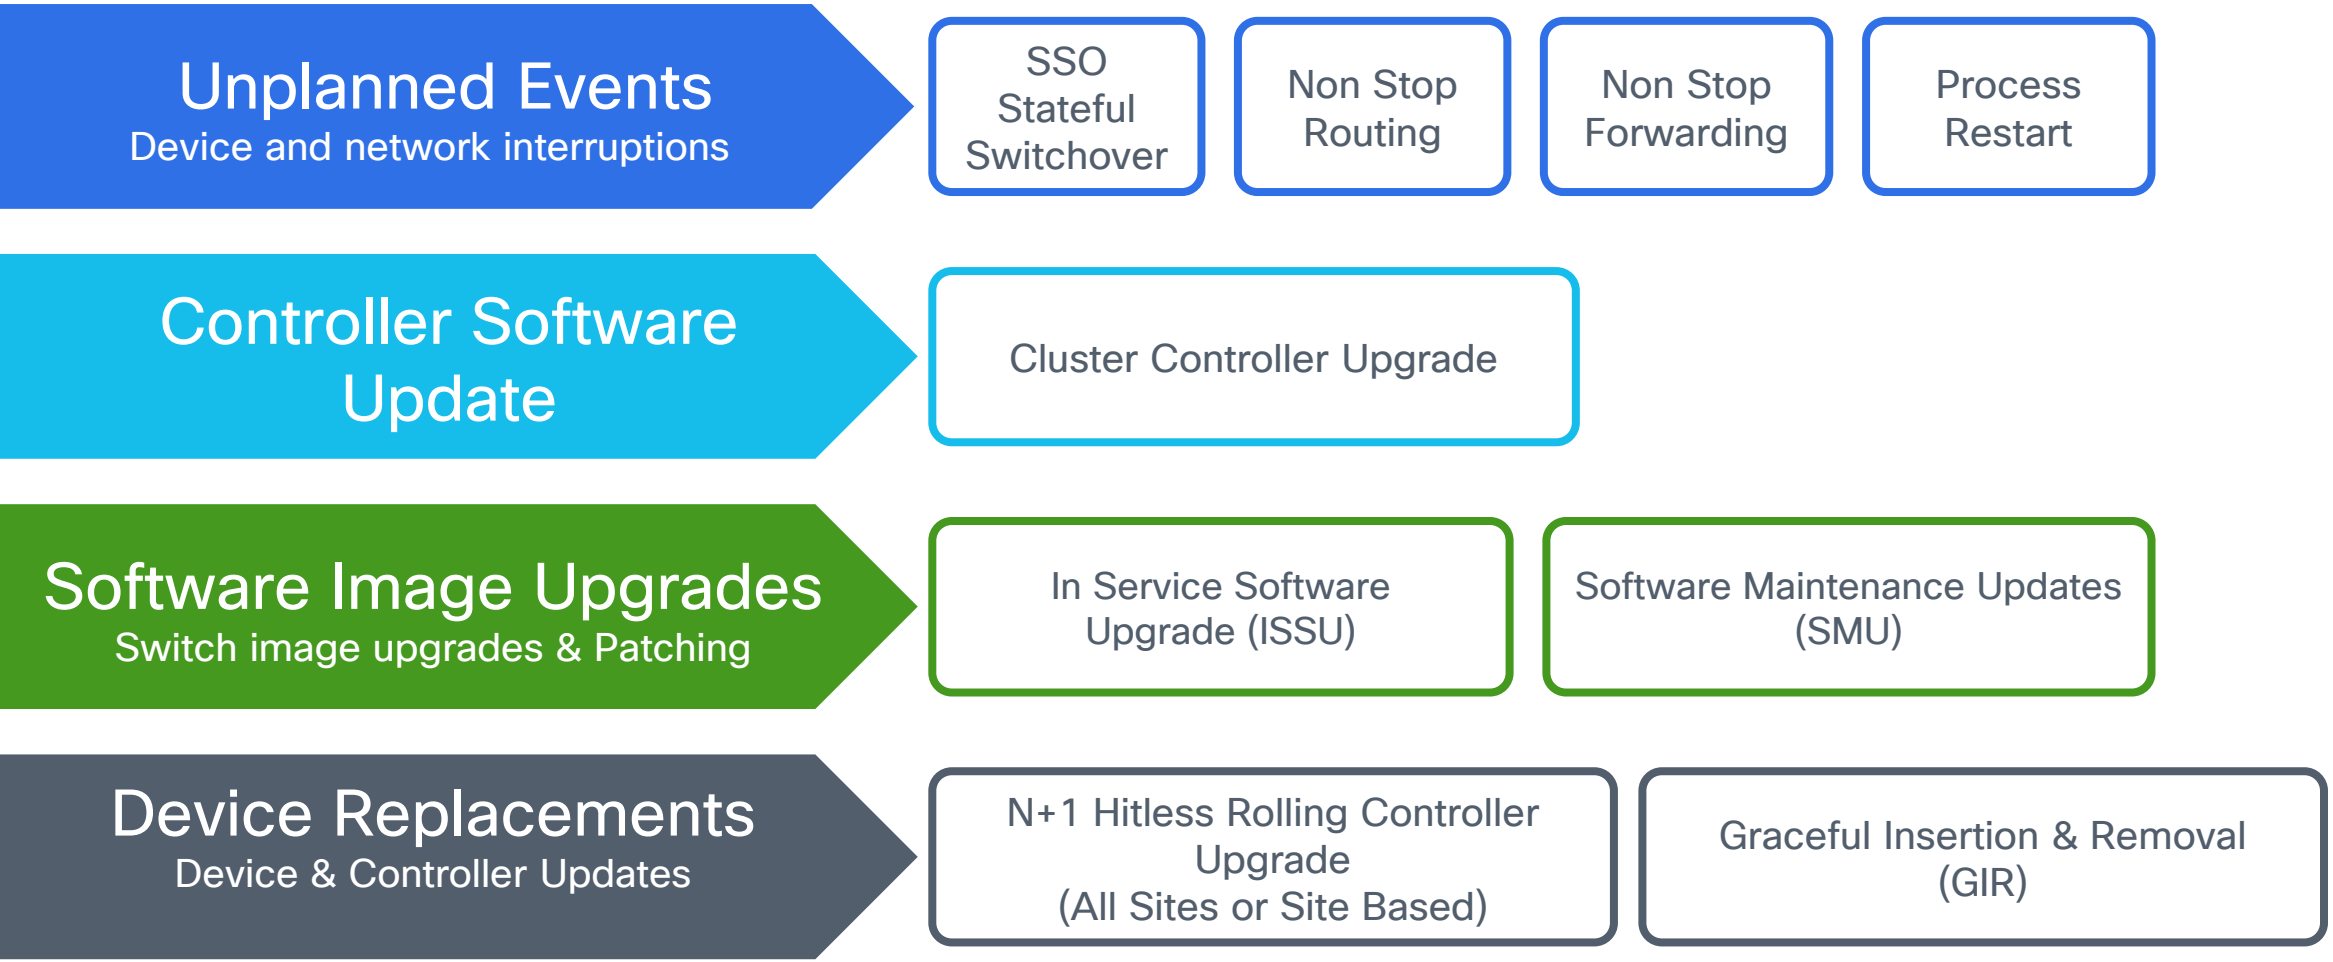

Planning is Key

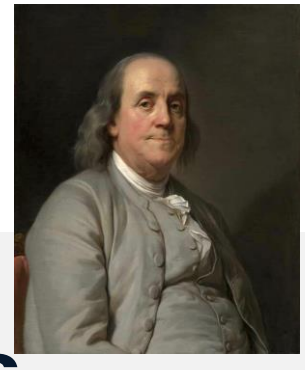

- **Which Services is more Important For my Business Processes?**
  - What HA level are you designing for?
- **What Can I Do to Mitigate Most of the Possible Risks?**
  - What is your budget?
  - Figure out where redundant components are needed
    - More is not always better
    - Keep it simple

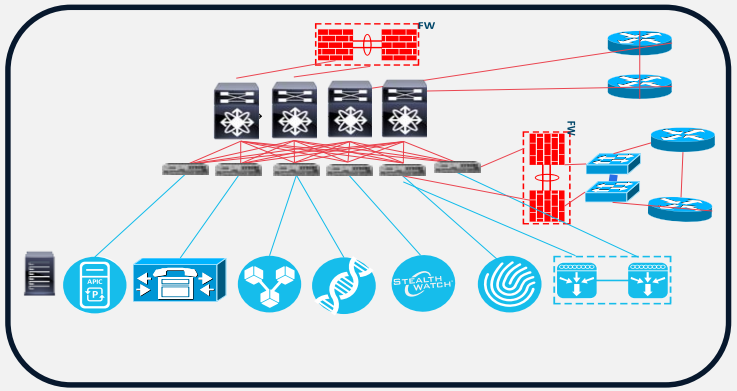

DC Network Design HA Principle

High availability using modularity, hierarchy, and structure

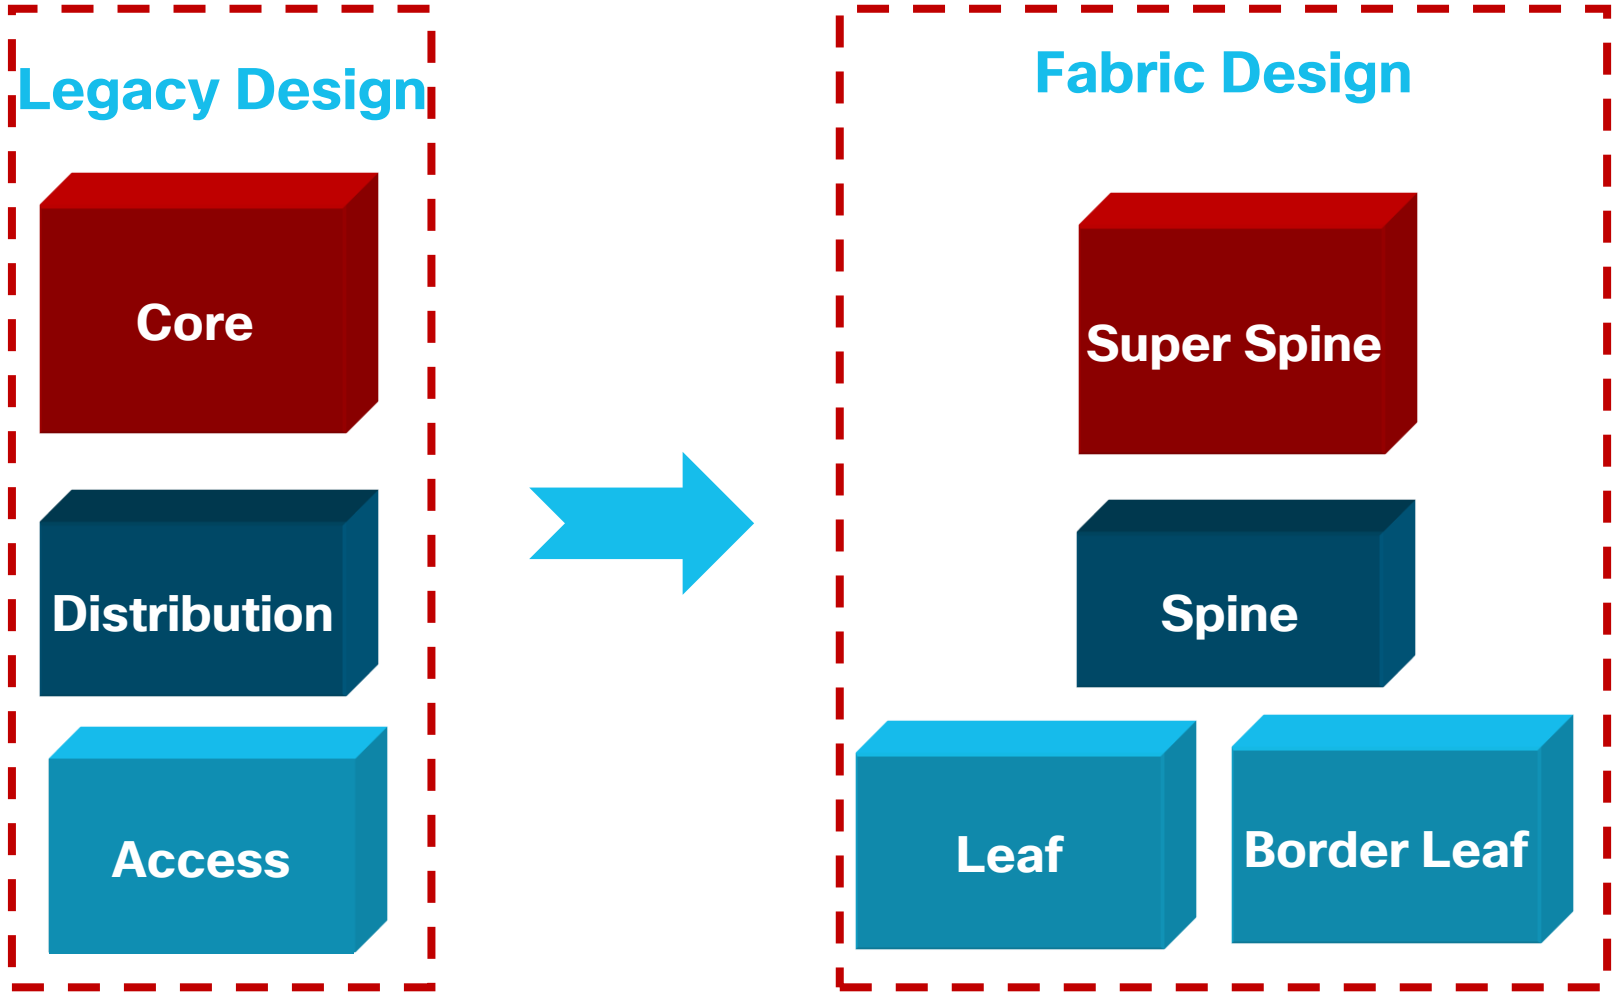

Unplanned

Physical Redundancy

Cisco Nexus Switches

	Nexus 9300 Nexus 9400	Nexus 9500 Nexus 9800
Type	1,2 and 4 RU	Modular Chassis
Physical Redundancy	Fan Tray and Power Supply Redundancy	Power supply, Fan tray, Switch fabric, System controller, Supervisor module

Redundant Hardware – Front Panel

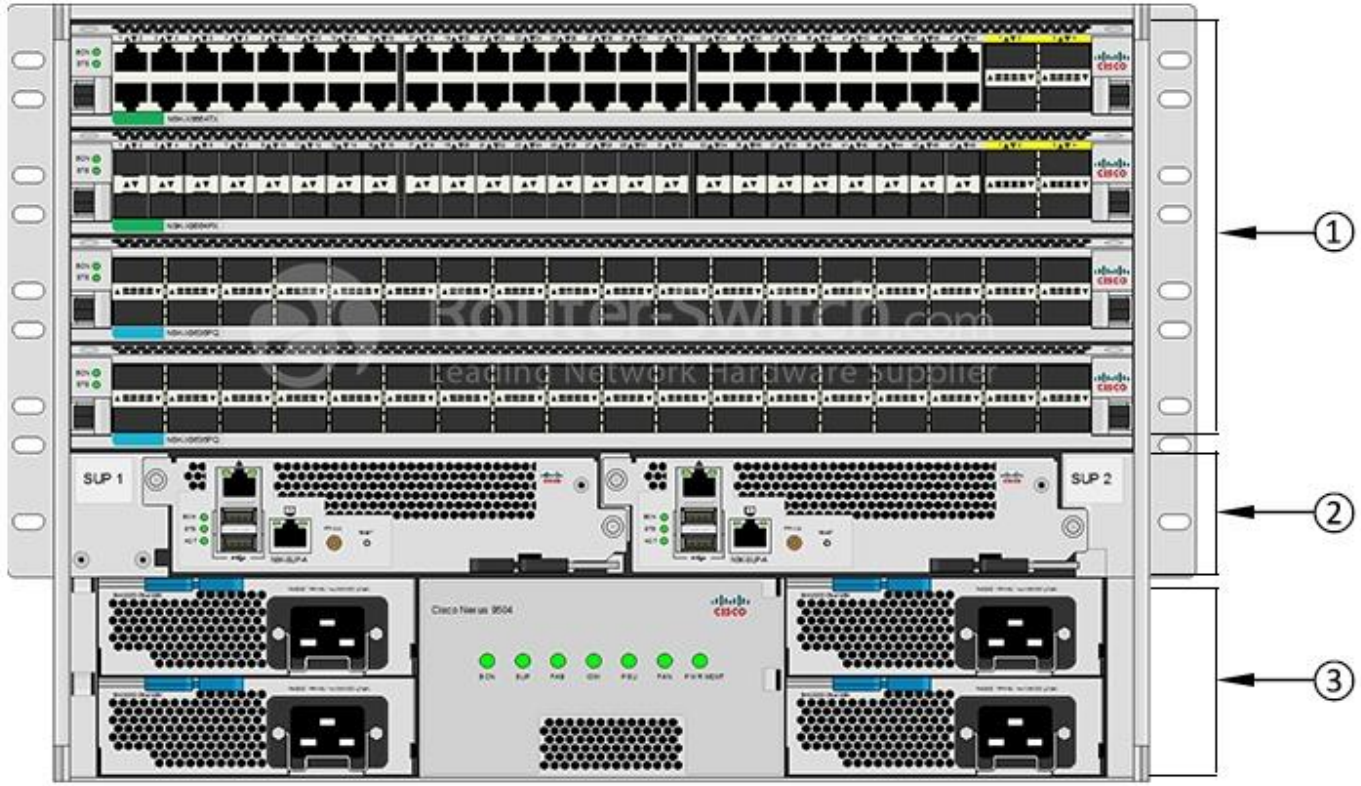

Line Cards

Redundant Supervisors

Power Supplies

Highly Available Architecture with Nexus Switches

Nexus 9500 Modular Series

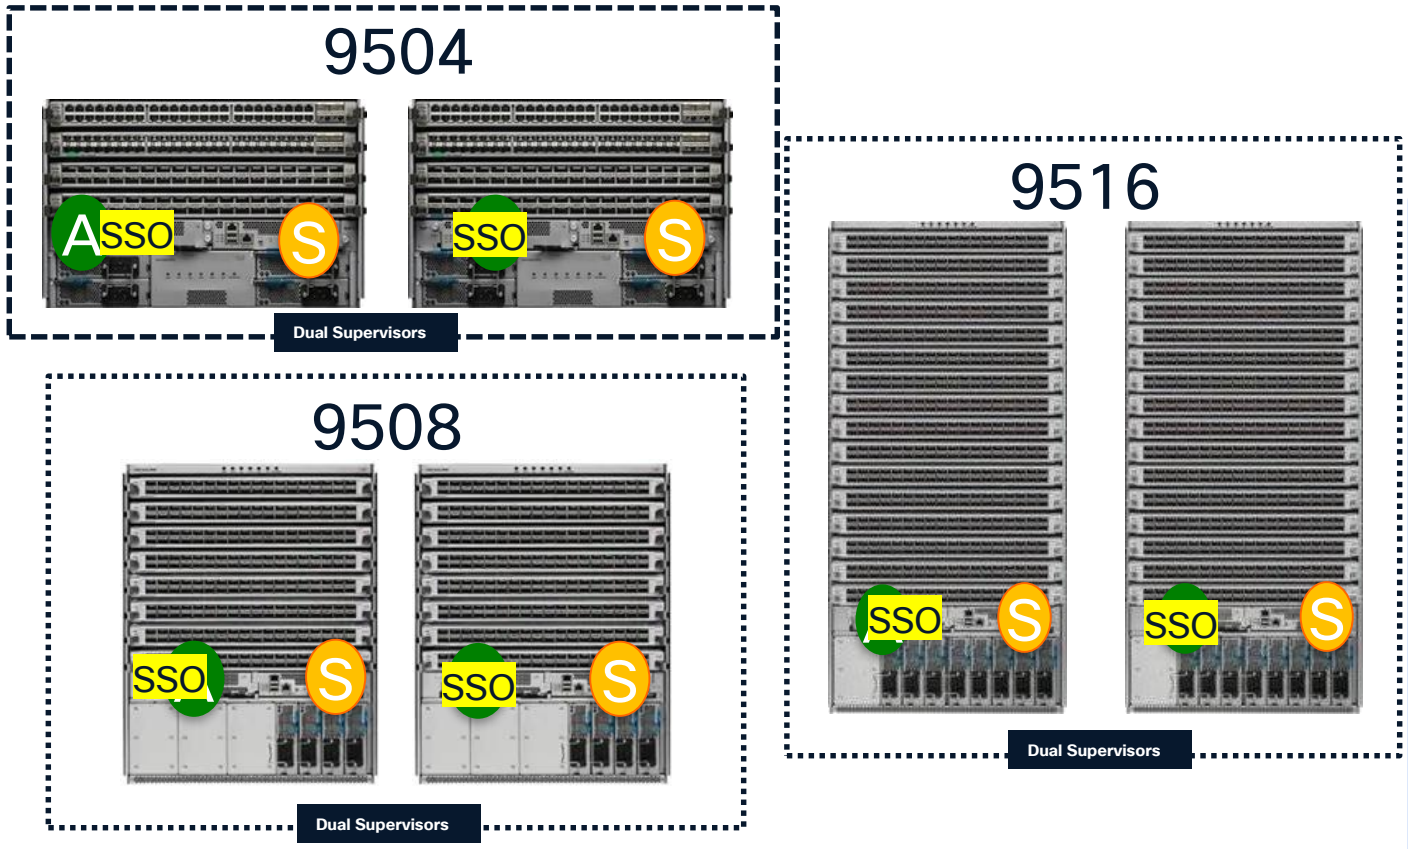


```
switch# show system redundancy status

Redundancy mode
---------------
          administrative: HA
             operational: HA

This supervisor (sup-1)
-----------------------
          Redundancy state: Active
          Supervisor state: Active
            Internal state: Active with HA standby

Other supervisor (sup-2)
------------------------
          Redundancy state: Standby
          Supervisor state: HA standby
            Internal state: HA standby
```

Power Mode Architecture with Nexus Switches

Combined Power Mode

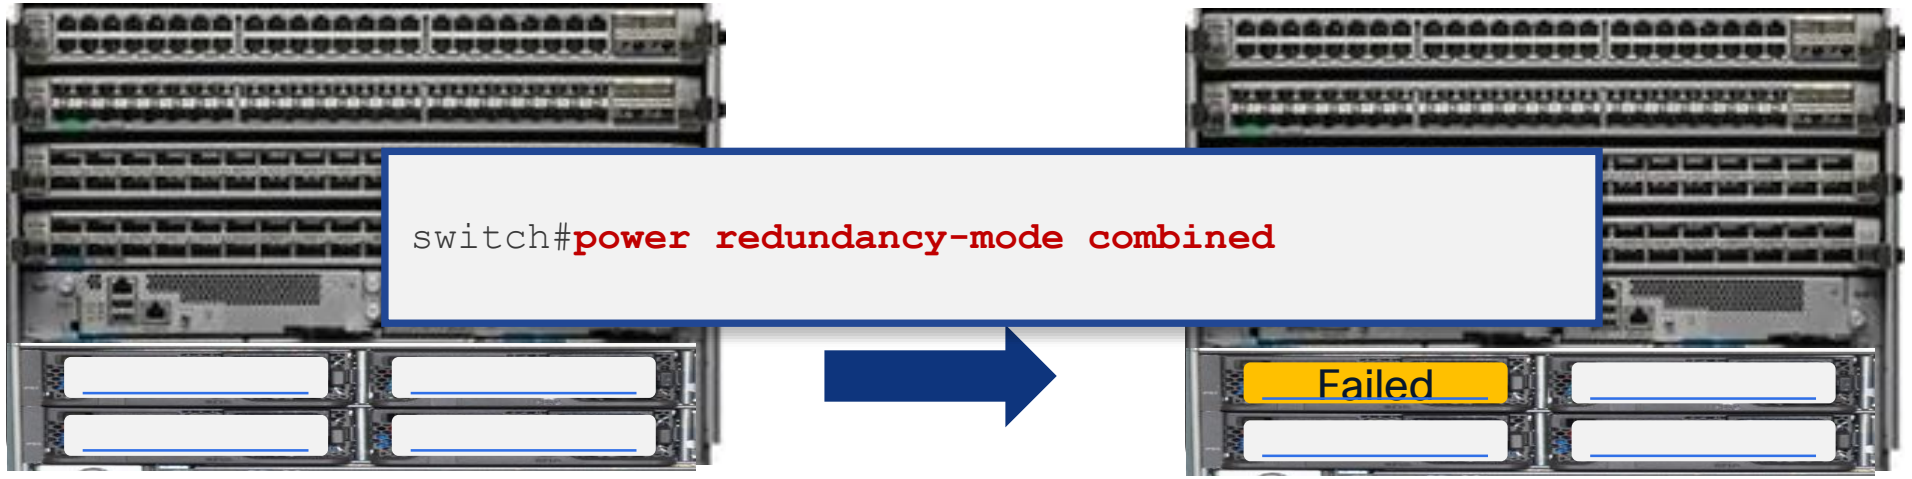

Load sharing on all PSs

Load sharing on all Active PSs

Redundant Power Mode

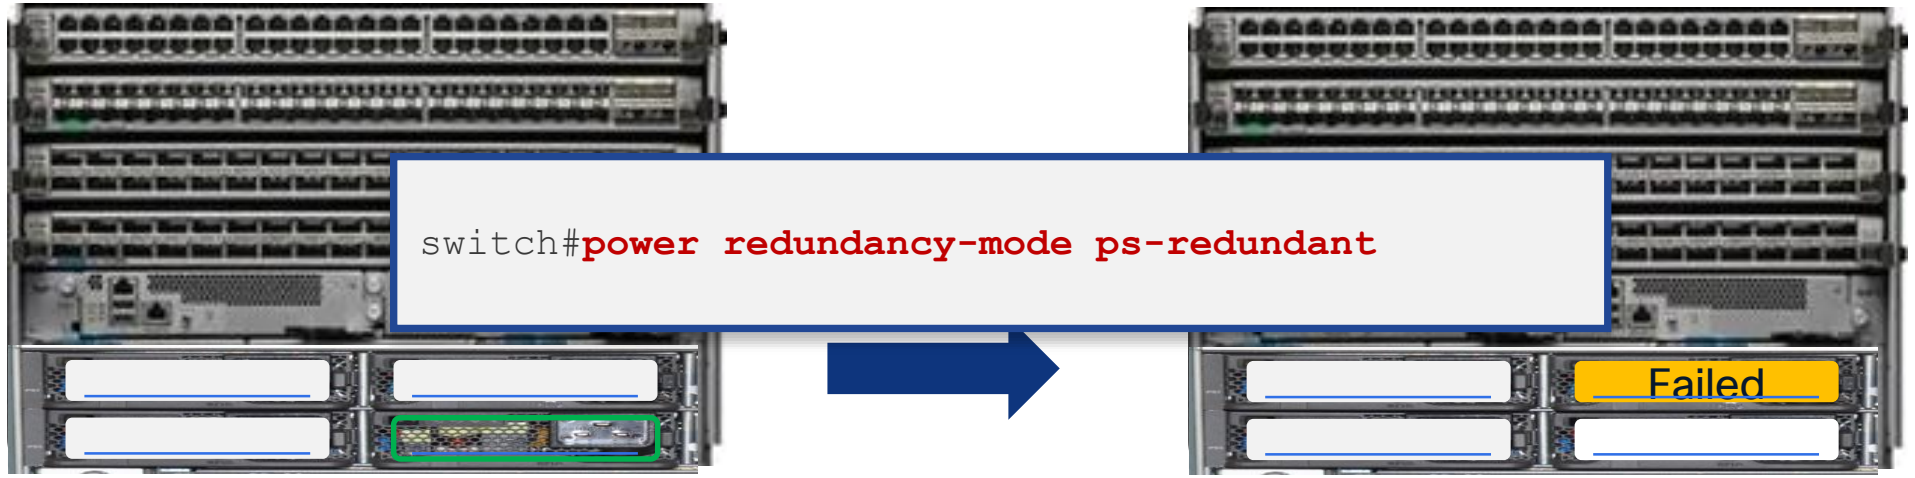

Load sharing on all Active PSs
Standby PS is disabled

Standby PS becomes Active

Grid Redundancy

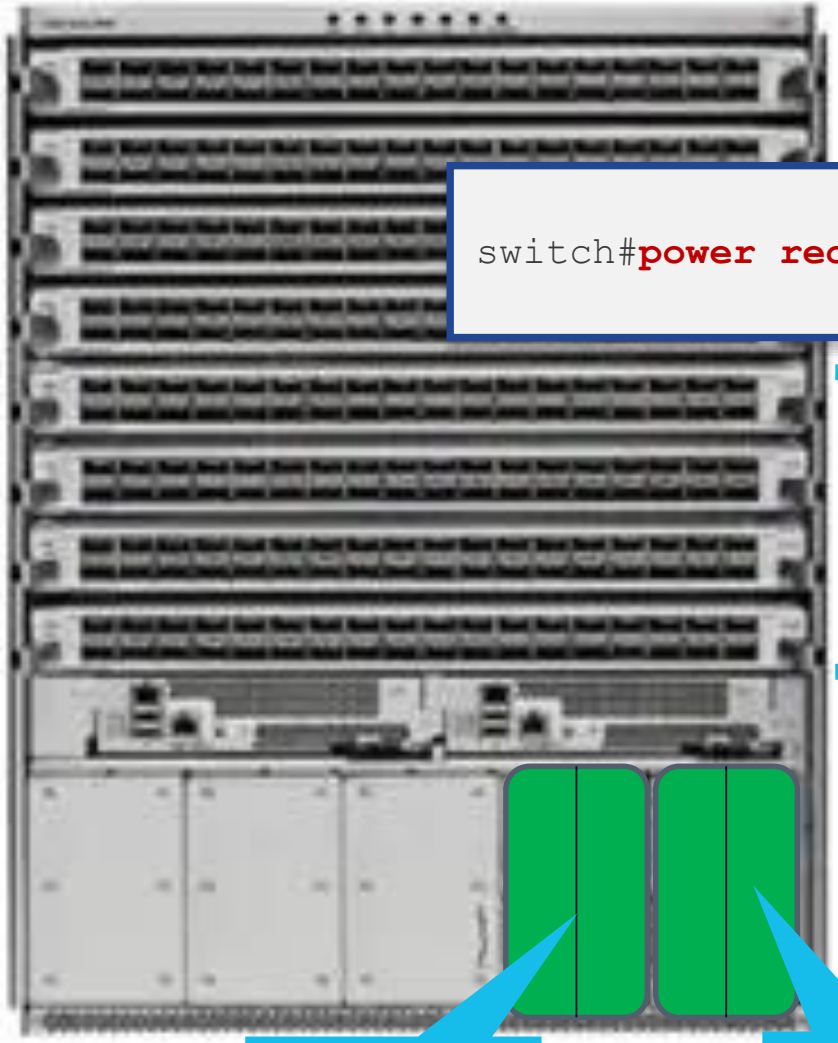

LCs 1-4 connect to Grid A

```
switch#power redundancy-mode insrc_redundant
```

LCs 5-8 connect to Grid B

Grid A

Grid B

Power sharing and redundancy are available only within individual Grid

Redundant Hardware – Back Panel

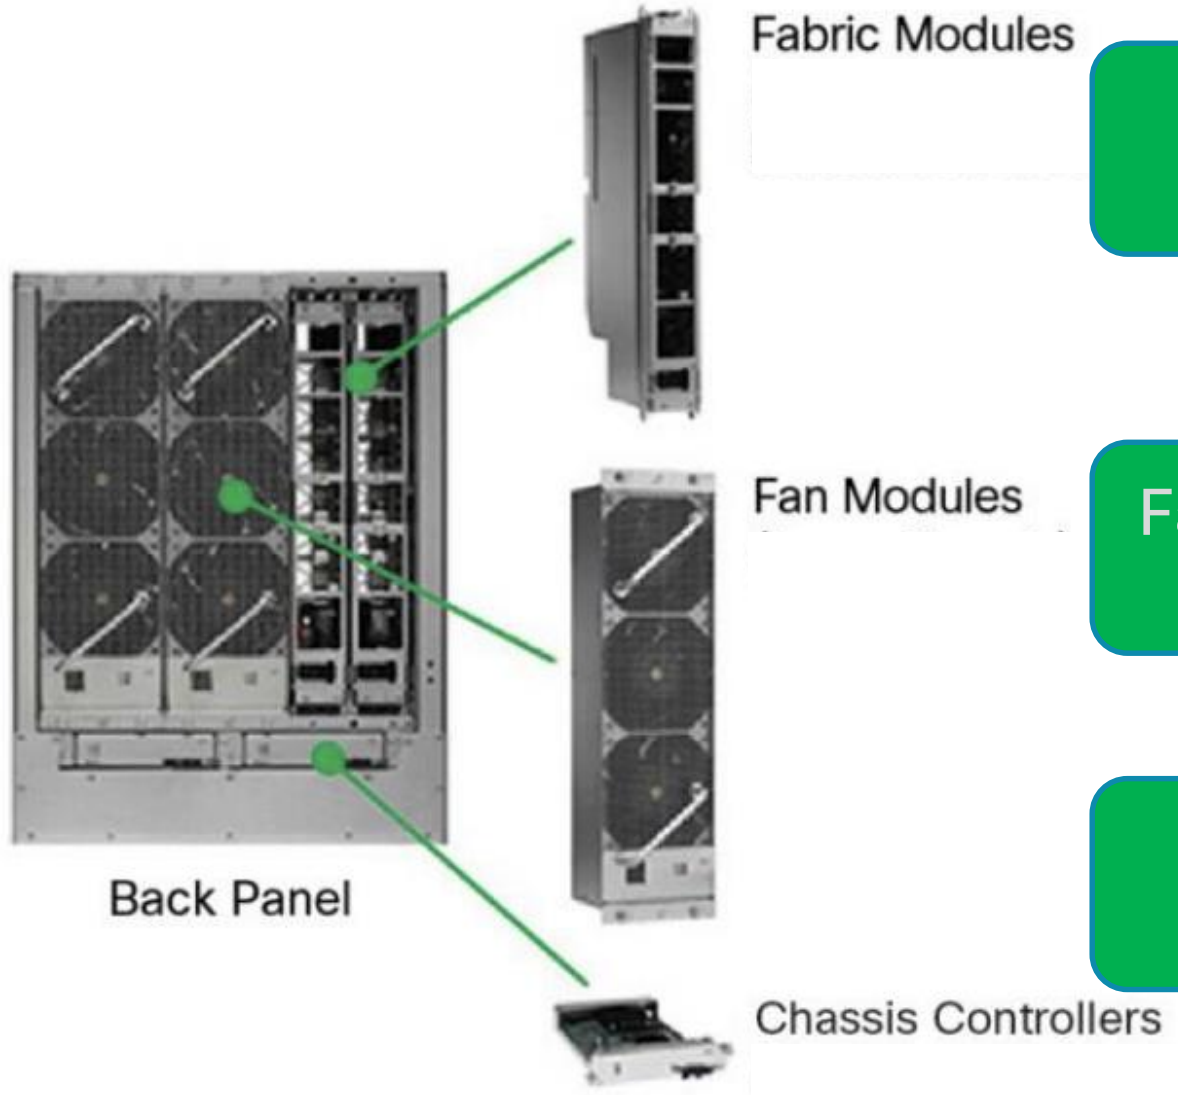

Help in distribution of traffic

Fan Speed Adjust depending on fan failures occurring

Offload the power and fan management
And handle EOBC as well

Show commands outputs

command output (NX 9504)

```
switch# show module

Mod  Ports     Module-Type                            Model              Status
---  -----     -------------------------------------  -----------------  ----------
1    32   32p 40G Ethernet Module                     N9K-X9432PQ        ok
2    52   48x1/10G SFP+ 4x40G Ethernet Module         N9K-X9564PX        ok
5    52   48x1/10G SFP+ 4x40G Ethernet Module         N9K-X9464PX        ok
6    36   36p 40G Ethernet Module                     N9K-X9536PQ        ok
7    36   36p 40G Ethernet Module                     N9K-X9536PQ        ok
10   32   32p 40G Ethernet Module                     N9K-X9432PQ        ok
11   52   48x1/10G-T 4x40G Ethernet Module            N9K-X9564TX        ok
12   52   48x1/10G-T 4x40G Ethernet Module            N9K-X9464TX        ok
15   52   48x1/10G SFP+ 4x40G Ethernet Module         N9K-X9464PX        ok
21   0    Fabric Module                               N9K-C9516-FM       ok
22   0    Fabric Module                               N9K-C9516-FM       ok
23   0    Fabric Module                               N9K-C9516-FM       ok
24   0    Fabric Module                               N9K-C9516-FM       ok
25   0    Fabric Module                               N9K-C9516-FM       ok
26   0    Fabric Module                               N9K-C9516-FM       ok
27   0    Supervisor Module                           N9K-SUP-A          ha-standby
28   0    Supervisor Module                           N9K-SUP-A          active *
29   0    System Controller                           N9K-SC-A           active
30   0    System Controller                           N9K-SC-A           standby
```

Show commands outputs

command output

```
switch# show environment power
Power Supply:
Voltage: 12 Volts
Power                        Actual             Actual        Total
Supply    Model              Output             Input         Capacity     Status
                             (Watts )           (Watts )      (Watts )
-------  ----------  ----------------  ------  ----------  ----------  ----------
1        NXA-PAC-1100W-PE2        0 W                0 W         0 W    Shutdown
2        NXA-PAC-1100W-PE2      254 W              262 W      1100 W       Ok


Power Usage Summary:
--------------------
Power Supply redundancy mode (configured)              PS-Redundant
Power Supply redundancy mode (operational)             Non-Redundant

Total Power Capacity (based on configured mode)             1100.00 W
Total Grid-A (first half of PS slots) Power Capacity           0.00 W
Total Grid-B (second half of PS slots) Power Capacity       1100.00 W
Total Power of all Inputs (cumulative)                      1100.00 W
Total Power Output (actual draw)                             254.00 W
Total Power Input (actual draw)                              262.00 W
Total Power Allocated (budget)                                 N/A
Total Power Available for additional modules                   N/A

switch09-n9324#
```

NXOS SW Architecture

NX-OS HA Architecture

- Fully distributed modular design
- Control-plane & data-plane separation
- Service restart-ability
- Non-disruptive SSO* & ISSU
- Open & Model Driven
- App Hosting Capability

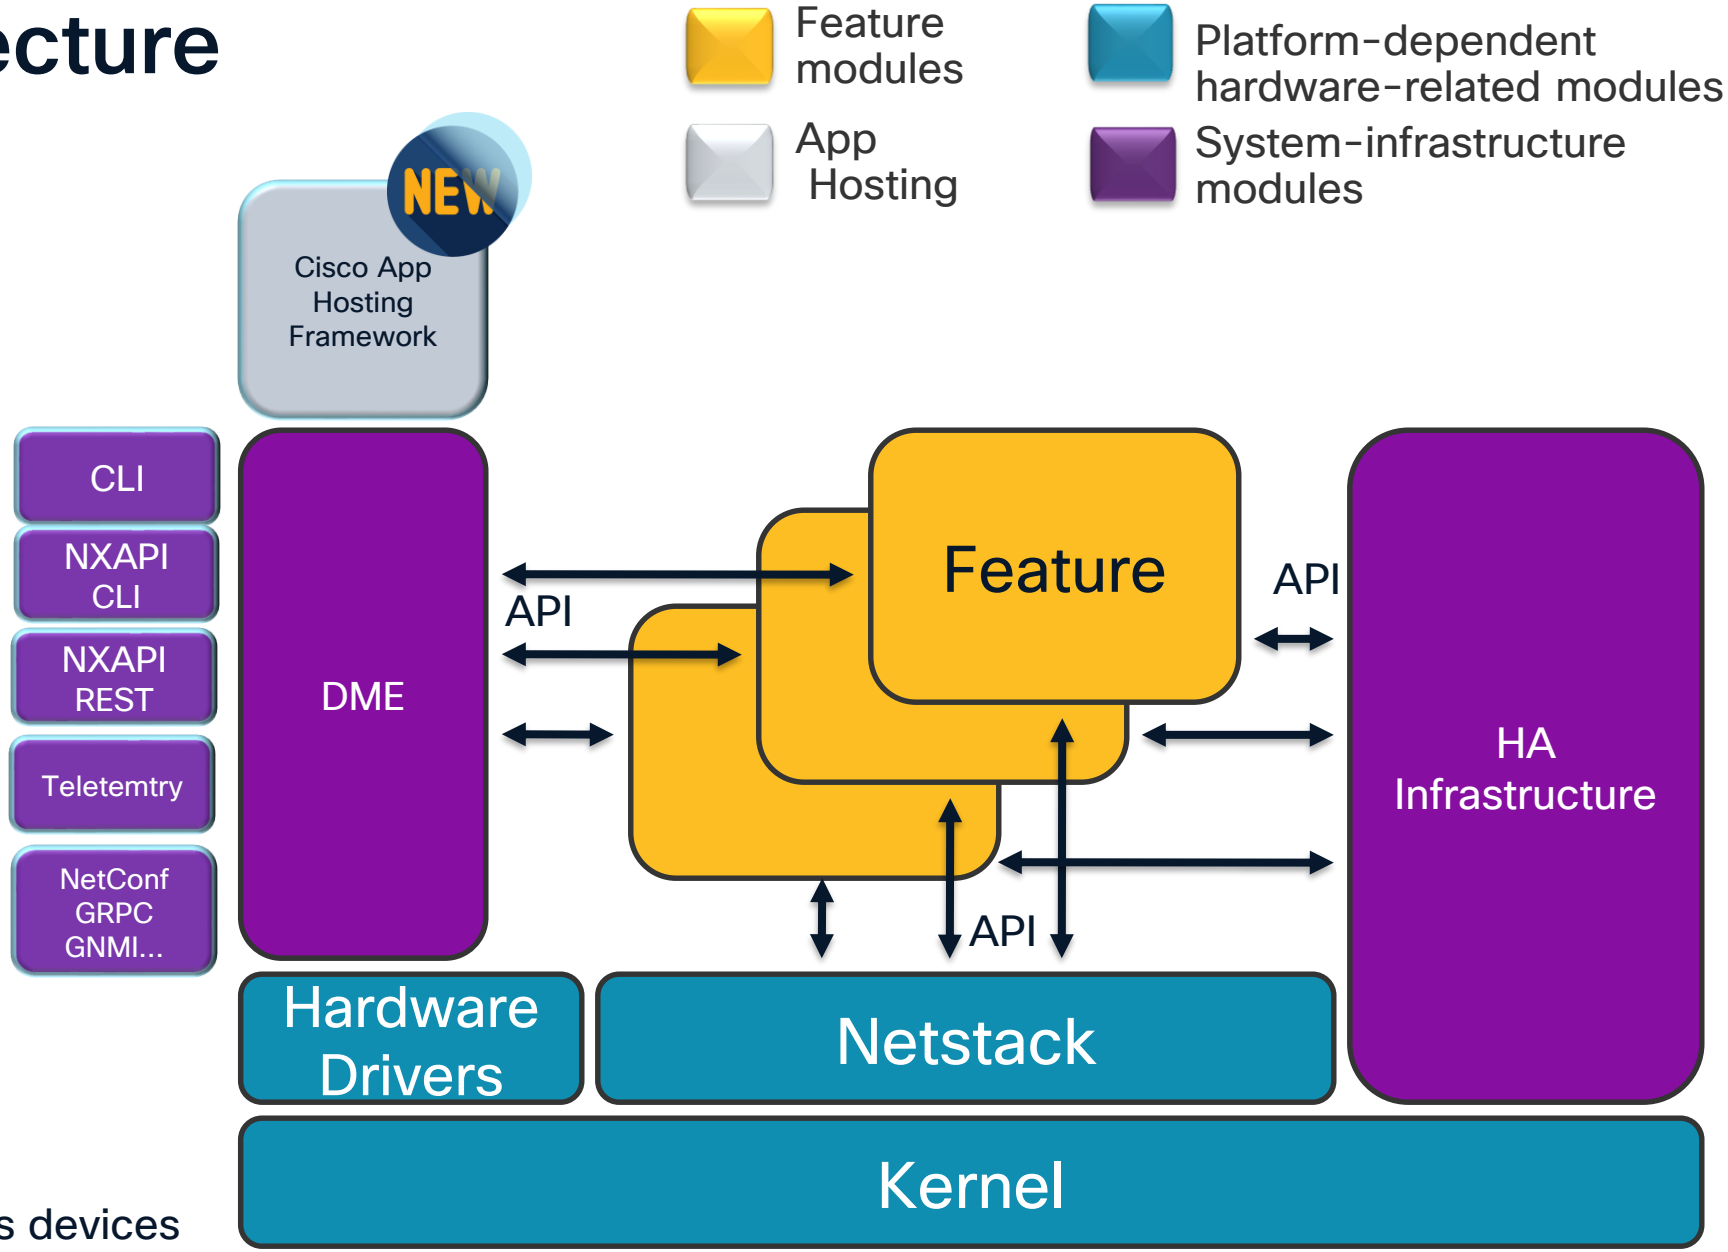

*SSO only available on dual-sup Nexus devices

ACI Architecture

Controller Integration

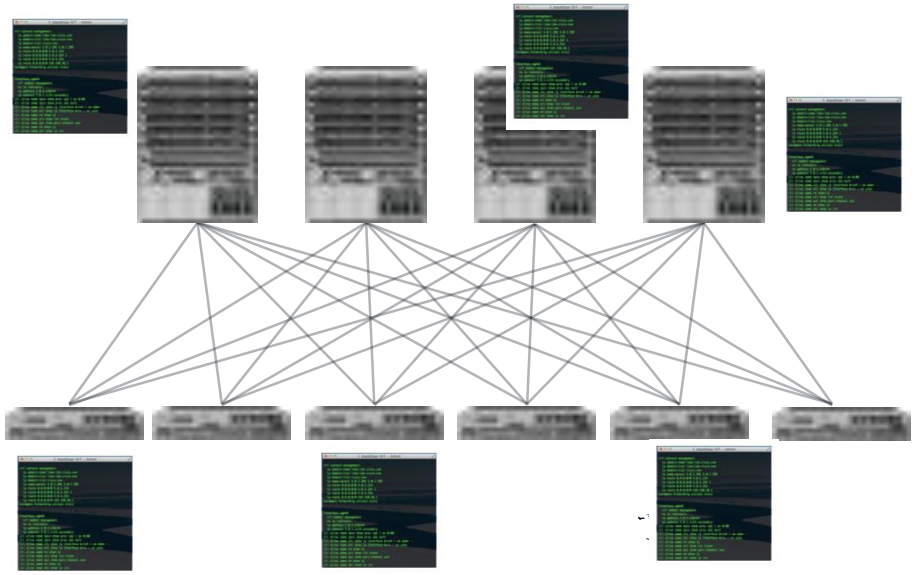

APIC becomes single point of management for the entire fabric … with a policy-based model

Interfaces, protocols, TCAM, etc … all represented in an object model, and ALL accessible through an XML/JSON API and CLI

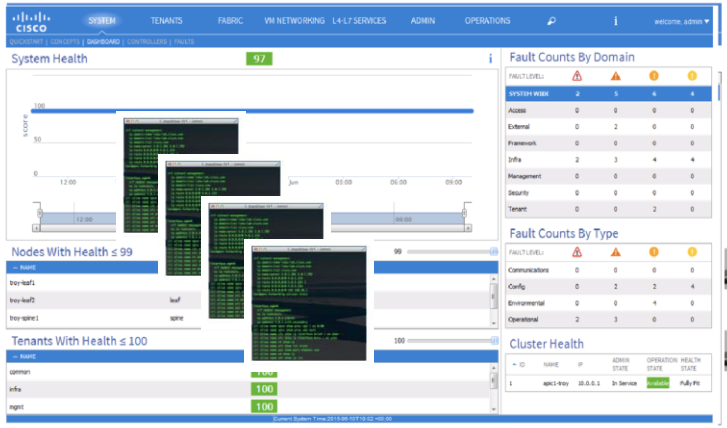

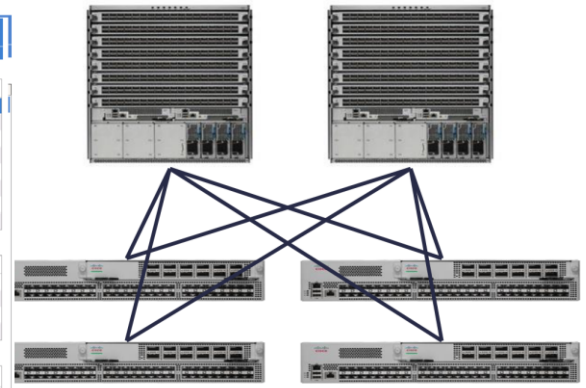

Now let's imagine a network switch … at the moment, largely configured on the CLI

Unplanned

Process Restart

NX-OS Stateful Process Restart

- NX-OS services Checkpoint their runtime state to the Persistent Storage Service (PSS) for Recovery in the event of a failure

If a fault occurs in a process...

- HA manager determines best Recovery Action (restart process, switchover to redundant supervisor)
- Process restarts with No impact on data plane
- State is recovered, operation resumes
- Total Recovery Time ~10ms

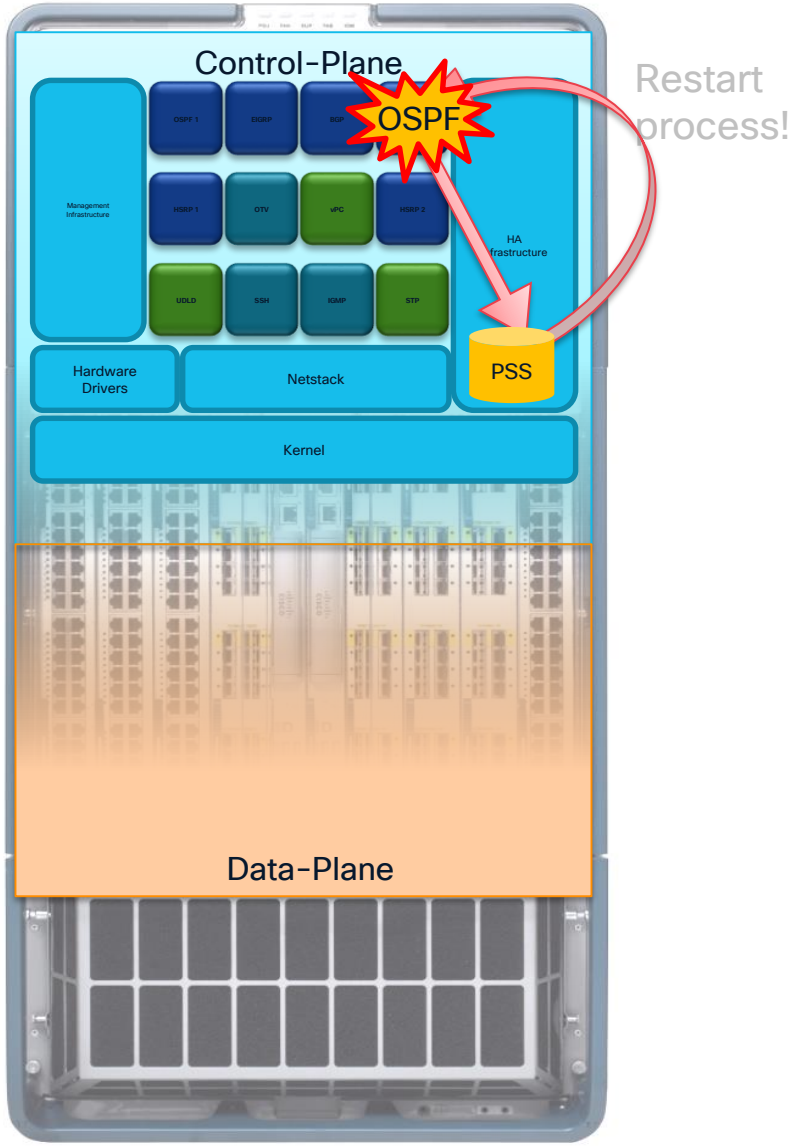

NX-OS Stateful Process Restart

- NX-OS services Checkpoint their runtime state to the Persistent Storage Service (PSS) for Recovery in the event of a failure

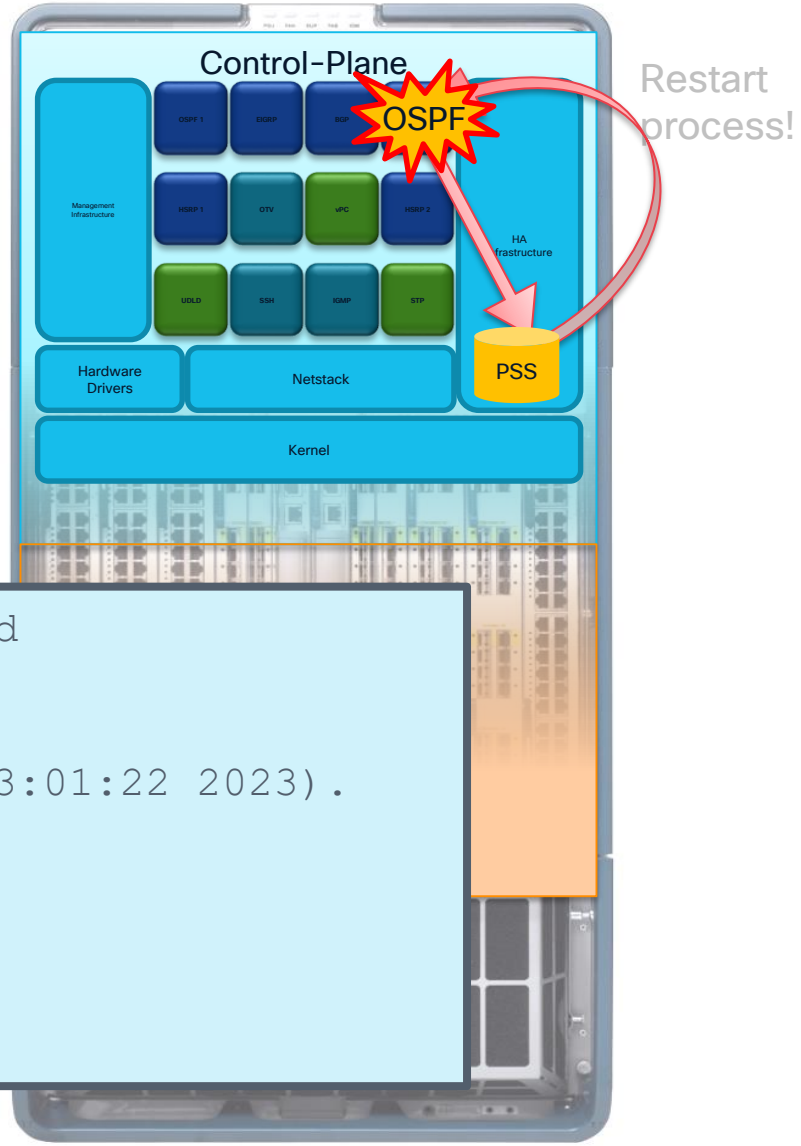

If a f

- HA (re sup
- Pro
- Sta
- To

```
tor-maibock2# show system internal sysmgr service name bfd
Service "bfd" ("bfd", 9):
UUID = 0x2C2, PID = 27798, SAP = 493
State: SRV_STATE_HANDSHAKED (entered at time Fri Apr 21 23:01:22 2023).
Restart count: 1
Time of last restart: Fri Apr 21 23:01:15 2023.
The service never crashed since the last reboot.
Tag = N/A
Plugin ID: 1
```

HA Policies

Service HA Policies	Action (single SUP)	Action (dual SUP)
Stateful restart	Stateful restart with running config	Stateful restart with running config
Stateless restart	Stateless restart with startup, default or no config	Stateless restart with startup, default or no config
No restart	Don't restart the service.	Don't restart the service.
Reset	Reload this SUP	Switchover
Switchover	N/A	Switchover

- **Each service defines its own HA policy**
- **Cascade HA policies.** If one fails (lightweight recovery), try another one. An example: fail to restart 3 times, try stateless or reset.

HA policies for a given service/application

Cascaded Policies Example

	Service/Application		

Single Sup (TOR)	ha_policy {
	restart=stateful;
	maxretry = 2;
	}
		ha_policy {
		restart=stateless;
		maxretry = 2;
		};	

Dual Sup (Modular)	ha_policy {
	restart=stateful;
	maxretry = 2;
	}
		ha_policy {
		restart=stateless;
		maxretry = 2;
		}
			ha_policy {
			
			restart=switchover;
			}

Unplanned

Non Stop Forwarding

NXOS Non-Stop Forwarding (GR)

Default Configuration

- The Nexus devices are “NSF Capable” by default for ALL the routing protocols in ALL NX-OS software releases.
- No additional configuration is required unless you need to modify the default NSF timers.

```
Switch# show running-config ospf all

❑  !Command: show running-config ospf all
!Time: Thu May 2 16:13:31 2024

version 9.3
feature ospf

<snip>

router ospf 1
  graceful-restart
  graceful-restart grace-period 60
  area 0.0.0.0 authentication message-digest
```

GR NXOS interoperability

- The Graceful Restart extensions used in NXOS are based on the IETF RFCs except for EIGRP, which is Cisco proprietary and can ONLY interoperate with Cisco NSF.
- This implies that for others like OSPF, BGP, the GR extension are compatible with versions of other devices that use the RFC based extensions.

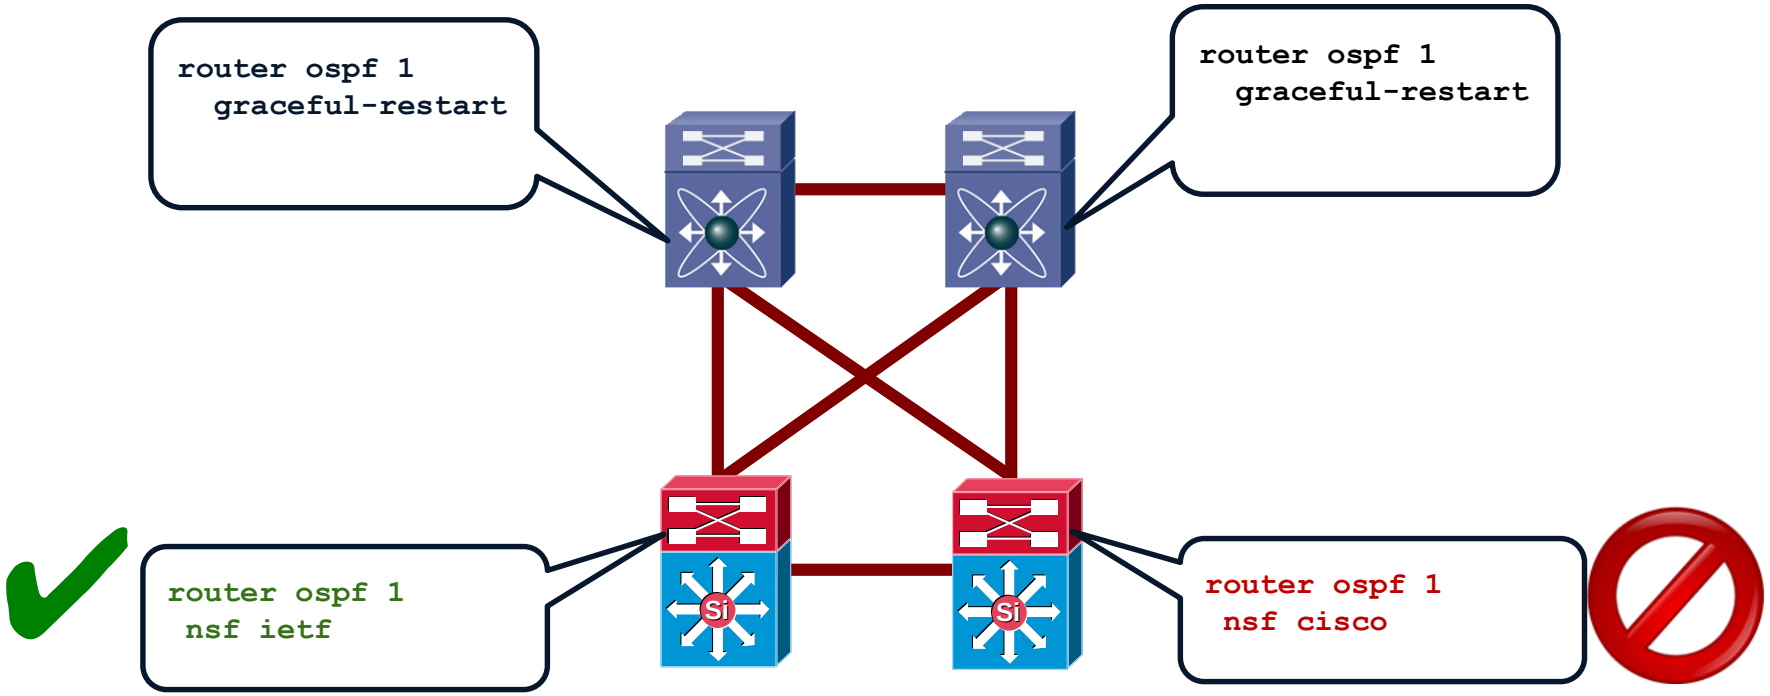

GR/NSF NXOS Operations

LSA exchanges

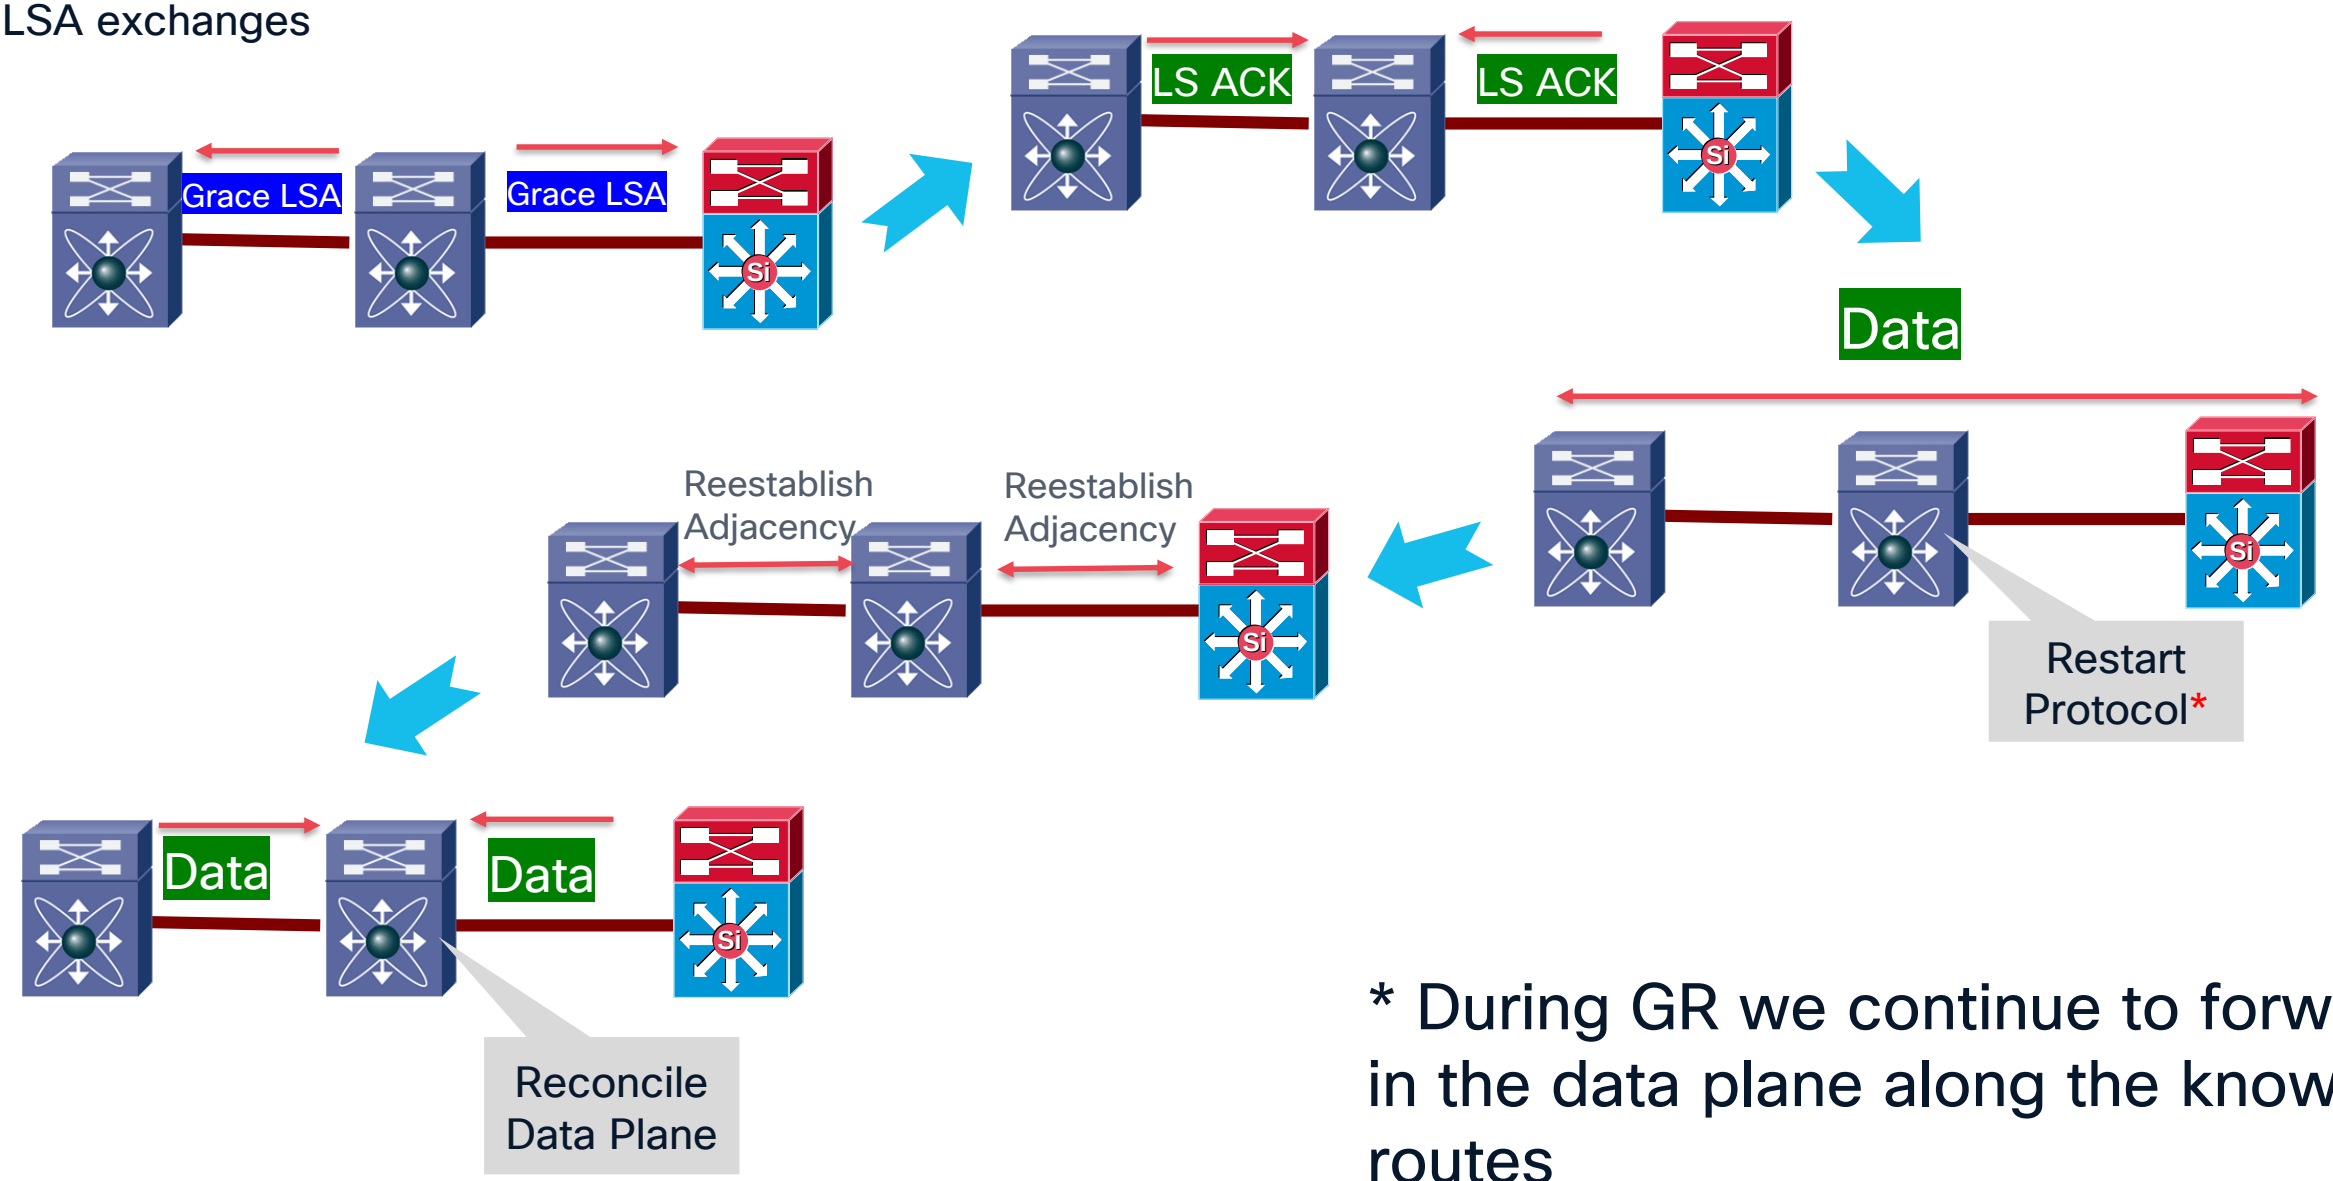

* During GR we continue to forward in the data plane along the known routes

NSF/GR Support Per Protocol

For Your Reference

Routing Protocol	NSF/GSR Definition
OSPFv2 Cisco	draft-nguyen-ospf-restart-06 (informational)
OSPFv2 IETF	RFC 3623
OSPFv3	RFC 5187
EIGRP	Cisco
BGP	RFC 4724
Intermediate System -Intermediate System	RFC 3847
Label Distribution Protocol	RFC 3478

Planned

Patching & ISSU

Validating SMU Impact

SMU Types	Impact
Non-Reload (Hot)	Process Restart
Reload (Cold)	Box Reload

Each SMU comes with a README file

```
##########################################################################
#
# Readme for SMU nxos64.CSCvz98895-n9k_ALL-1.0.0-10.3.1.167.lib32_64
#
# Copyright (c) 2021 by Cisco Systems, Inc.
# All rights reserved.
#
##########################################################################################
Name:           nxos64.CSCvz98895-n9k_ALL
Version:        10.3.1.167
Compressed size:89146
DDTS:           CSCvz98895
Description:    Kr2f-122: nxpython3 core while collecting show-tech-detail > icam script is failing
Type:           non-reload
MD5 :           a64fa50587d045828421860b840cc366
```

SMU Impact

Nexus SMU install commands

Adding SMU

Activating SMU

Committing SMU

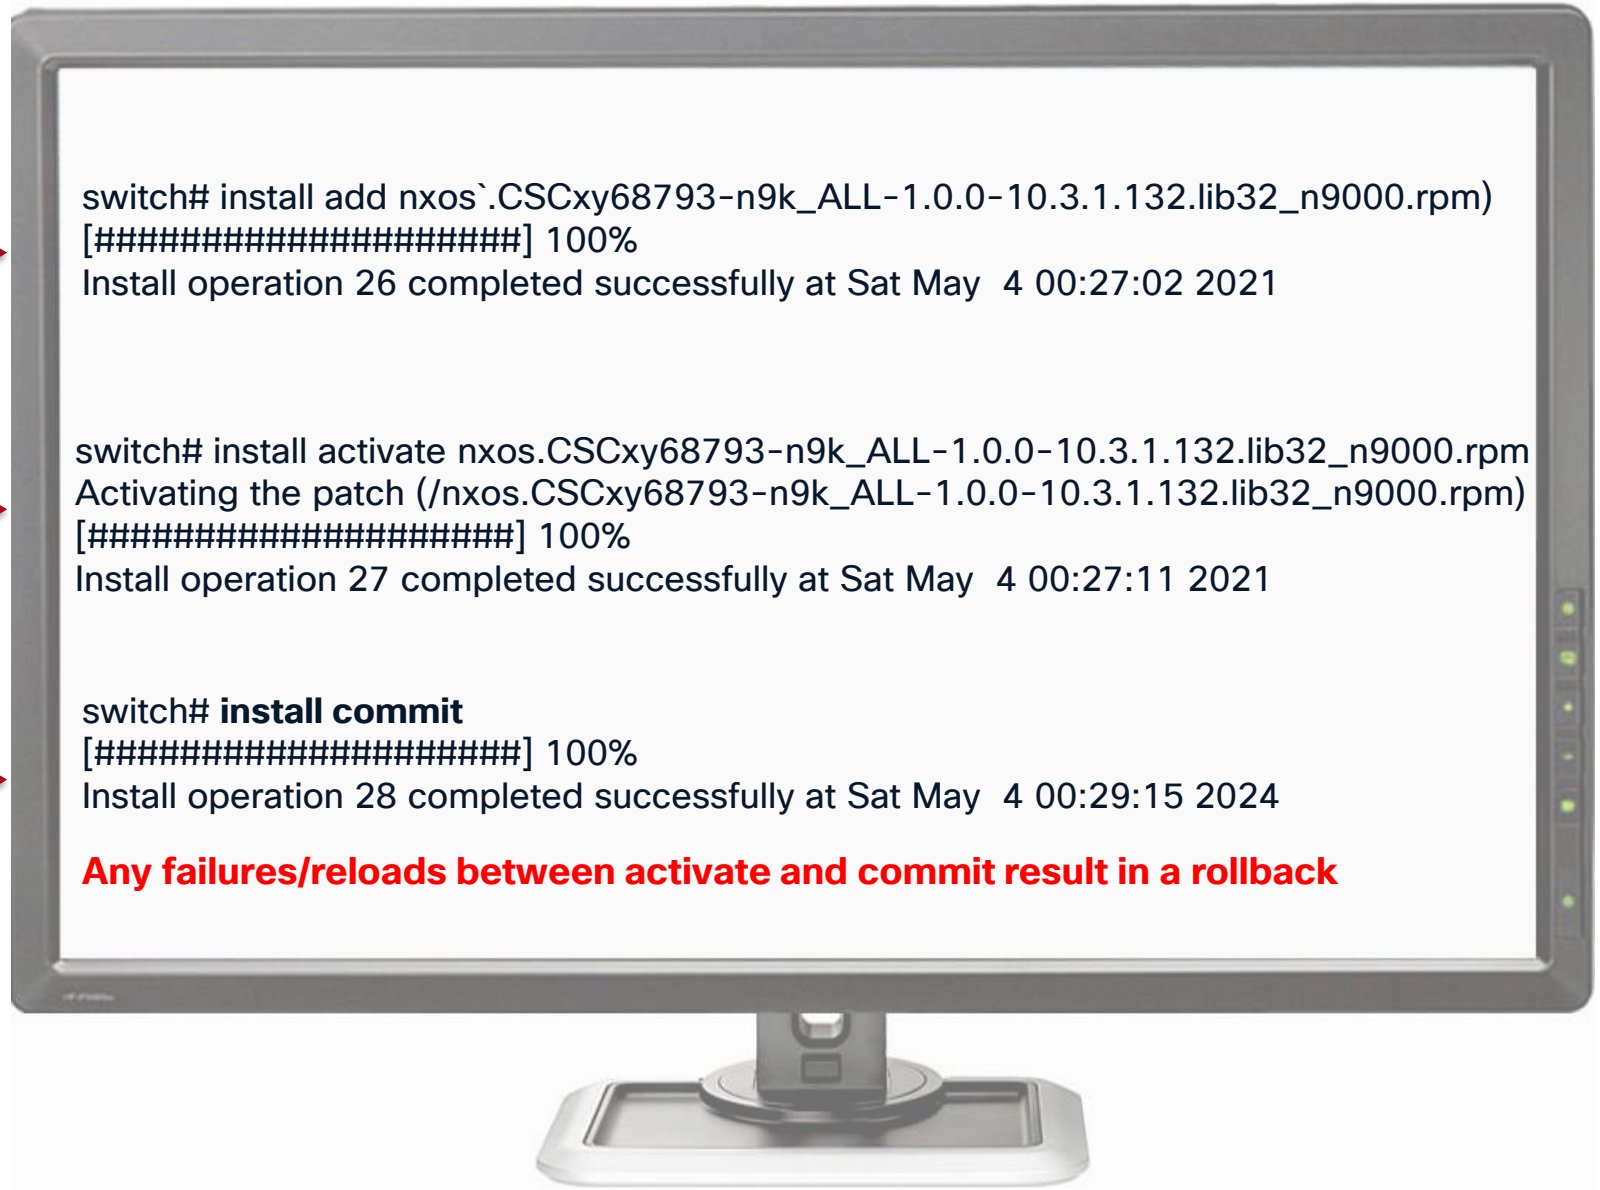

Non-reload(HOT) SMU update on Nexus 9000

- Services Checkpoints their state in PSS
- Installer asks SysMgr to stop the process
- Installer updates File System with files from SMU
- Installer asks SysMgr to restart the process
- No impact on data plane
- State is recovered, operation resumes
- Total Recovery Time ~10ms

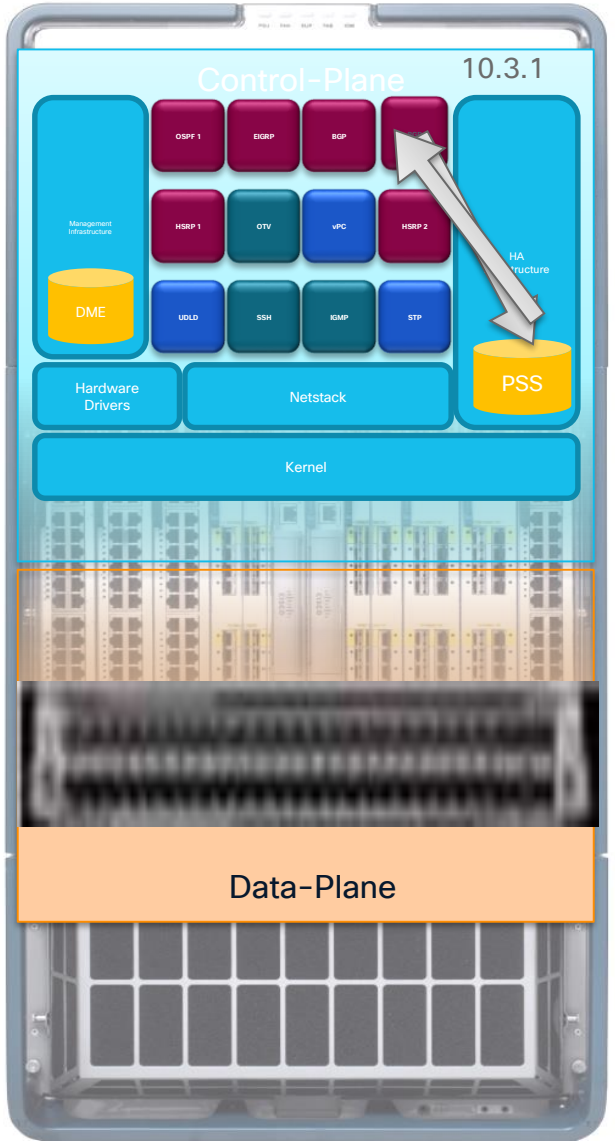

NXOS ISSU

Running Manual Switch Checks (Upgrades)

- First step is downloading the image to switch

#show install all impact nxos...

New Image Version

```
switch# show install all impact nxos bootflash:nxos64.10.3.1.F.bin
Installer will perform impact only check. Please wait.

Verifying image bootflash:/nxos64.10.3.1.F.bin for boot variable "nxos".
[####################] 100% -- SUCCESS

Verifying image type.
[####################] 100% -- SUCCESS

Preparing "nxos" version info using image bootflash:/nxos64.10.3.1.F.bin.
[####################] 100% -- SUCCESS

Preparing "bios" version info using image bootflash:/nxos64.10.3.1.F.bin.
[####################] 100% -- SUCCESS

Notifying services about system upgrade.
[####################] 100% -- SUCCESS
```

Running Manual Switch Checks (Upgrades)

#show install all impact nxos...

New Image Version

The impact

Reason

```
switch# show install all impact nxos bootflash:nxos64.10.3.1.F.bin
Installer will perform impact only check. Please wait.

Verifying image bootflash:/nxos64.10.3.1.F
[####################] 100% -- SUCCESS

Verif
[####

Prepa
[####

Prepa
[####

Notif
[####
```

```
Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes      disruptive         reset  default upgrade is not hitless
    27       yes      disruptive          none  default upgrade is not hitless


Images will be upgraded according to following table:
Module       Image                  Running-Version(pri:alt)           New-Version  Upg-Required
------  ----------  ------------------------------------------  -------------------  ------------
     1       lcn9k                                   10.2(2)              10.2(2)            no
    27        nxos                                   10.2(2)              10.2(2)            no
    27        bios     v05.45(07/05/2021):v05.39(08/30/2019)    v05.45(07/05/2021)            no
```

Hitless ISSU on the Nexus TOR N9K

- Single supervisor
- Hitless Upgrade
  - Control plane is inactive during reload while Data plane is forwarding

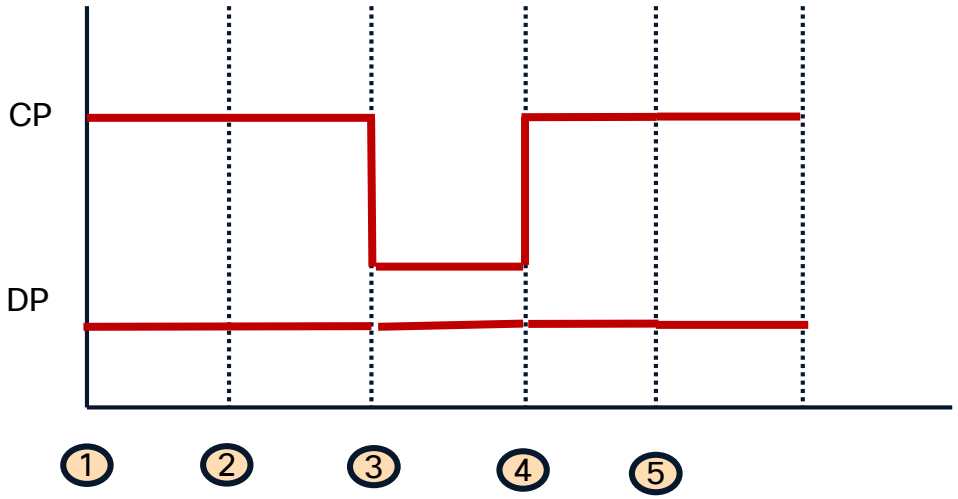

1. Pre Upgrade Check
   - Config Locked
   - Stable Network
2. Save State
3. kexec to new Image
4. Restore control plane from saved State
5. Reconcile with Data Plane

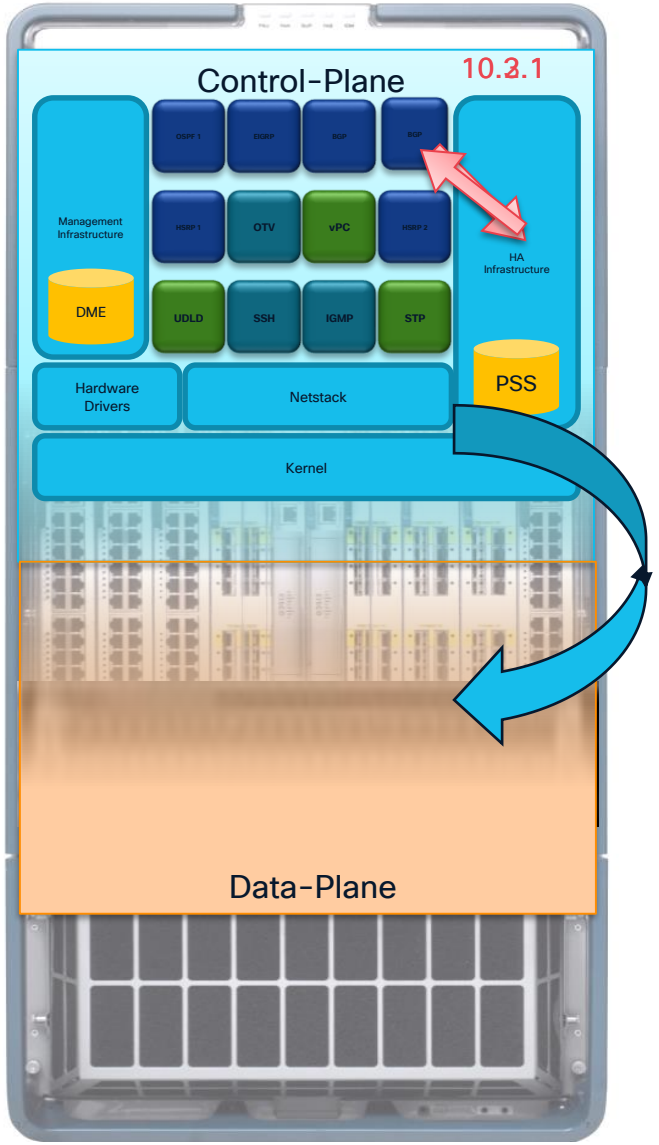

Enhanced ISSU

Enhanced ISSU on the Nexus TOR N9K

- Dual Container
- Control plane is only down for ~3-6 sec during switchover

If Kernel uprev is detected, we auto-revert to normal ISSU

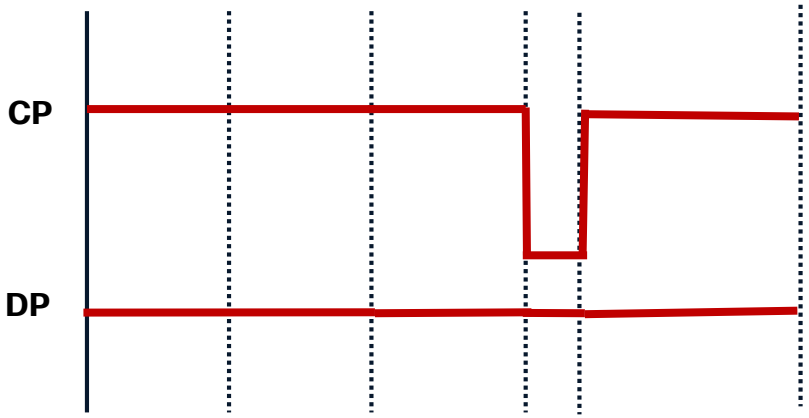

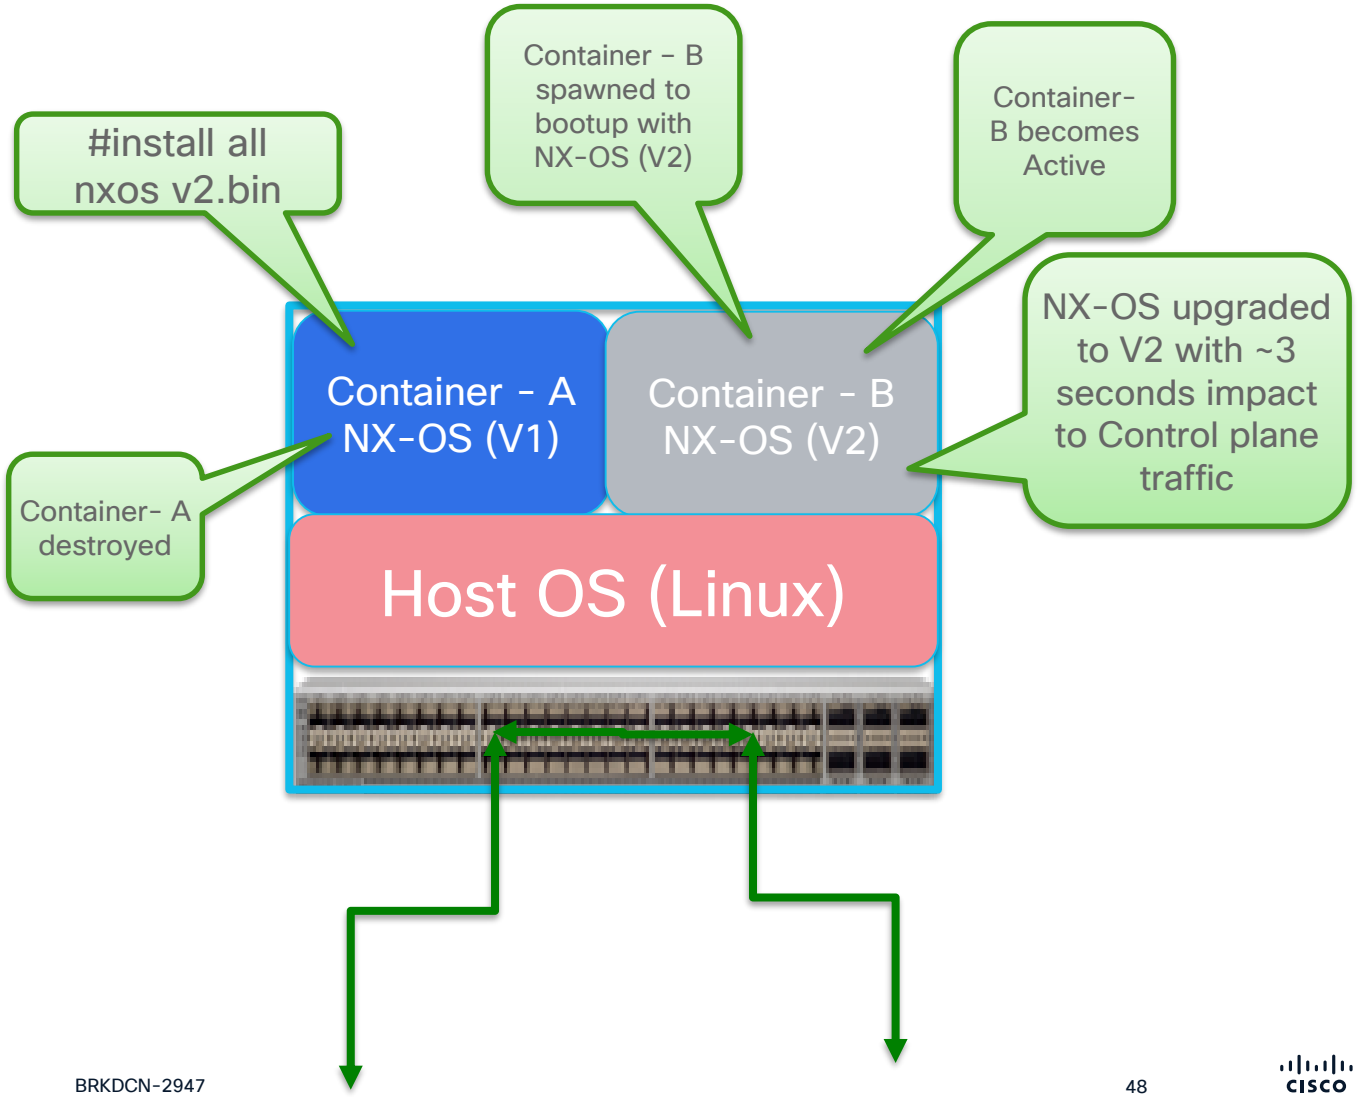

Show Module output

LXC enabled

```
tor3-hw-7# show module
Mod Ports             Module-Type                       Model         Status
--- ----- --------------------------------------------- ---------------------- --------
1   54   48x10M/100M/1G + 4x10/25G + 2x40/100G Eth Mod    N9K-C9348GC-FXP      ok
27  0    Virtual Supervisor Module                  N9K-C9348GC-FXP      active *

Mod  Sw                  Hw    Slot
---  ----------------------- ------ ----
1    10.3(2)IOV9(0.198)      0.1080 NA
27   10.3(2)IOV9(0.198)      0.1080 VSUP


Mod  MAC-Address(es)                    Serial-Num
---  -------------------------------------  ----------
1    2c-33-11-ec-5f-16 to 2c-33-11-ec-5f-5d  FDO21100DTW
27   2c-33-11-ec-5f-16 to 2c-33-11-ec-5f-5d  FDO21100DTW

Mod  Online Diag Status
---  ------------------
1    Pass
27   Pass
```

ISSU Software Upgrades via Controllers NDFC/APIC

Controller Upgrade Types

NDFC & APIC

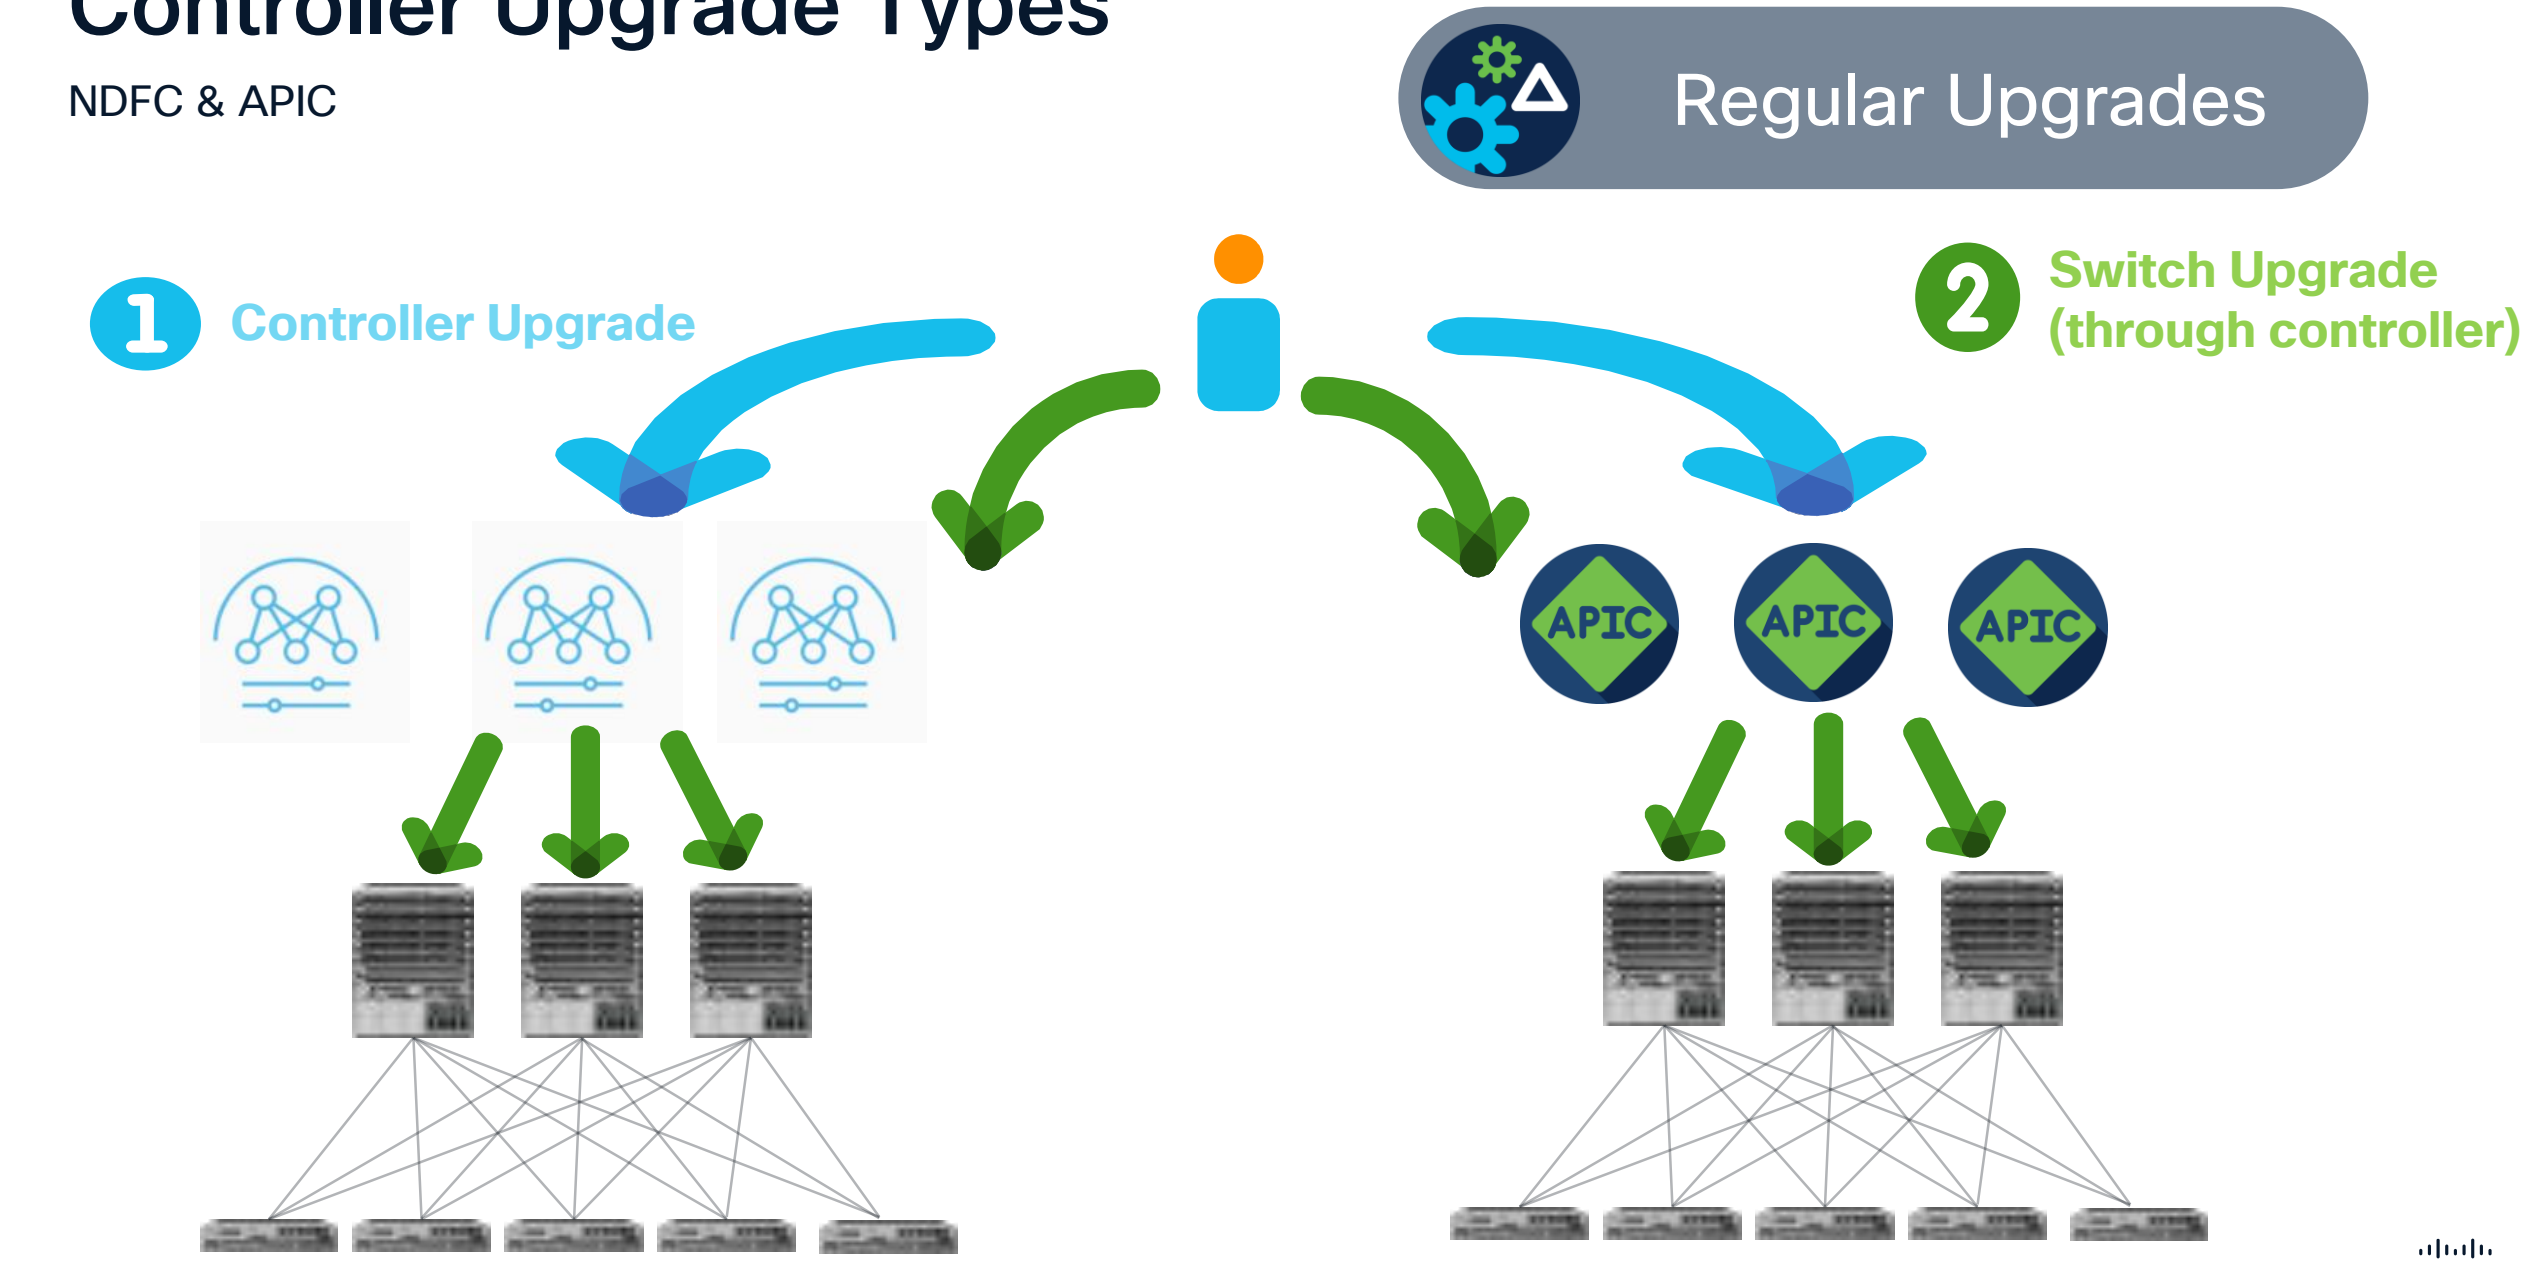

Controller Upgrade Types

NDFC & APIC

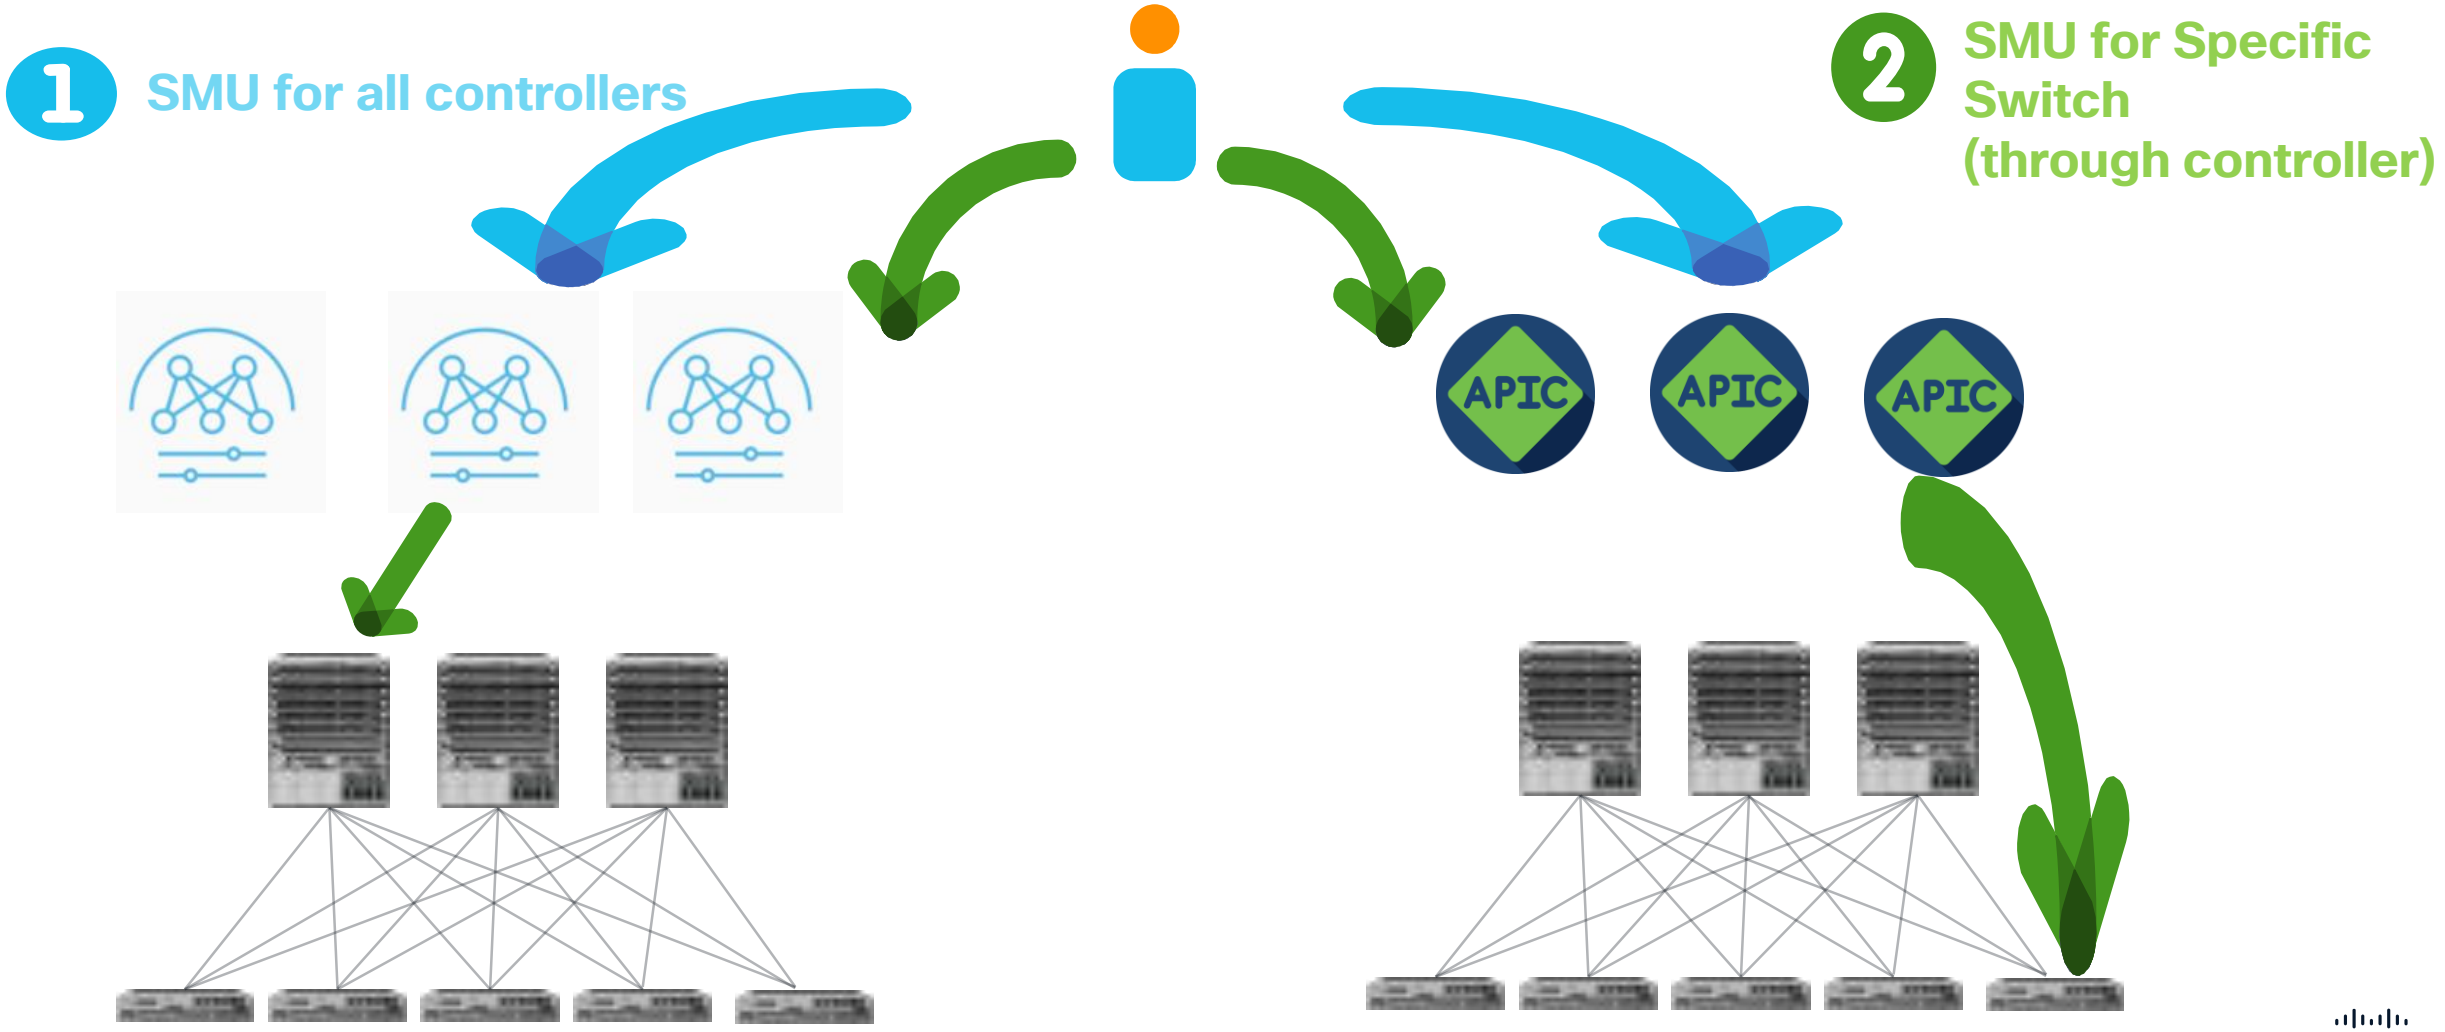

Controller Upgrade Types

NDFC & APIC

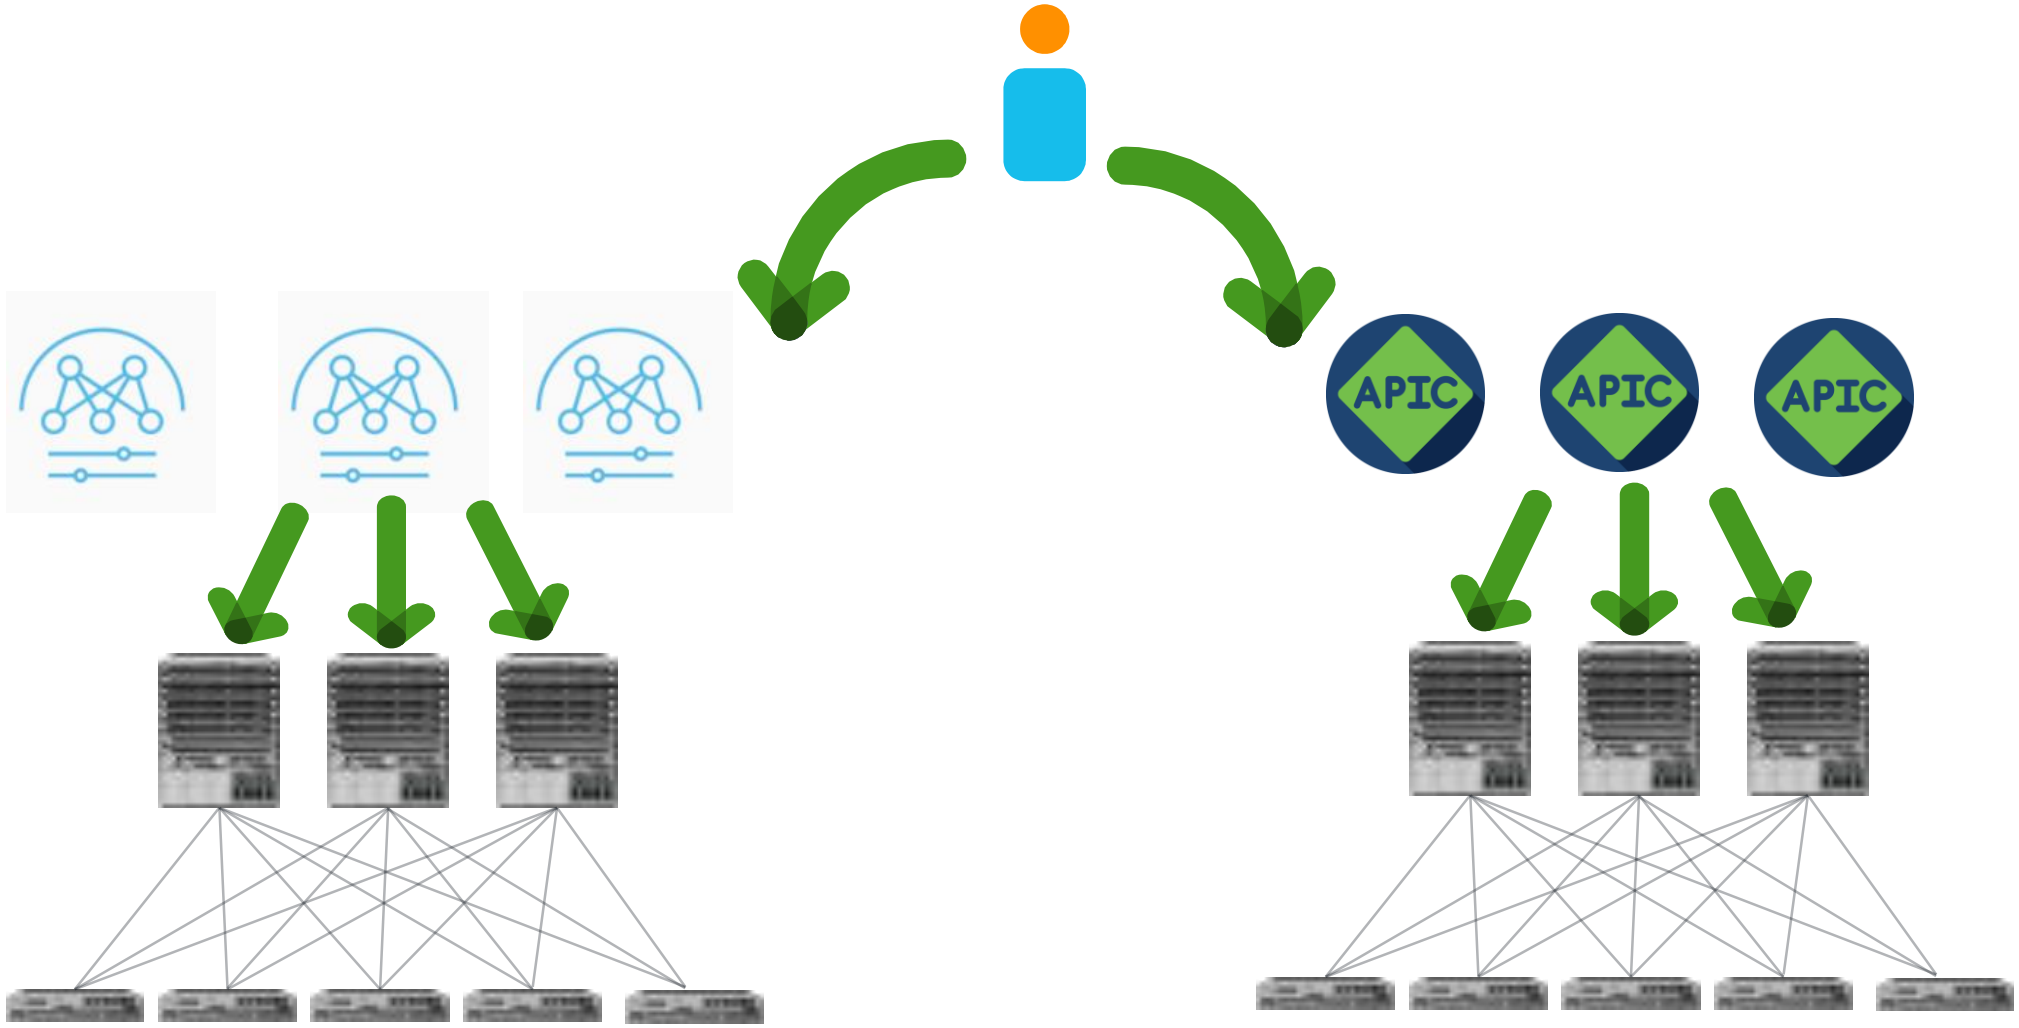

Simplifying Upgrade Process

Non-Disruptive NXOS upgrade

Defaulting all ND ISSU to enhanced ISSU on GX2A and GX2B platforms

10.3.3 GX, FX3

Reducing Number of reloads required during upgrade

Combining NXOS and EPLD upgrade avoids additional reload in case of EPLD – 9.3(8) and 10.1(2)

Apply SMU with NXOS and EPLD upgrades - 10.2(1)

Prevent downtime for Kernel Patch

Providing the ability to perform Kernel patch without reload 10.2(2)

Best Practices for Nexus Upgrades

Controller or Switch

Check Upgrade Path
https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/ISSUmatrix/index.html

Download Upgrade/Patch Images prior to maintenance window

Do not make Configuration Changes during upgrades

Run Pre-Upgrade Validation Checks

Compare state Before and After Upgrades

Using NAE

3rd party Apps

Network Level HA

Legacy Data Center Design and Tradeoffs

- Well- known placement of STP root and secondary root switches,
- Well-known placement of (HSRP) primary and secondary devices.
- Forwarding and blocking links from the access layer to the aggregation layer
- Various hardening features to guarantee deterministic spanning-tree behavior upon link failures.

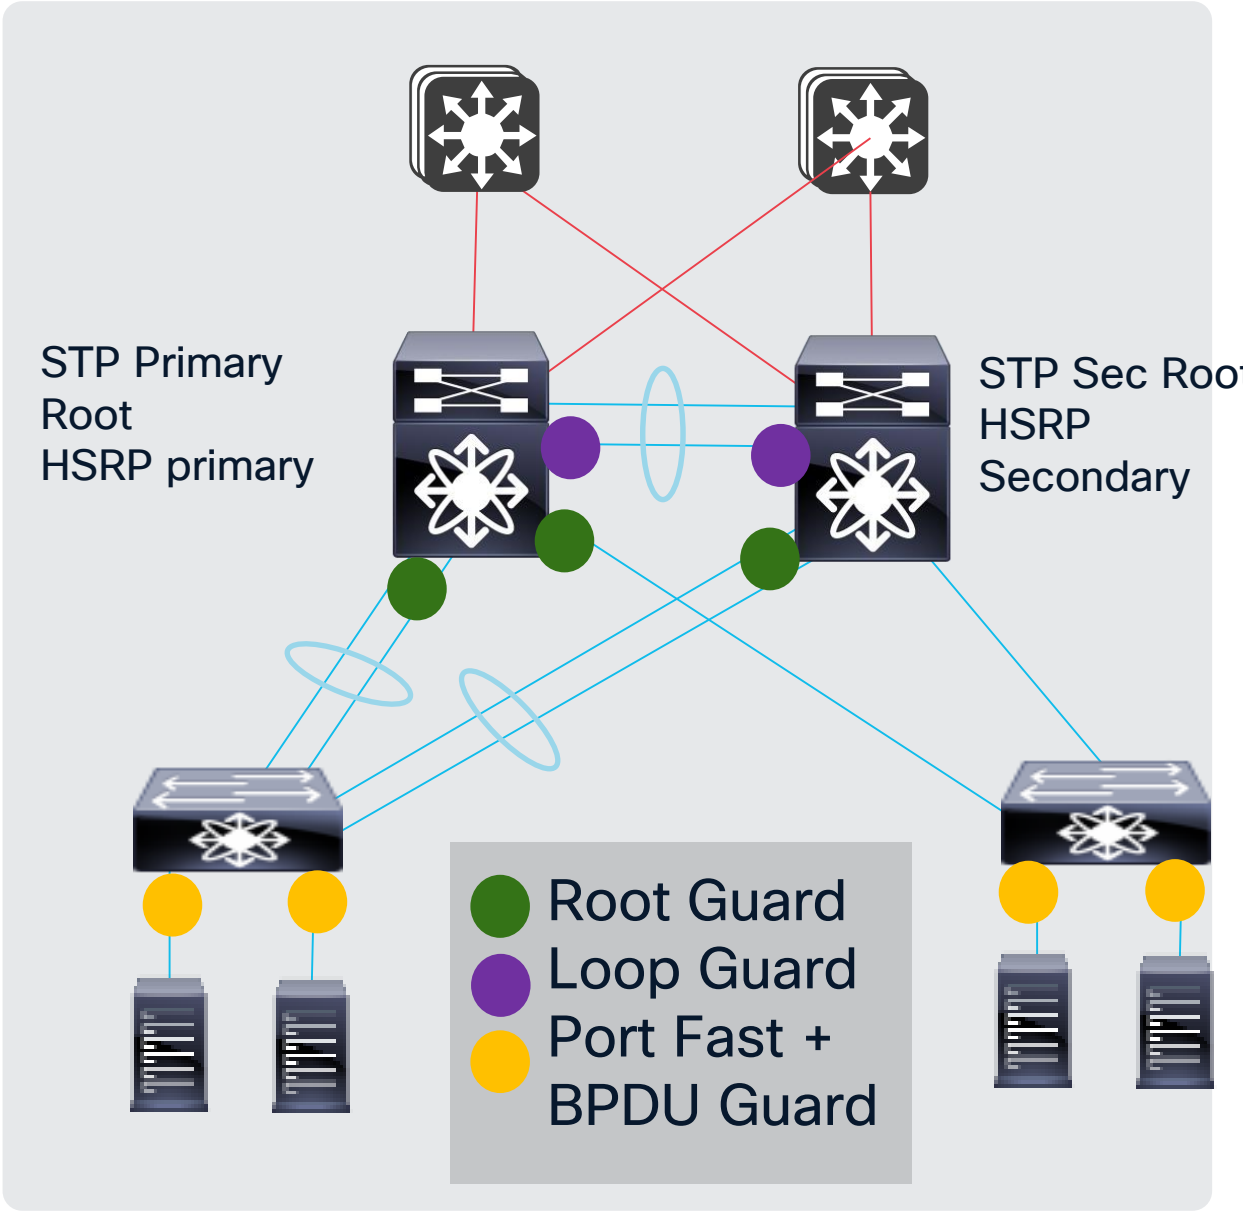

The Paradigm

- A Leaf and Spine Topology
- A Scale Out Architecture
  - More Leaf = More Ports
  - More Spine = More Bandwidth

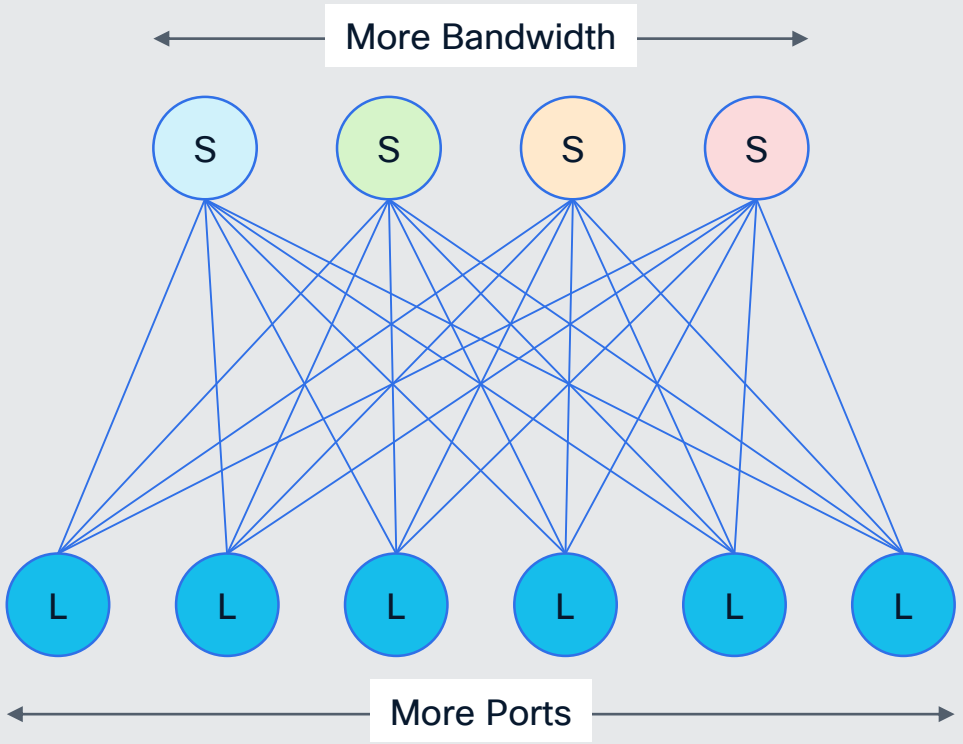

The Paradigm

- N+1 Redundancy
- Redundancy increases by Building out the Topology
- On Spine failure
  - 4 Spine = 25% impact
  - 8 Spine = 12.5% impact

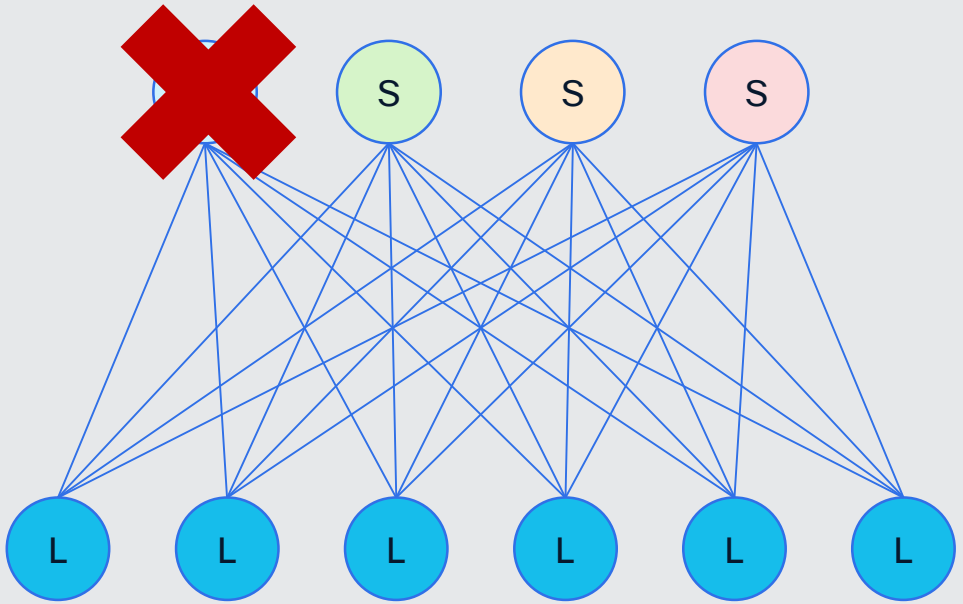

The Journey to Build Better and Further

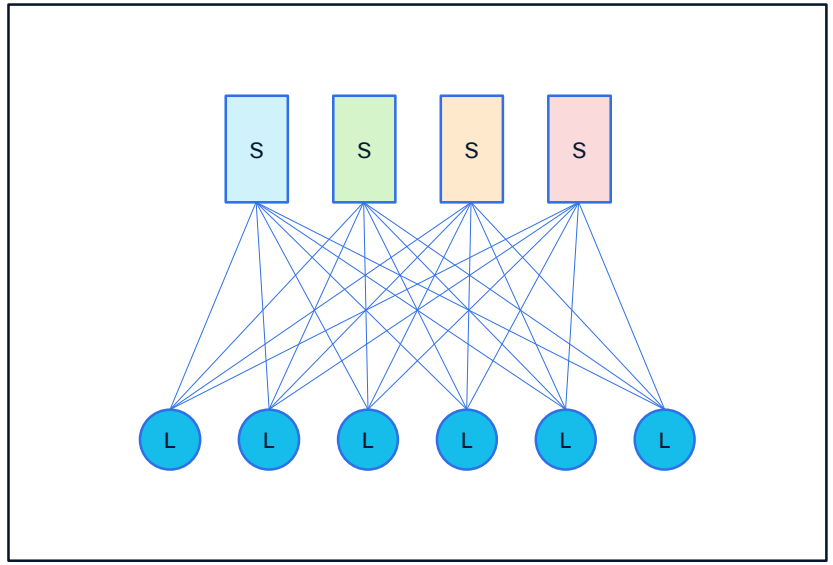

2 Tier Leaf Spine
(5 Stages)

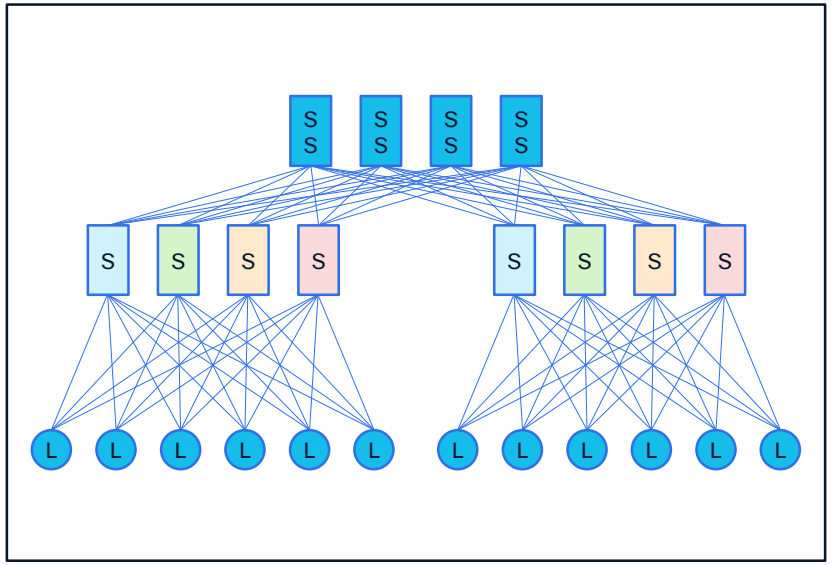

3 Tier Leaf-Spine-SuperSpine
(11 Stages)

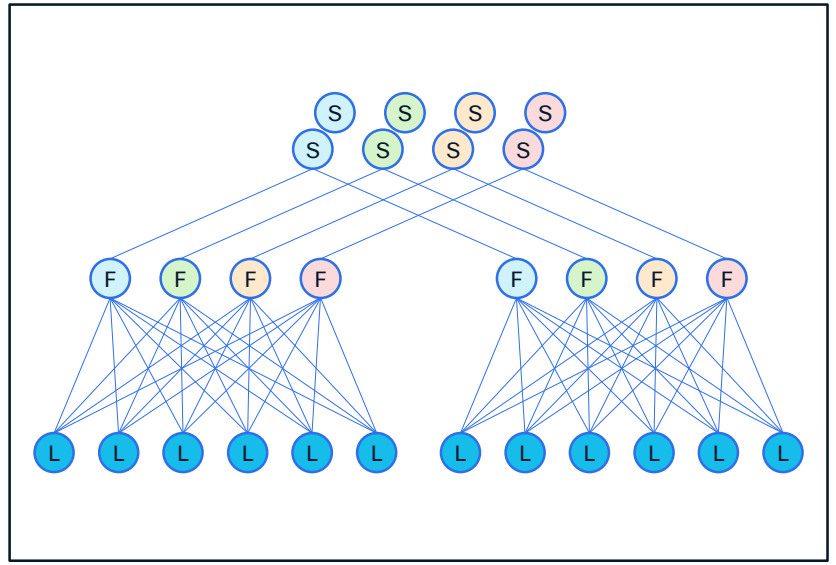

3 Tier Leaf-Fabric-Spine
(5 Stages)

The Journey to Build Better and Further

2 Tier Leaf Spine
(5 Stages)

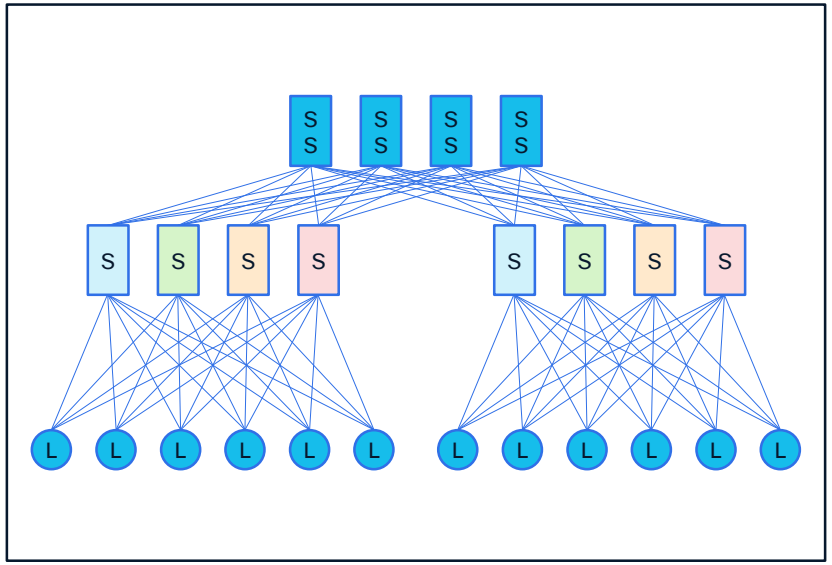

3 Tier Leaf-Spine-SuperSpine
(11 Stages)

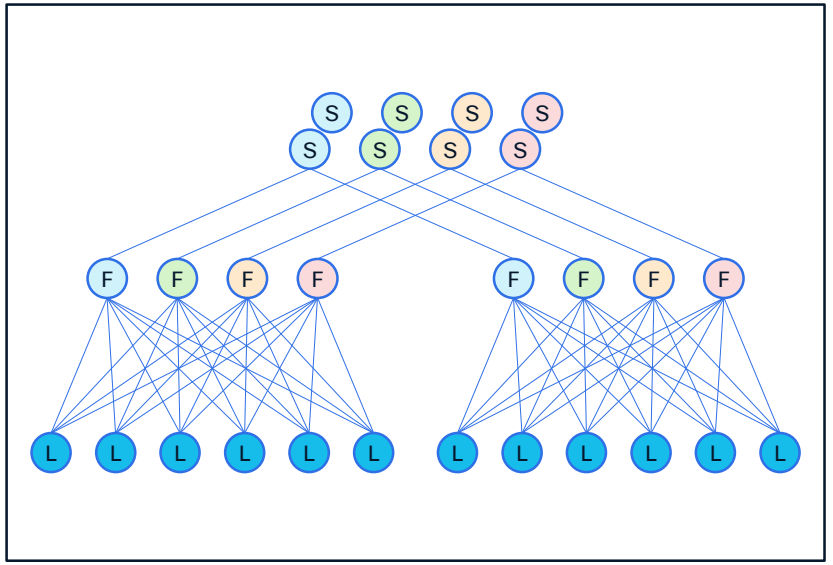

3 Tier Leaf-Fabric-Spine
(5 Stages)

Let's Move Forward

The Journey to Build Better and Further

2 Tier Leaf Spine
(5 Stages)

3 Tier Leaf-Spine-SuperSpine
(11 Stages)

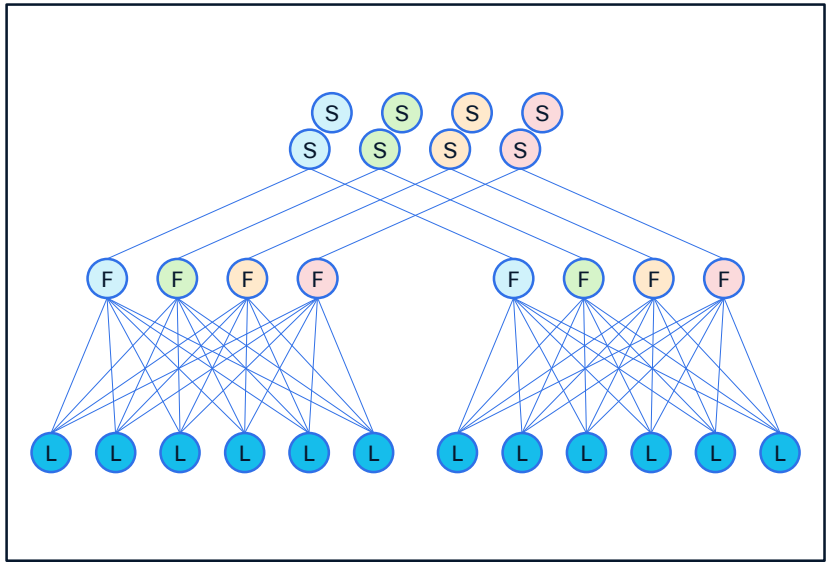

3 Tier Leaf-Fabric-Spine
(5 Stages)

Let's Move Forward

What we learned from the Cloud Titans

Building Scalable and Reliable DataCenter Networks

#1

Simplicity is Key

Simple Design Principals

#2

Scale as you Go

Scale is Never Finite

#3

Fail but Fail Fast

Reduce Brown-Out Exposure

#4

Redundant and Repeatable

Risk is Never an Option

Don’t Build Fabric for Maximum Leaf

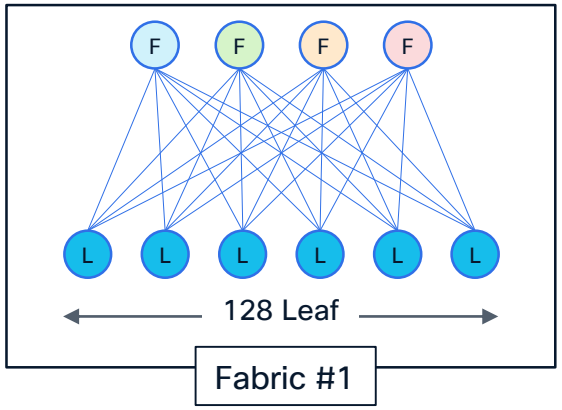

Repeat for Host Port Scale (Scale Out)

- Increasing Fabrics at need
  - of Host Port
  - of Oversubscription between Tier #2 and #3
- Result Defines Tier #2 to Tier #3 Uplinks
  - and respectively Tier #3 Requirements

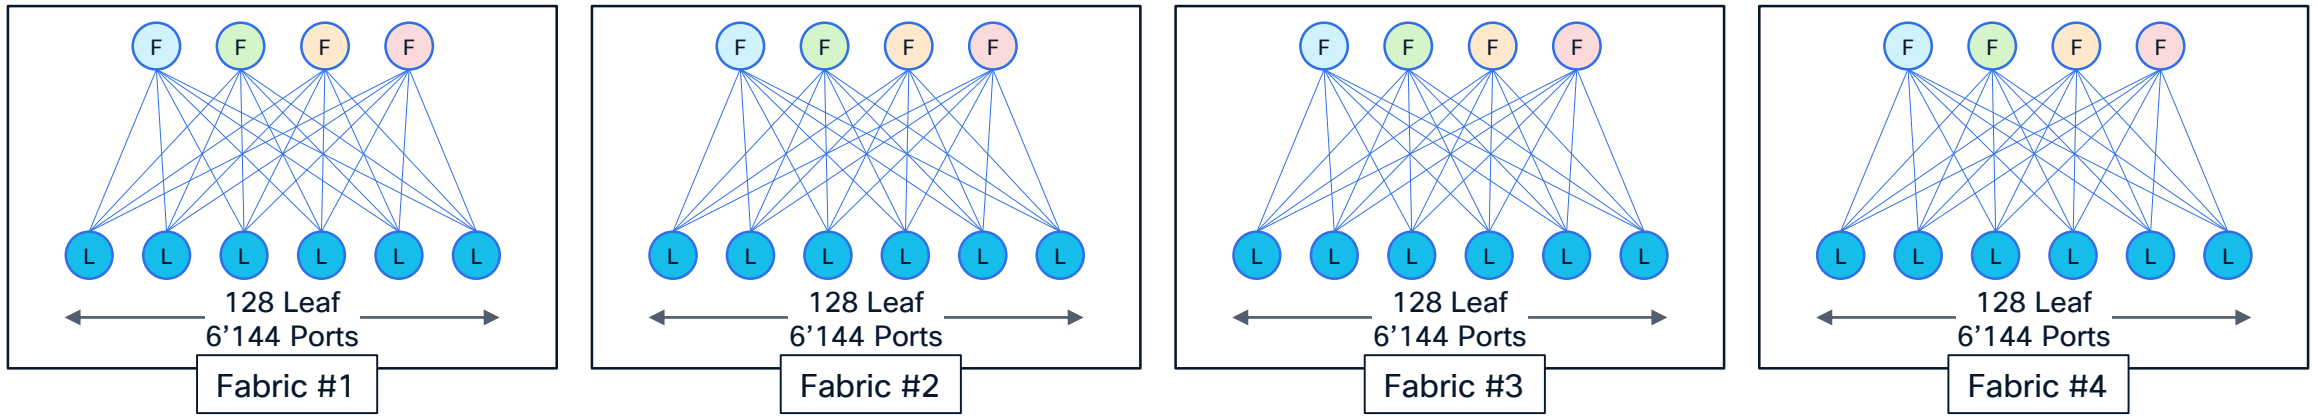

Designing Tier #3 (Single Link)

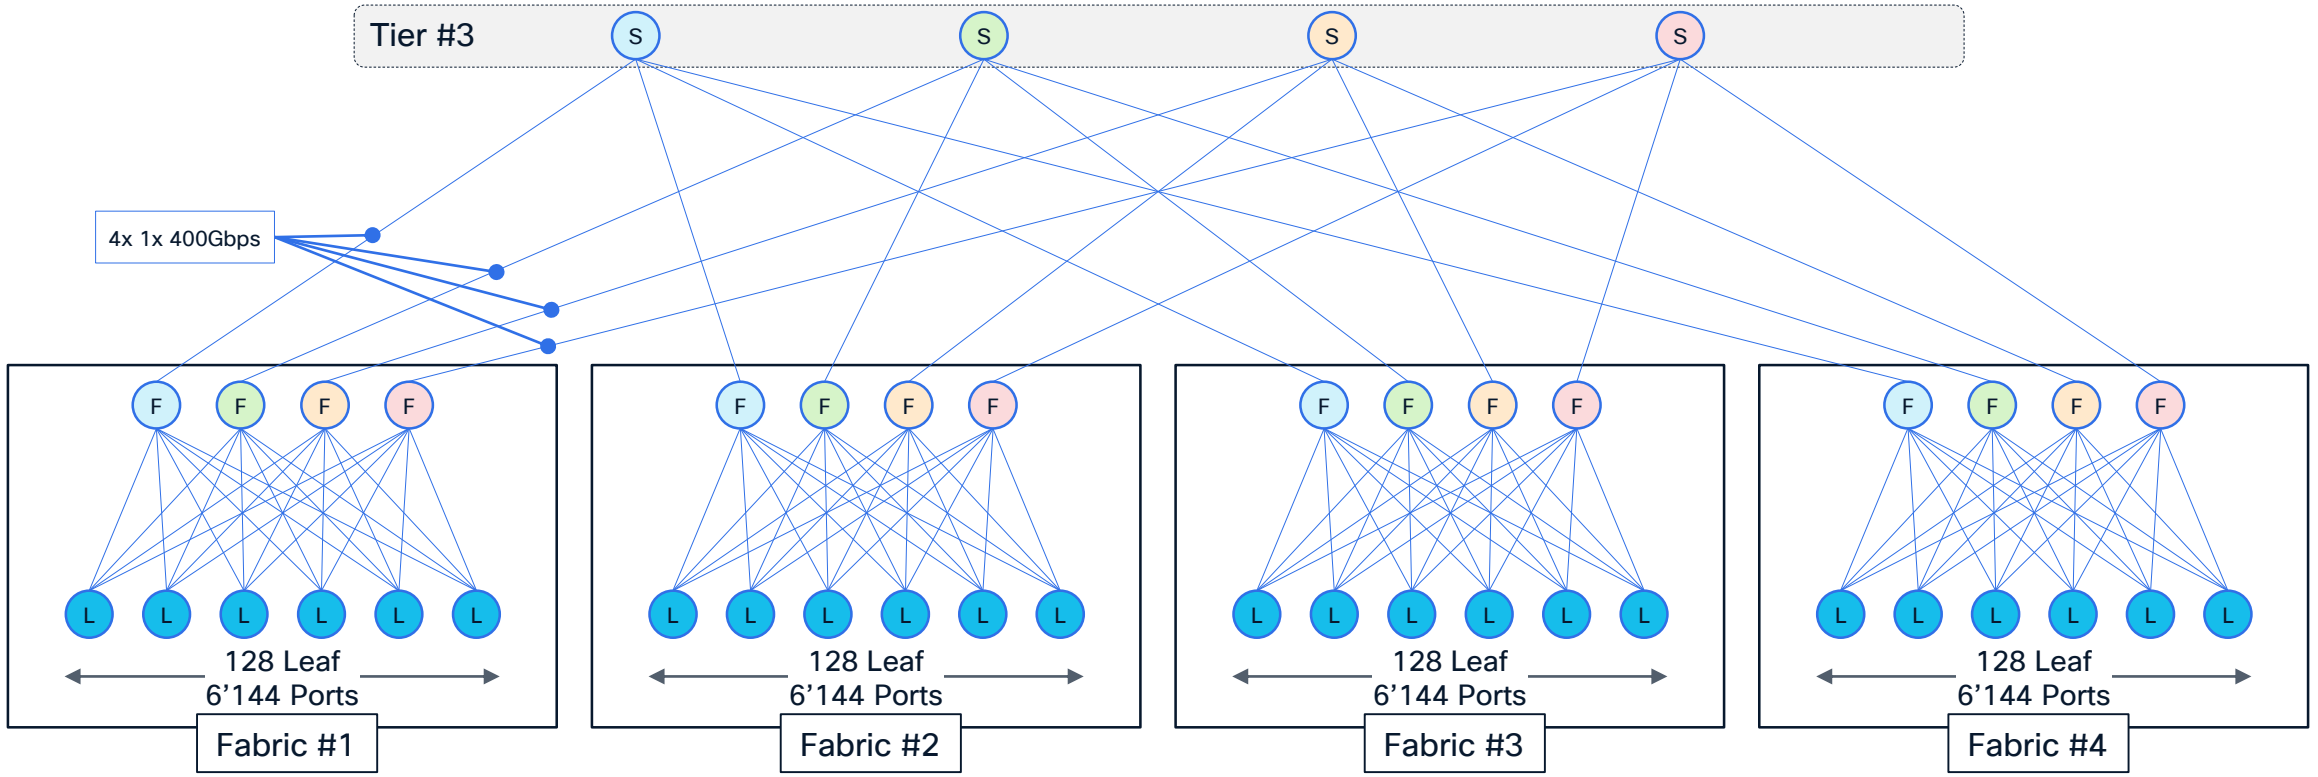

Increasing the Tier #3 Planes

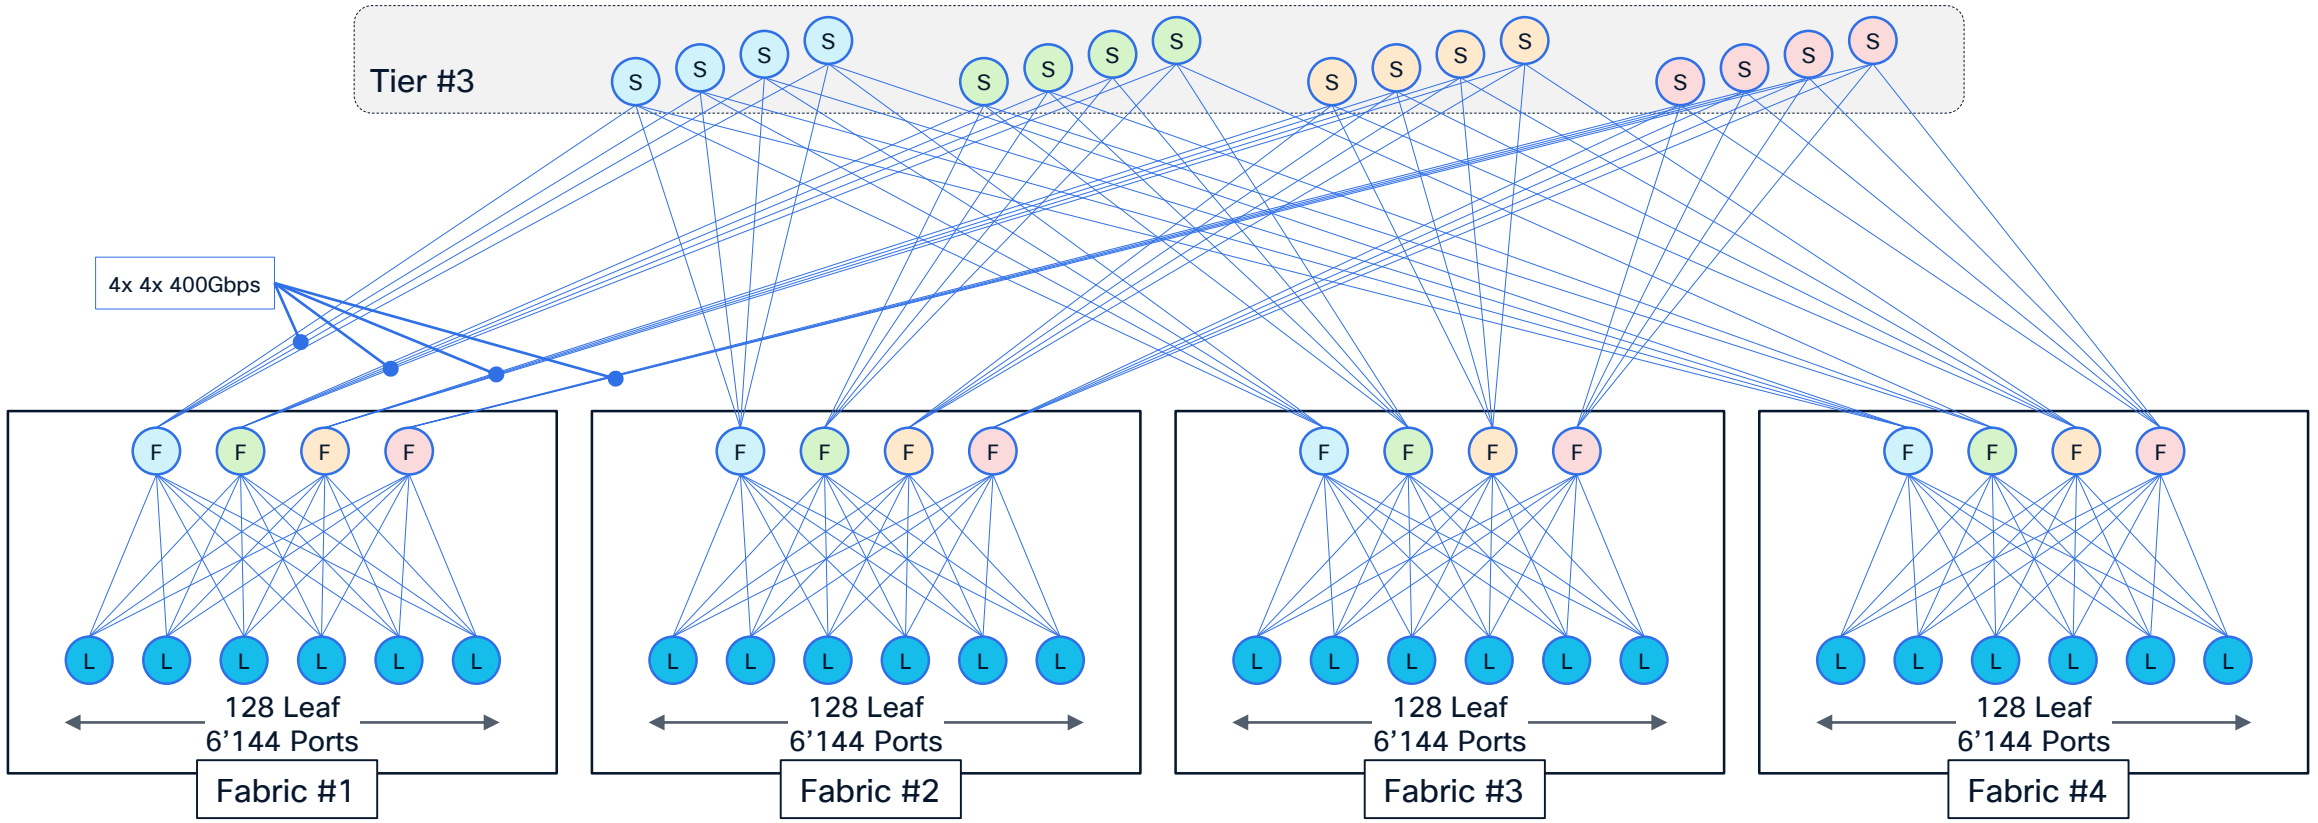

Key Takeaways

Start small and Grow to Large as needed

Bigger is not Always Better; Using Fixed-Form factor Switches is a Modern Practice

More switches != Higher Cost

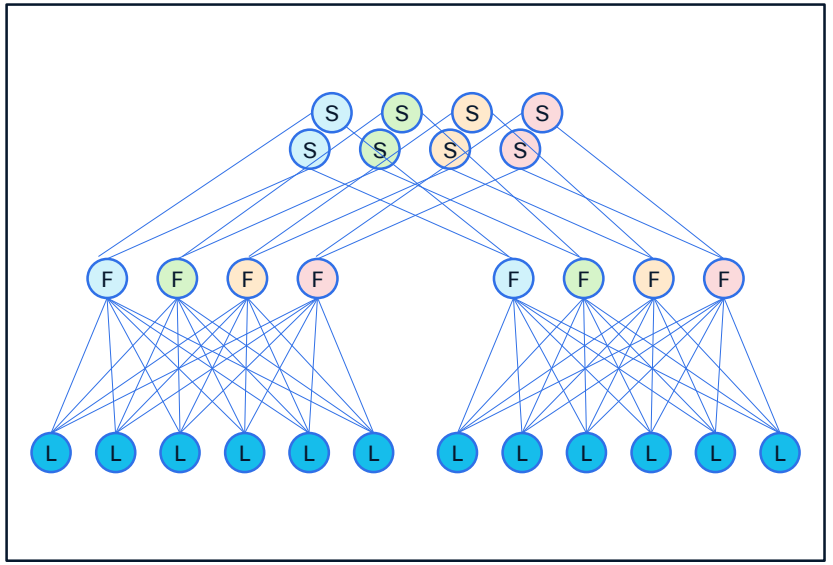

3 Tier Leaf-Fabric-Spine (5 Stages)

Recommendations for Route Server Placement

Inline vs Separate Route Servers

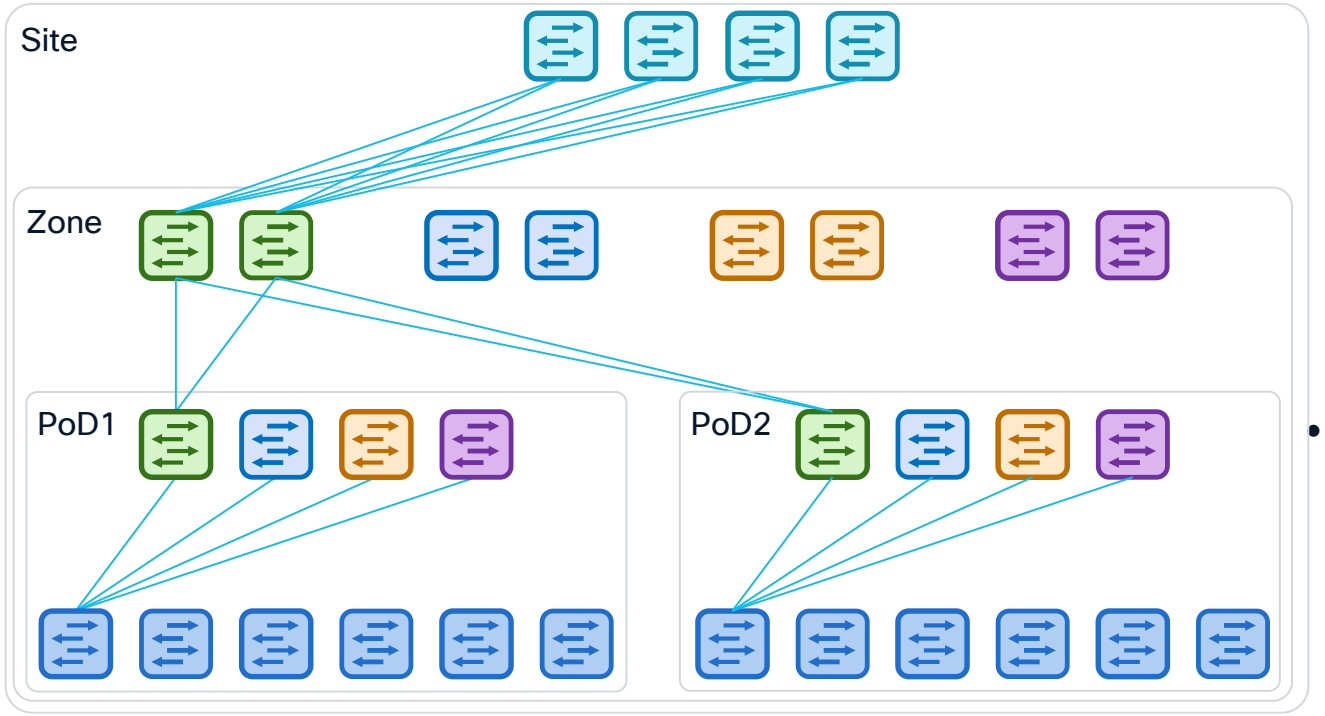

- **Inline RS**
  - Make two of Super spine or Fabric Spine as Route Server
  - Benefit
    - Simple Design and optimized hops
    - Data Plane and Control plane are in-lined
    - Easy to detect RS connectivity failure
- **Separate RS**
  - Connect two RS to Super Spine
  - Benefit
    - Can use 3rd party RS
  - Cons
    - Increased Hops from Leaf to RS
    - Rely on BGP keepalive to detect RS failure

Unplanned

Planned

Virtual Port Channels (vPC)

Achieving High Availability with vPC

Providing HW redundancy with vPC

- Dual Control Plane
- Active/Active Data Plane
- Cisco Fabric Service (CFS)
  - Protocol State Synchronization
  - System Configuration

  - The traffic flow is **symmetric**

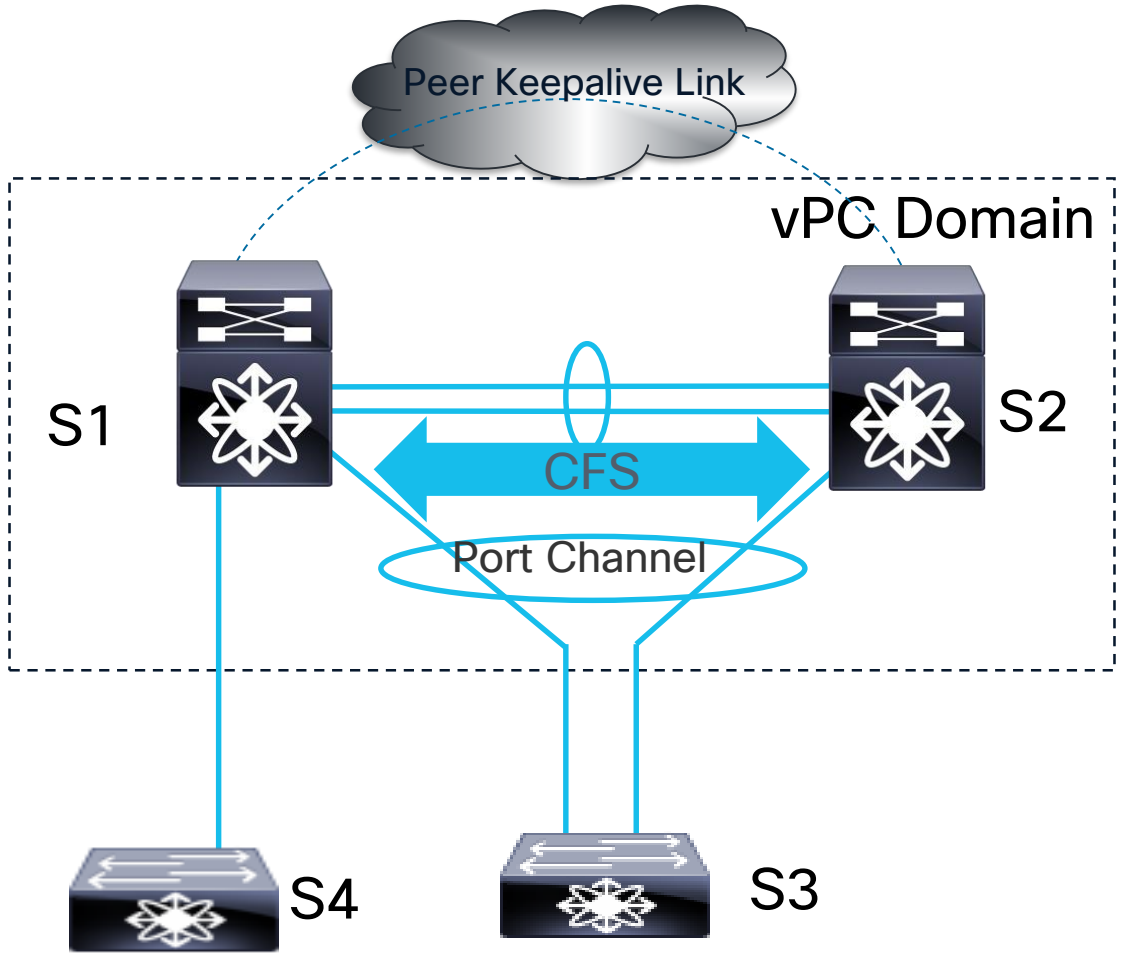

vPC Feature Overview

vPC Terminology

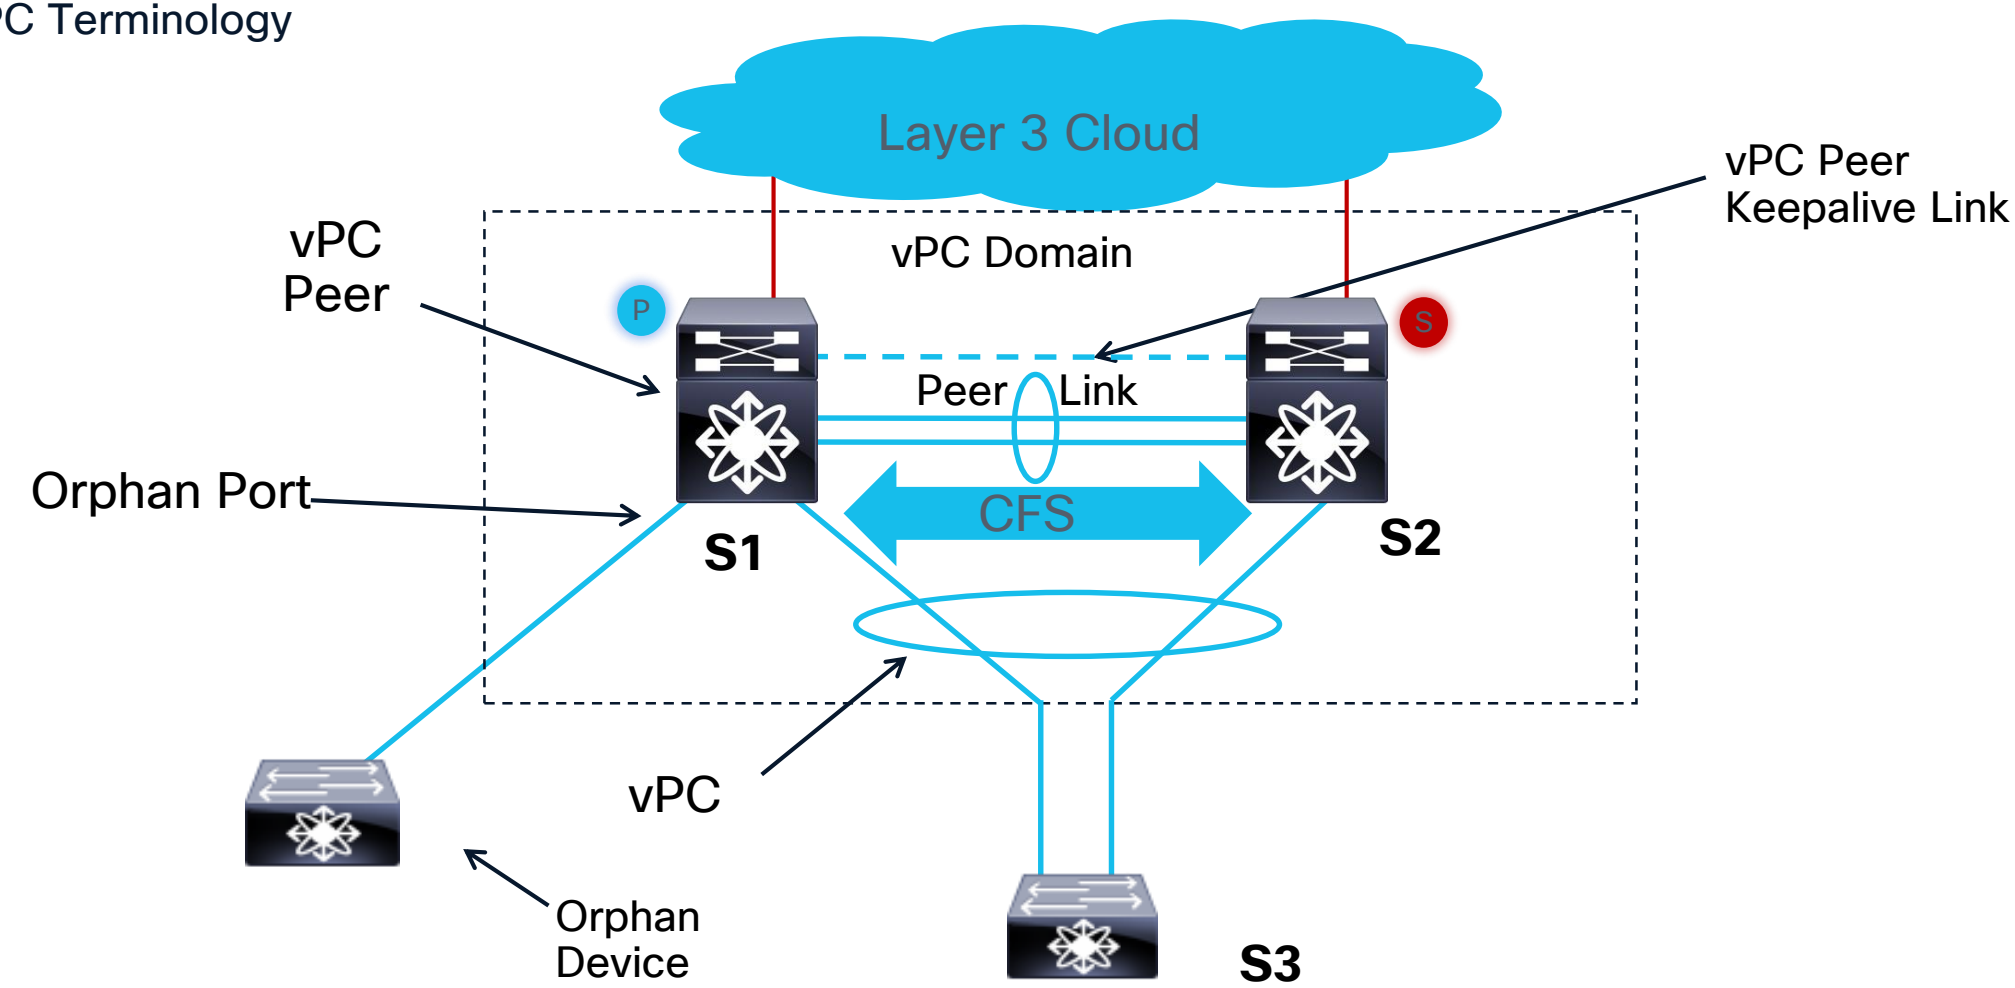

vPC is supported on All Cisco Nexus switches

vPC Roles

- Primary election is based on role priority
- Rules of election:
  - Higher priority wins
  - Lower system mac wins
- Role is non-preemptive
- Operational role may be different from the priorities configured under the domain
- vPC role defines which of the two vPC peers processes BPDUs

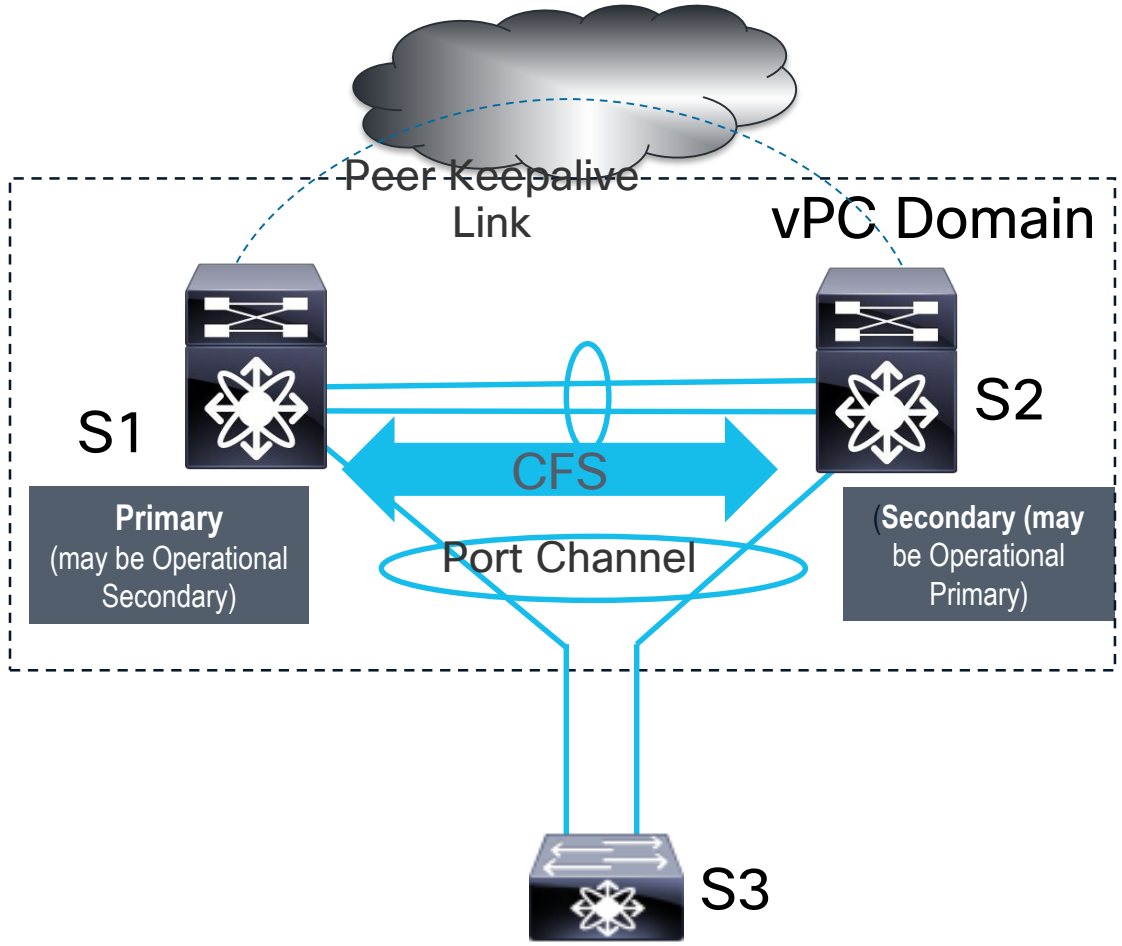

vPC Roles

- Primary election is based on role priority
- Rules of election:
  - Highe
  - Lowe
- Role is non-preemptive
- Operational role may be different from the priorities configured under the domain
- vPC role defines
  processes BPDU

```
switch(config-vpc-domain)# role priority ?
<1-65535>  Specify priority value
```

```
switch# sh vpc
<snip>
vPC role                : secondary, operational primary
```

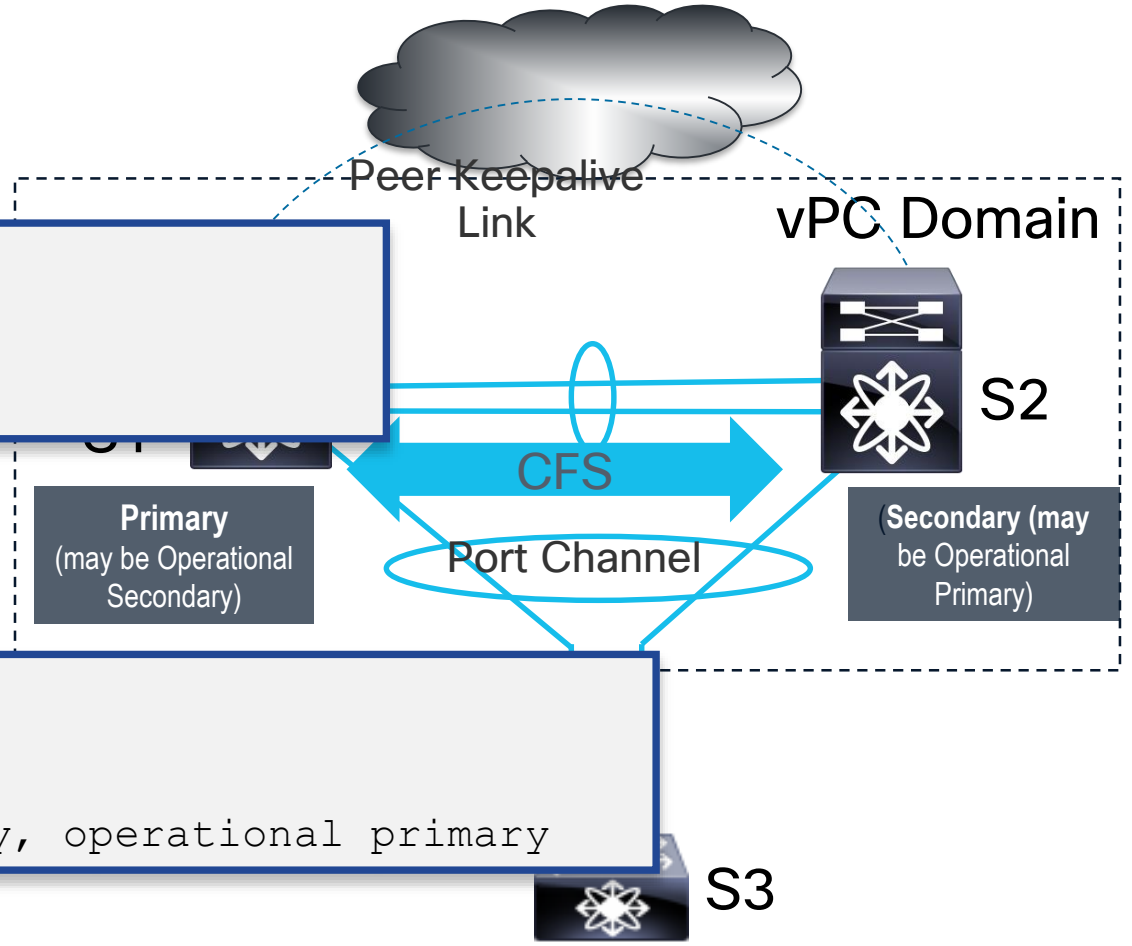

vPC in ACI Fabric

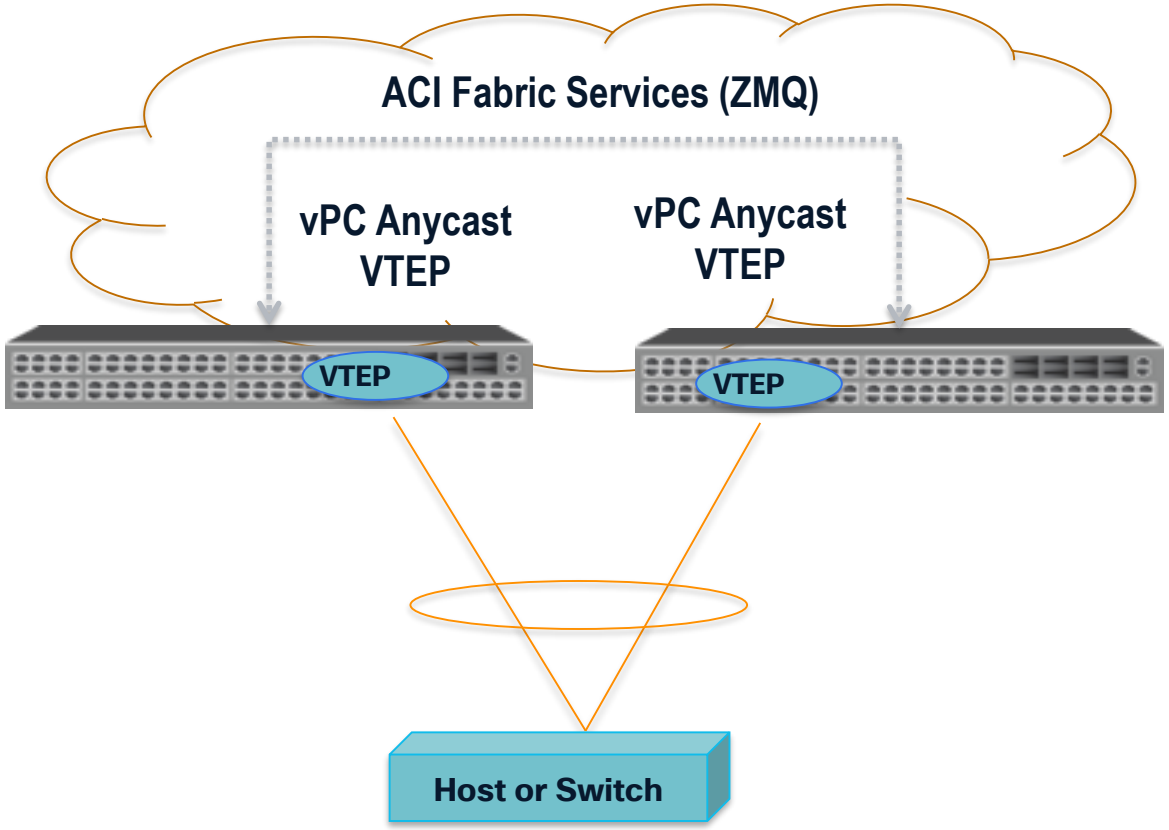

- Differences between ACI vPC and standard vPC
  - No Peer Link is required
  - Peer communication/path recovery happens via the Fabric
  - CFS (Cisco Fabric Services) is replaced by IFS (ACI Fabric Services) which is based on Zero Message Queue (ZMQ)
  - Forwarding selection (which peer will forward a frame
- Within the Fabric the vPC interfaces use an anycast VTEP which is active on both vPC peers

VSS vs vPC

Catalyst

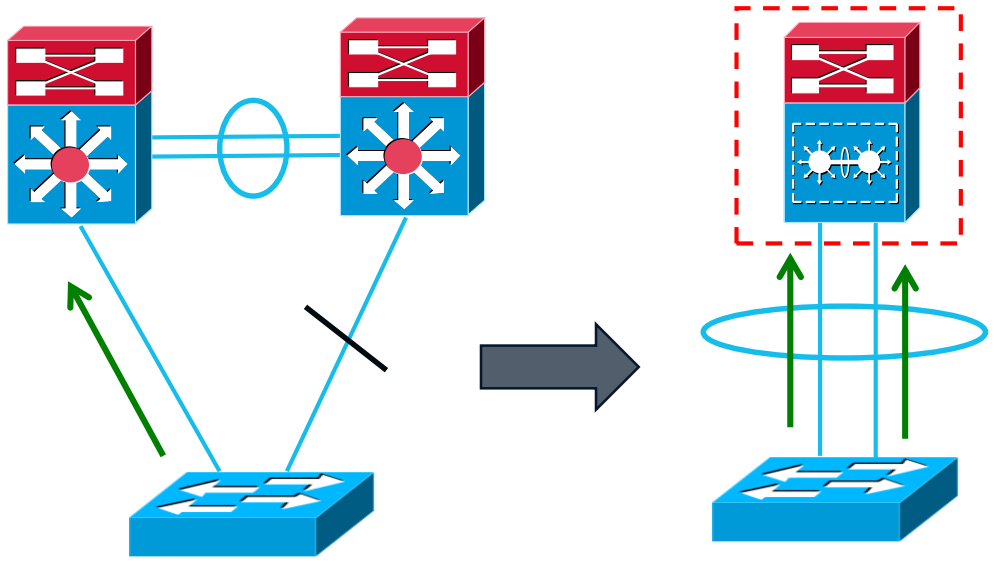

Non VSS

VSS

Merge Data and Control Plane

Nexus

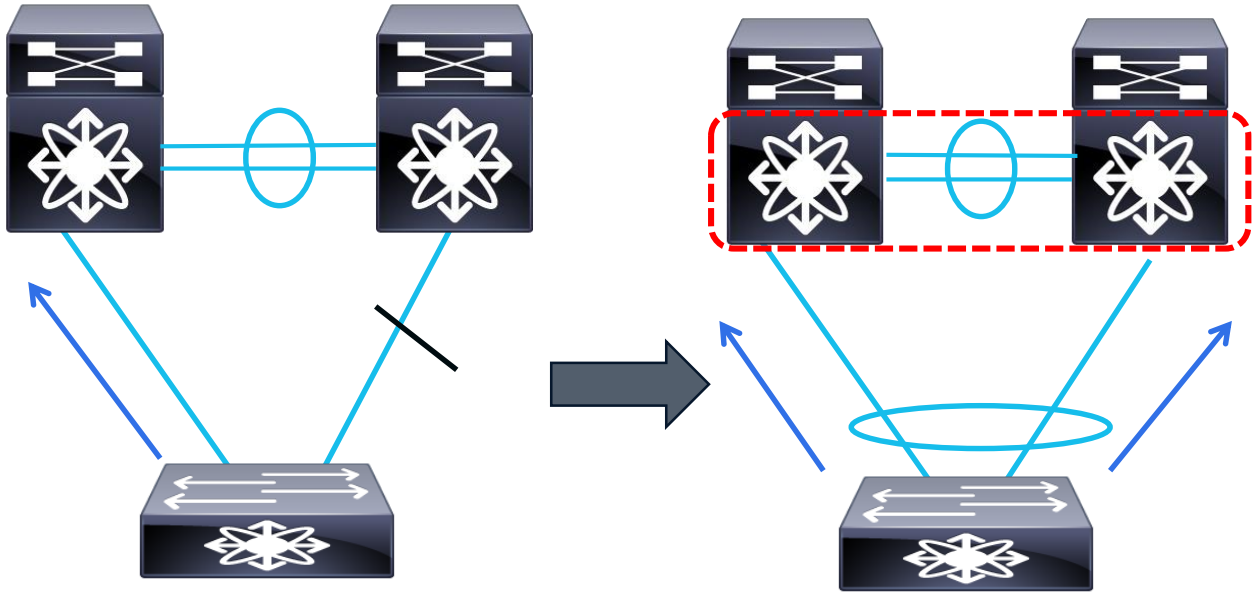

Non vPC

vPC

Merge Data Plane only!!
Control Plane still separate

vPC Failure Scenario

vPC Peer-Link down

- vPC peer-link failure (link loss)
- vPC peer-keepalive up
- Status of other vPC peer known
- Secondary vPC peer disables all vPC's
- Traffic forwarded to vPC primary

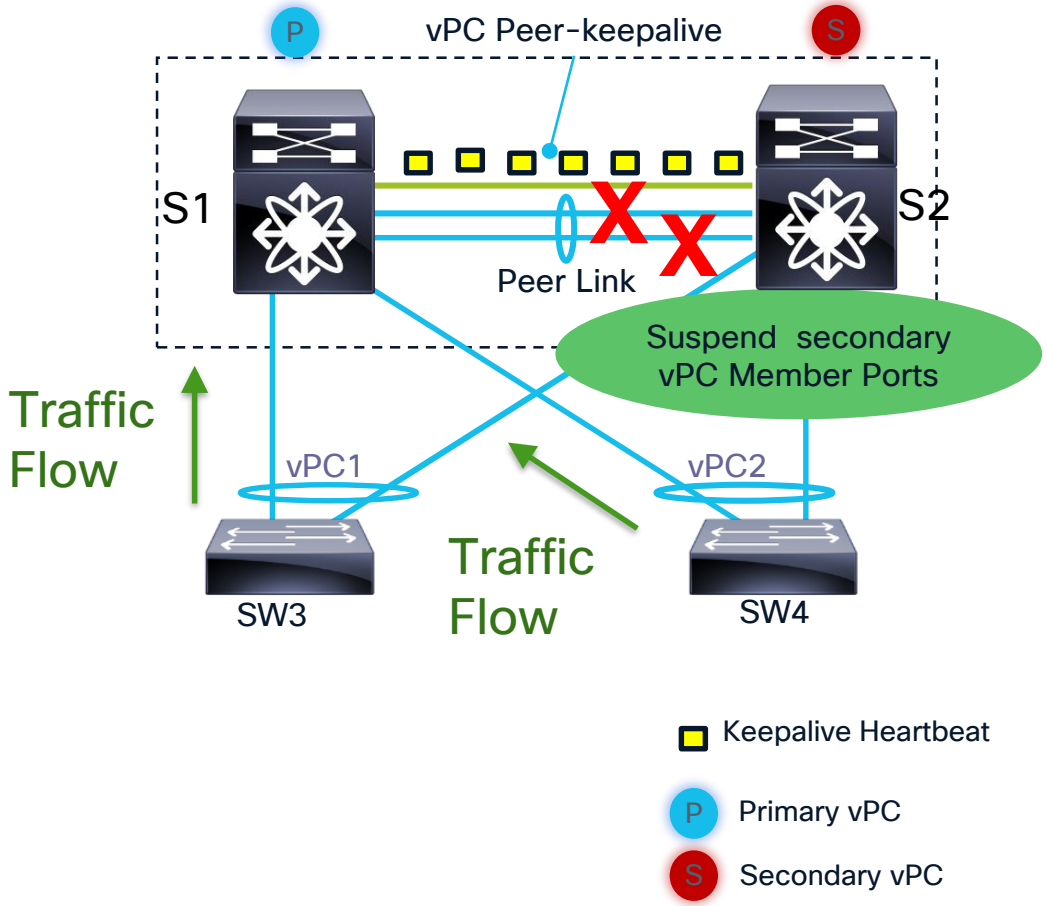

vPC Failure Scenario

vPC Peer KeepAlive Link down

- vPC peer-link is up
- vPC peer-keepalive goes down
- Status of other vPC peer known
- No Impact to the system
- Traffic continues to flow normally

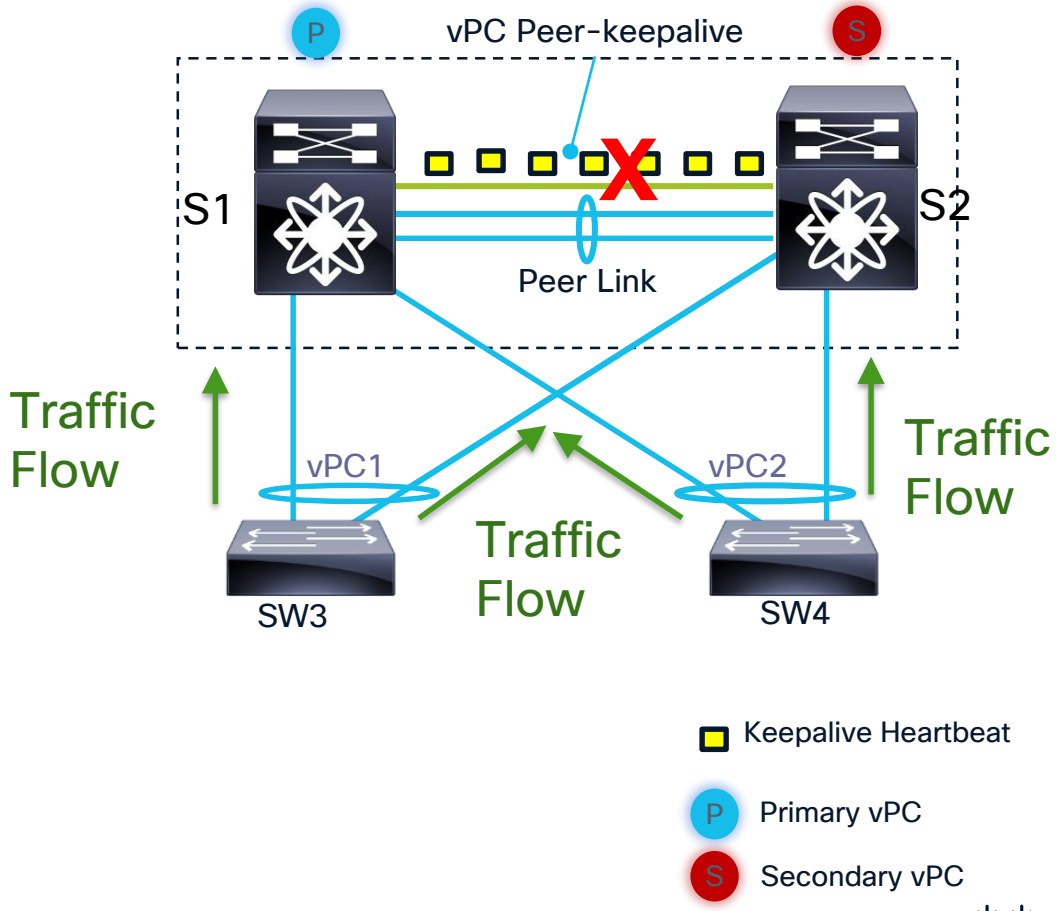

vPC Failure Scenario

Primary Peer Fails

- Primary Peer goes down
- Secondary Detects via Keep alive link
- Secondary Takes on Primary Role
- Traffic will be forwarded to new vPC primary

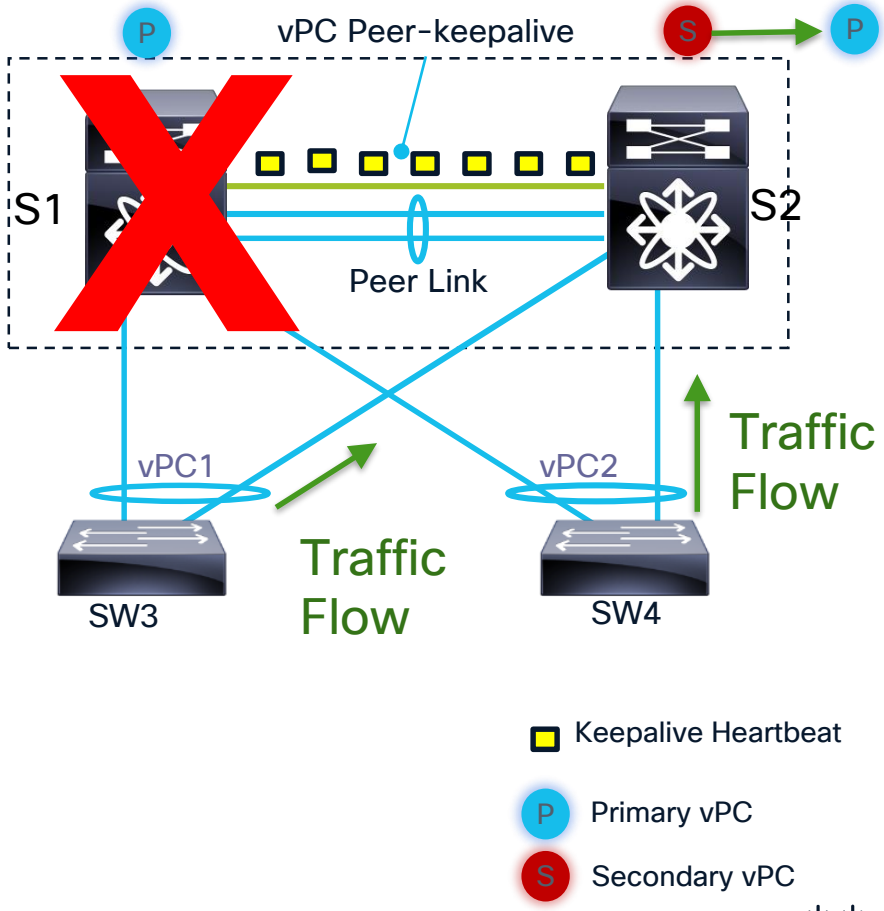

vPC Failure Scenario

Secondary Peer Fails

- Secondary Peer goes down
- Primary Detects via Keep alive link
- Traffic will be forwarded to vPC primary

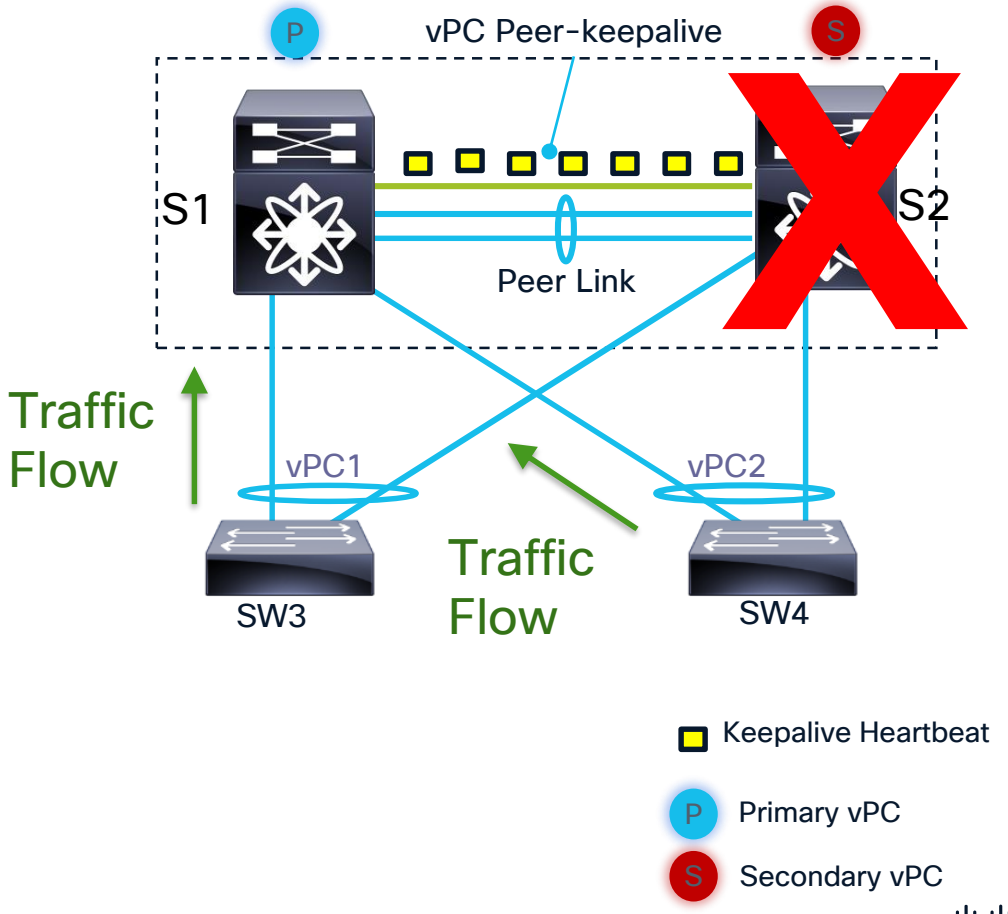

vPC Failure Scenario

Dual Active or Split Brain

- vPC peer-link goes down
- vPC peer-keepalive goes down
- Status of other vPC peer is UNKNOWN
- Secondary also becomes Primary
- Traffic Pattern is Uncertain

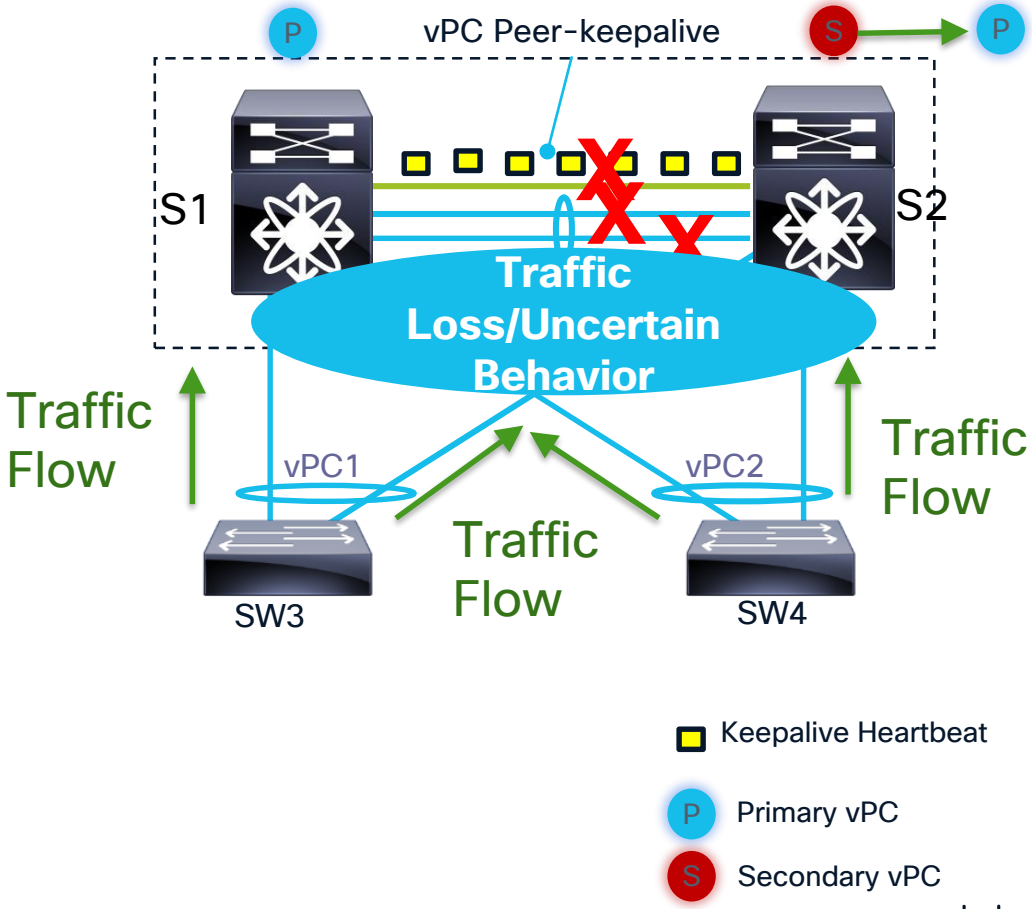

VSS Design Migration to vPC Design

Catalyst VSS

Nexus vPC

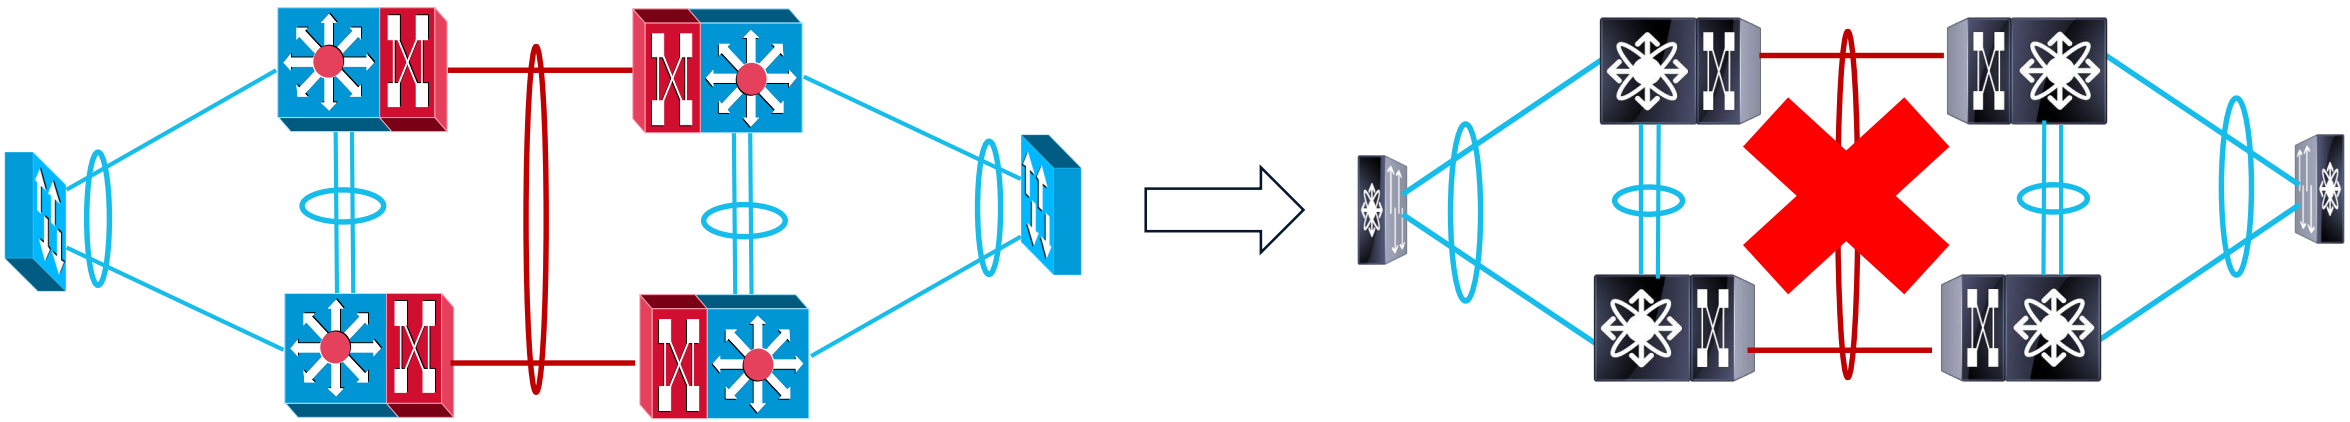

Don’t design VSS and vPC in same way for Layer 3!

VSS Design Migration to vPC Design

Catalyst VSS

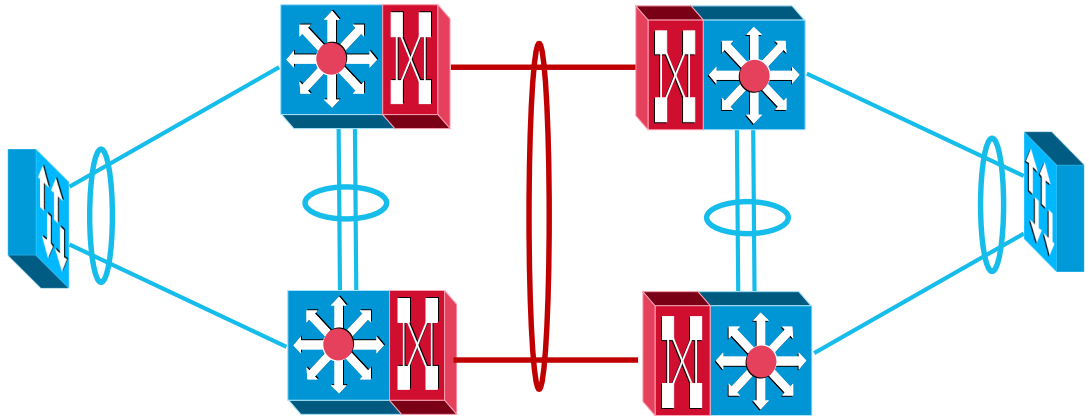

Nexus vPC

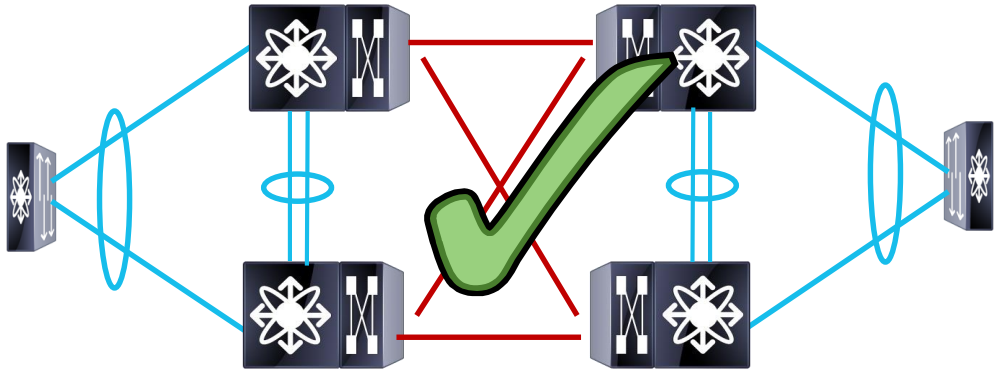

vPC Layer 3 routed uplink with ECMP

vPC Control Plane Recommendations

- For the Peer Link, it is good to use two 10GbE ports on separate line cards for resilient connectivity.
- Peer links should be point-to-point connection.
- The Peer Keepalive traffic should be over a separate keepalive link and not over the Peer Link.

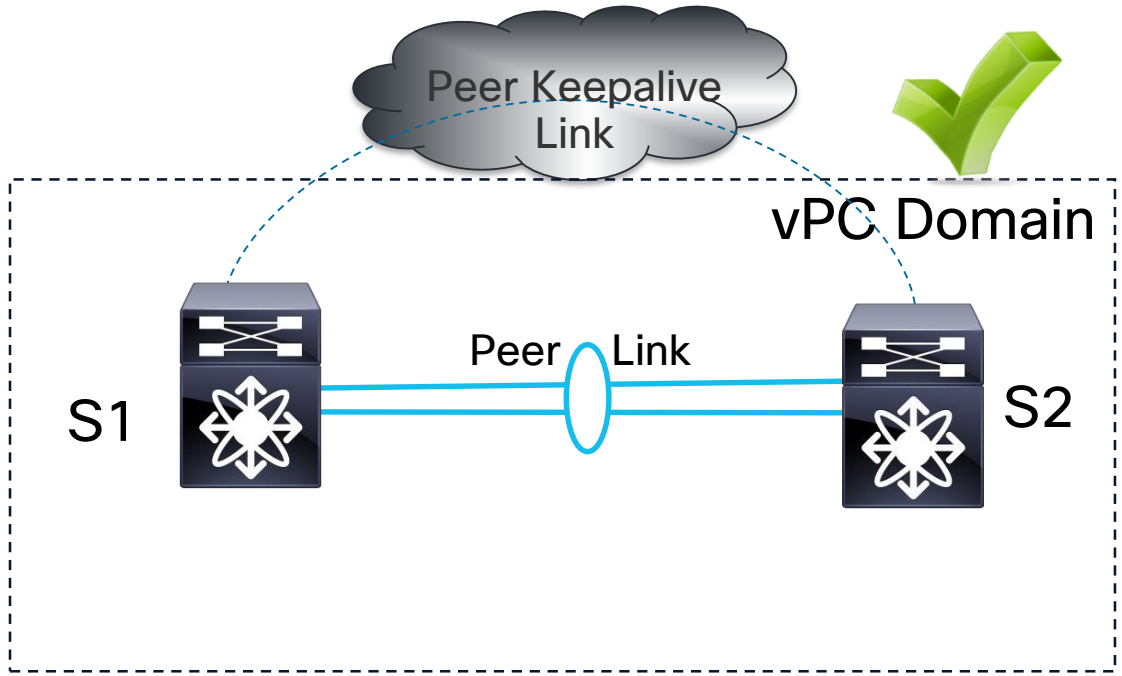

vPC Member connectivity Recommendations

Dual Home Devices

- Make sure to dual attach all devices into the vPC Domain
- It is best to use LACP for the Port Channels of vPC Member ports
- The configuration must match on both switches
- Use the same value for the port channel ID and the vPC ID

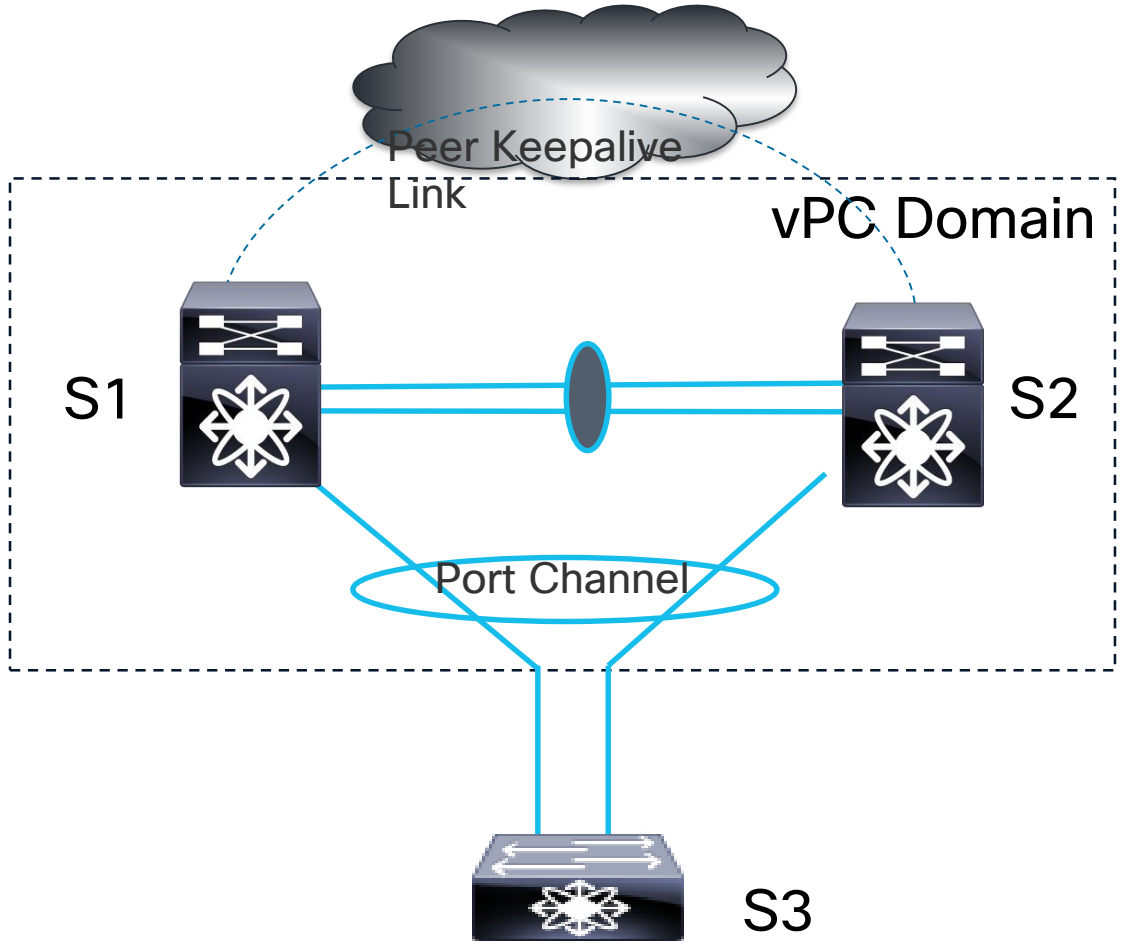

vPC Member connectivity Recommendations

Dual Home Devices

- Make sure to dual attach all devices into the
- It is best to use Port Channels of vPC Member ports
- The configuration must match on both switches
- Use the same value for the port channel ID and the vPC ID

```
VPC-Primary(config)# feature vpc
VPC-Primary(config)# vpc domain 1
```

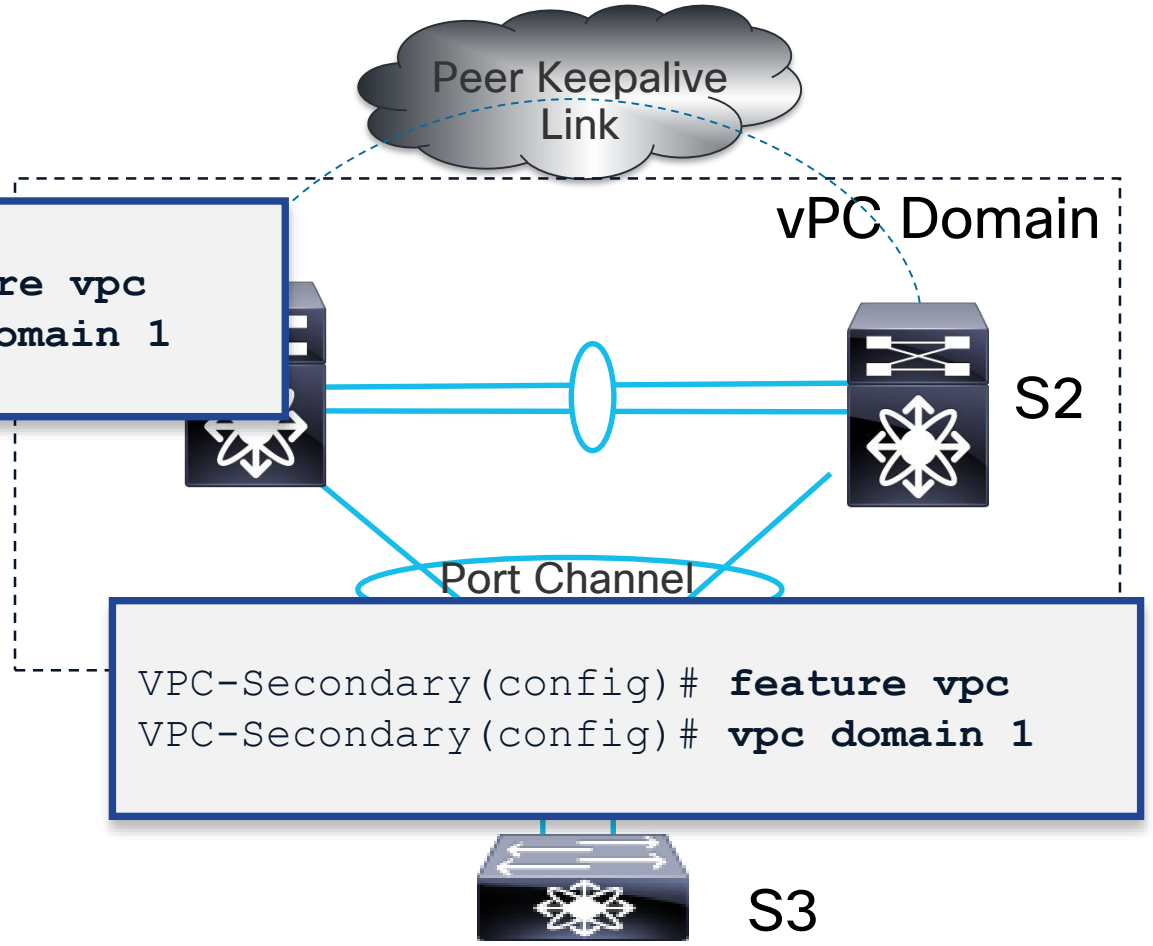

Planned

Graceful Insertion & Removal (GIR)

Graceful Insertion and Removal

Isolation of Switch from network

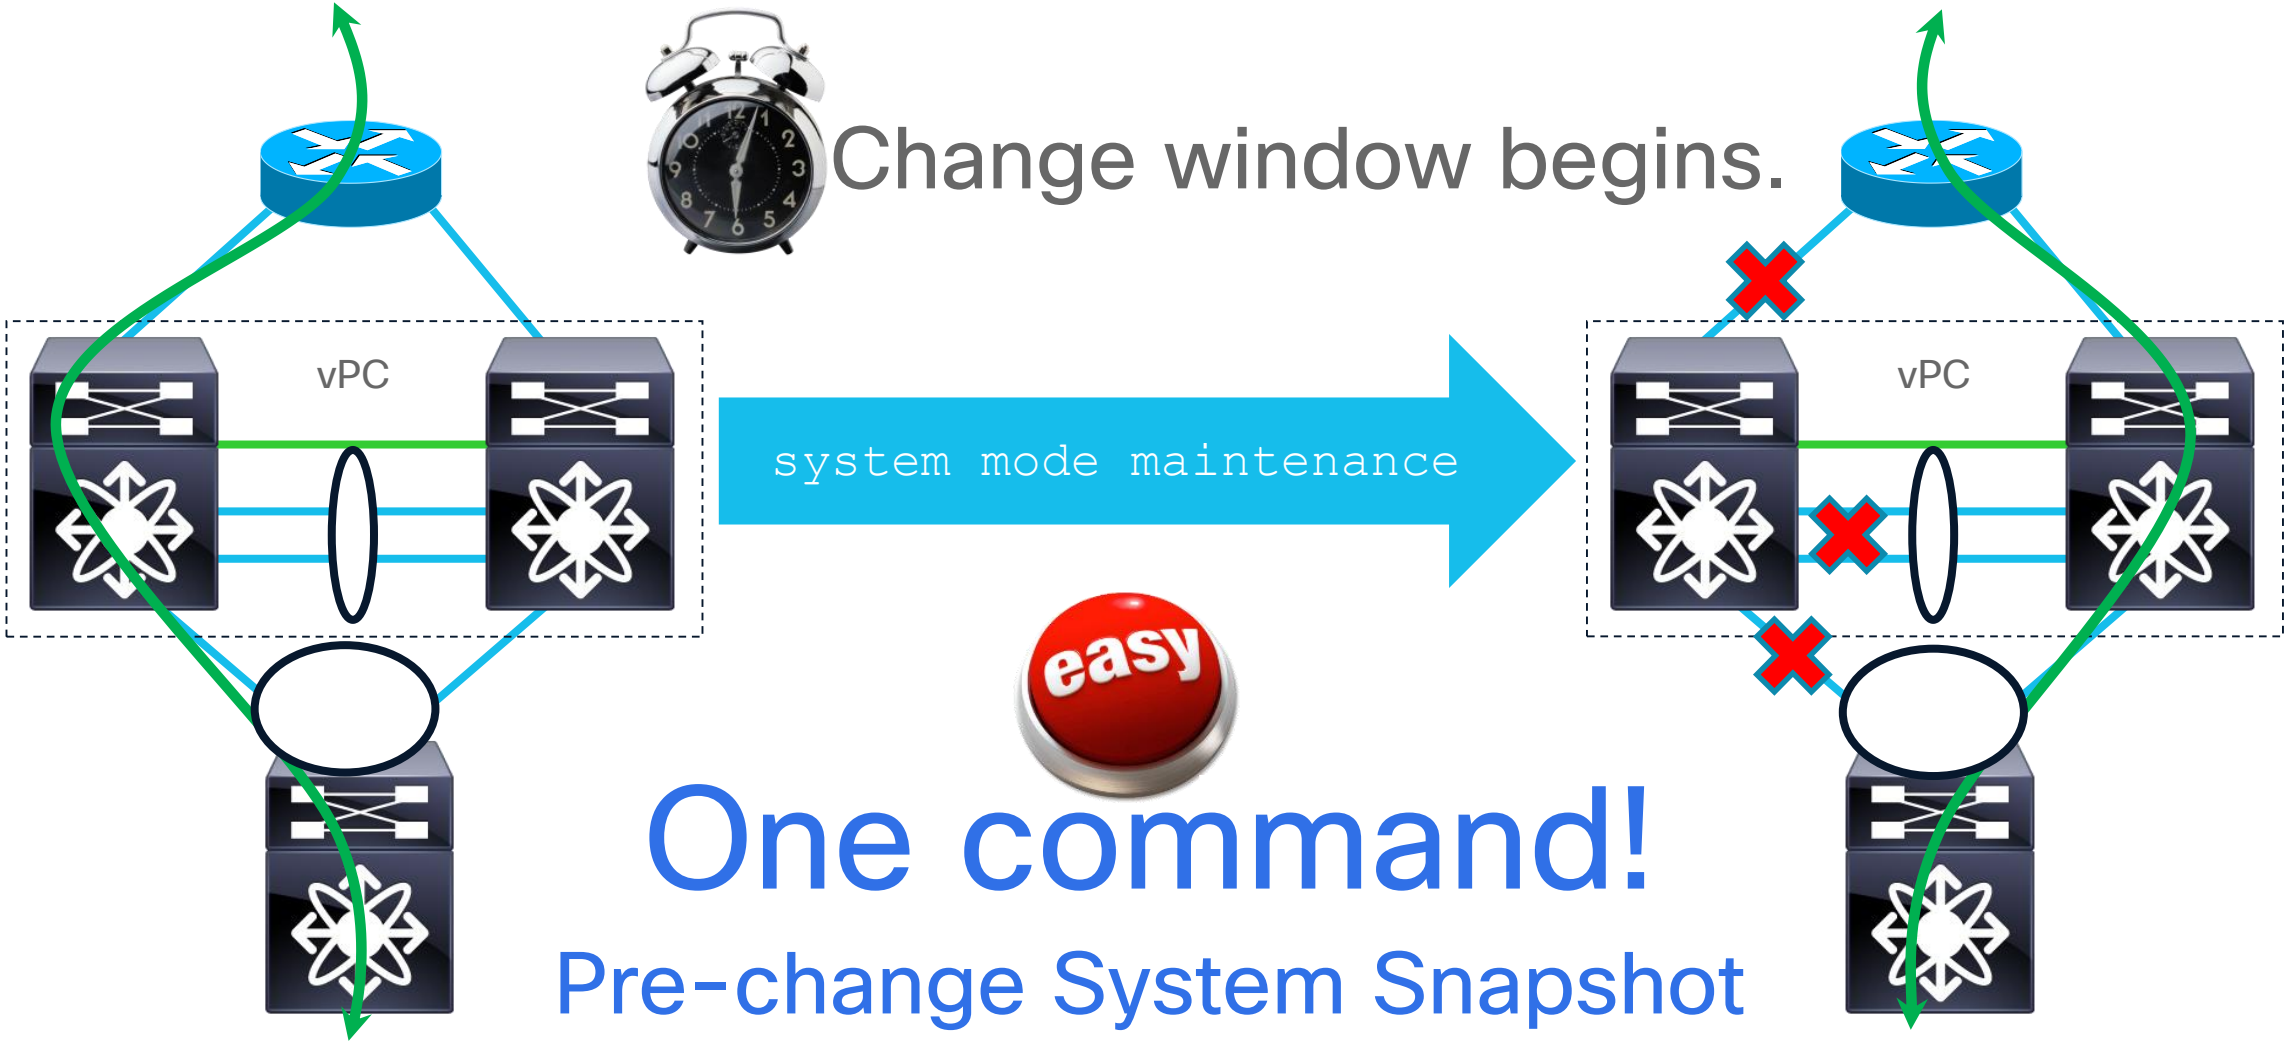

Graceful Insertion and Removal

Return of Switch into network

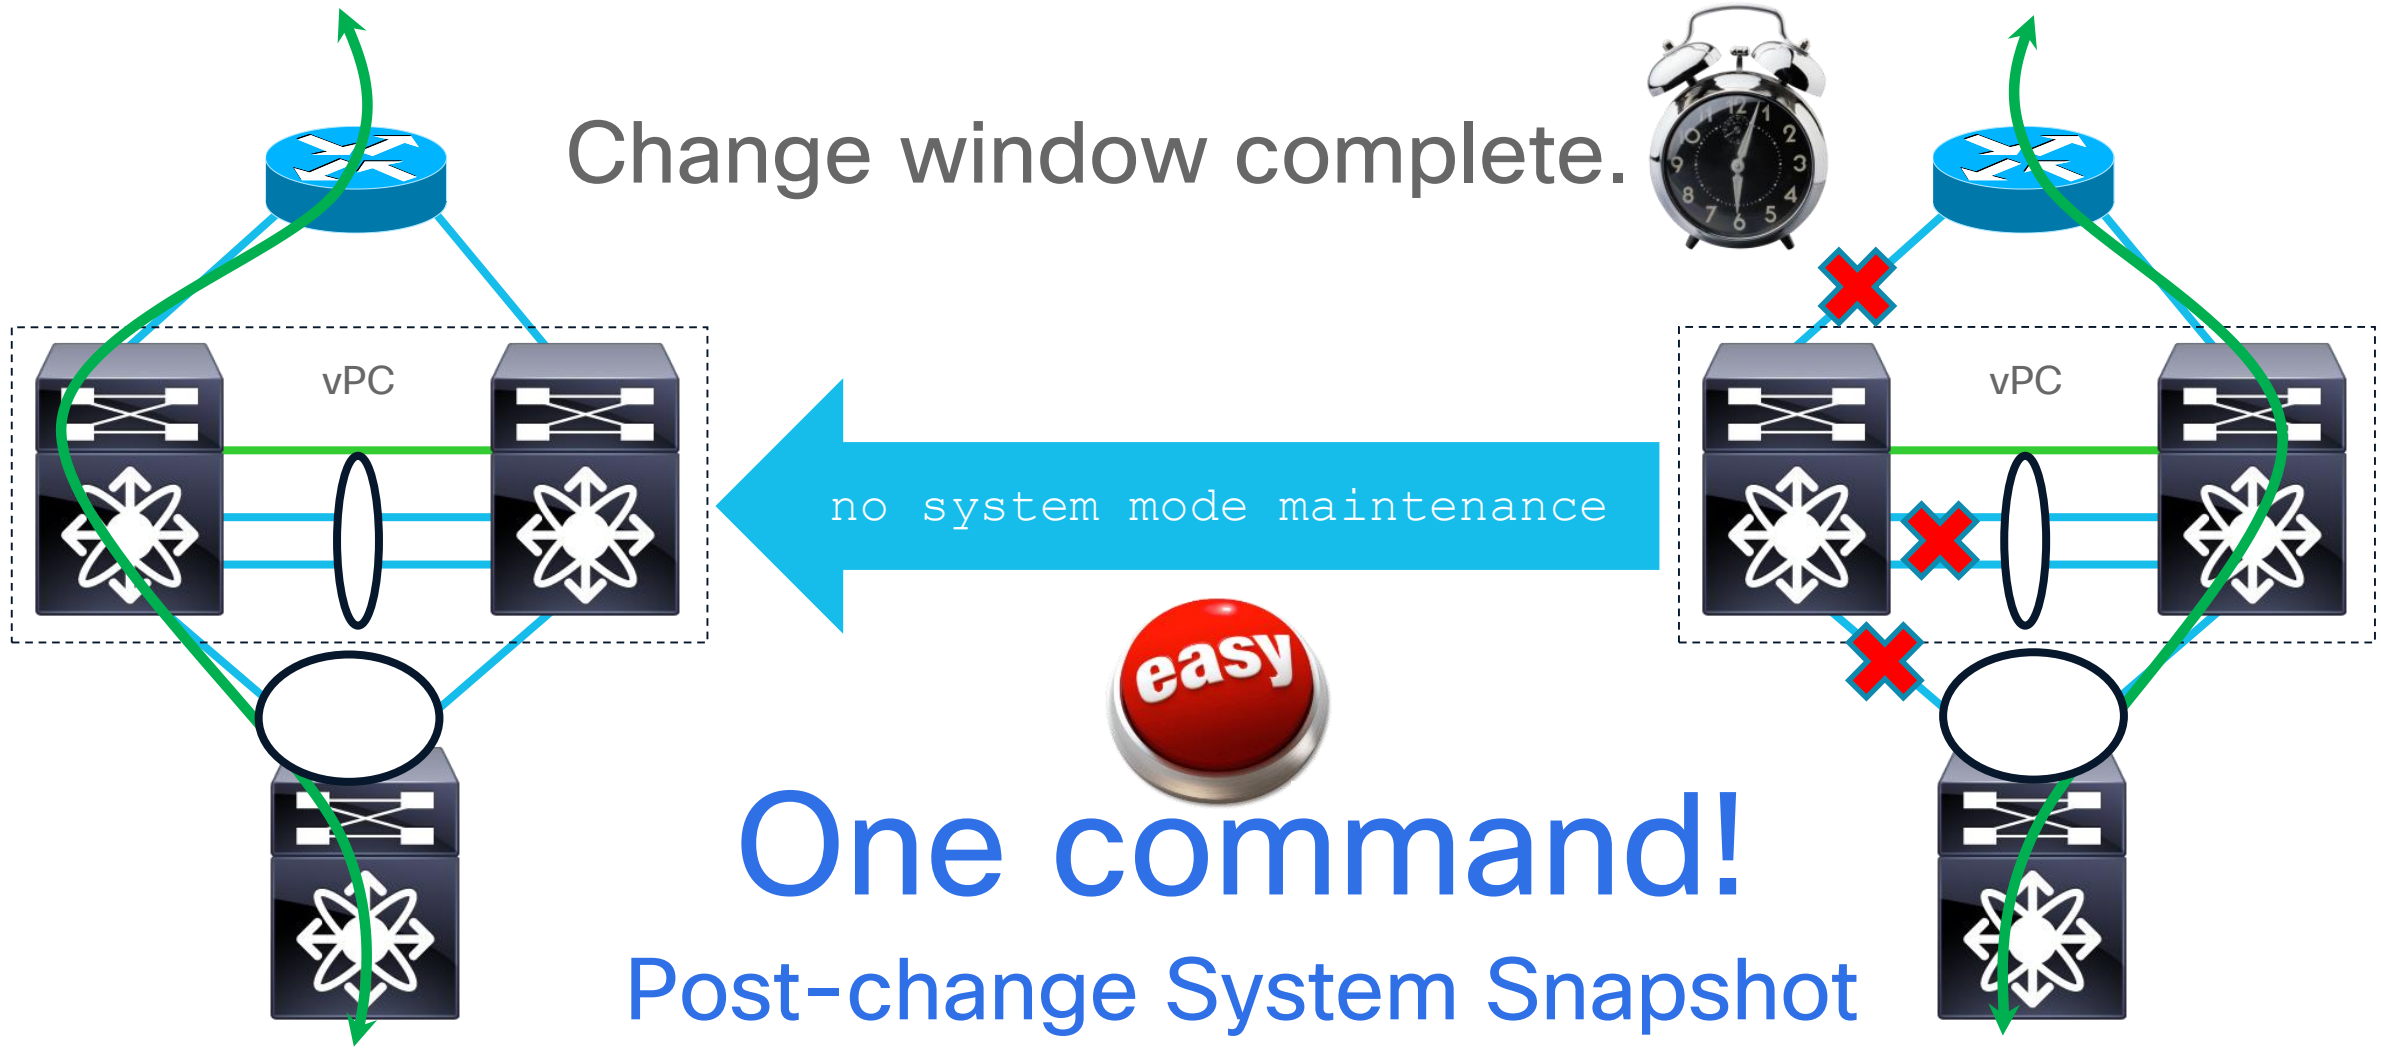

Graceful Insertion and Removal

- 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)

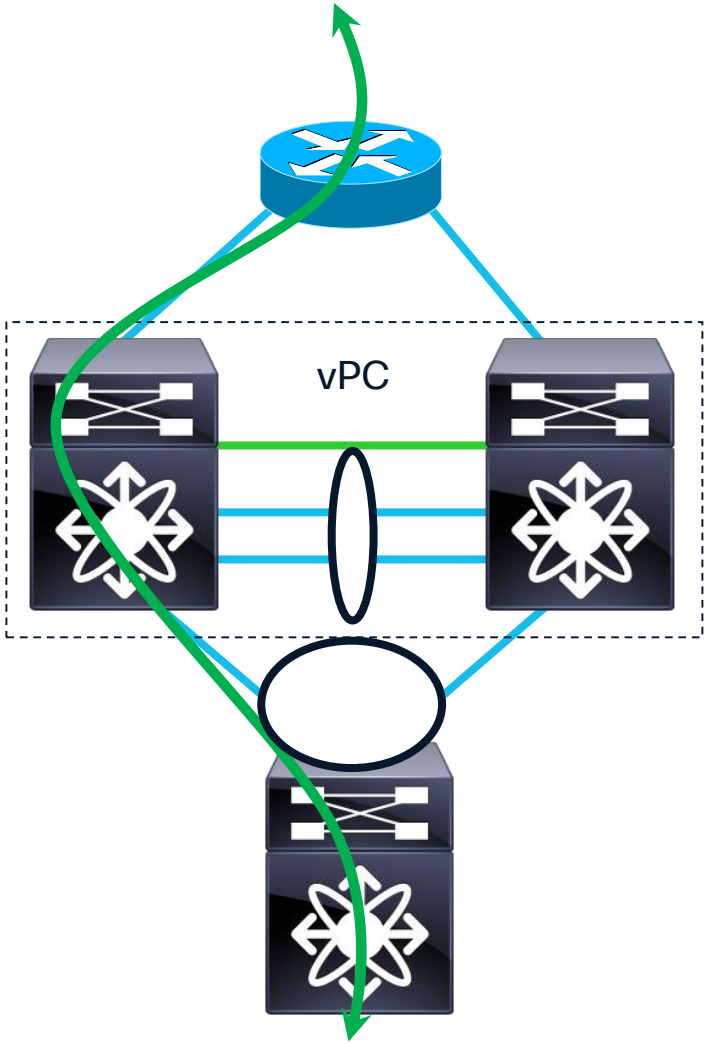

Graceful Insertion and Removal Recommendations

Maintenance Window Timer

- Can be configured to keep the switch in Maintenance Mode for period of time
  - Started on entry to maintenance mode
  - When configured time elapses, switch automatically returns to normal mode.

Recommendations: set the timer value to 60 minutes

Graceful Insertion and Removal Recommendations

Maintenance Window Timer

- Can be configured
  Maintenance Mod

```
switch# configure terminal
switch(config)# system mode maintenance timeout <Value>
        <5-65535>  Timer value in minute
```

  - Started on entry to maintenance mode
  - When configured time elapses, switch automatically returns to normal mode.

Recommendations: set the t
60 minutes

```
switch# show maintenance timeout
Maintenance mode timeout value: 20 minutes
```

Graceful Insertion and Removal Recommendations

Similar to vPC, 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 devices that are singly attached

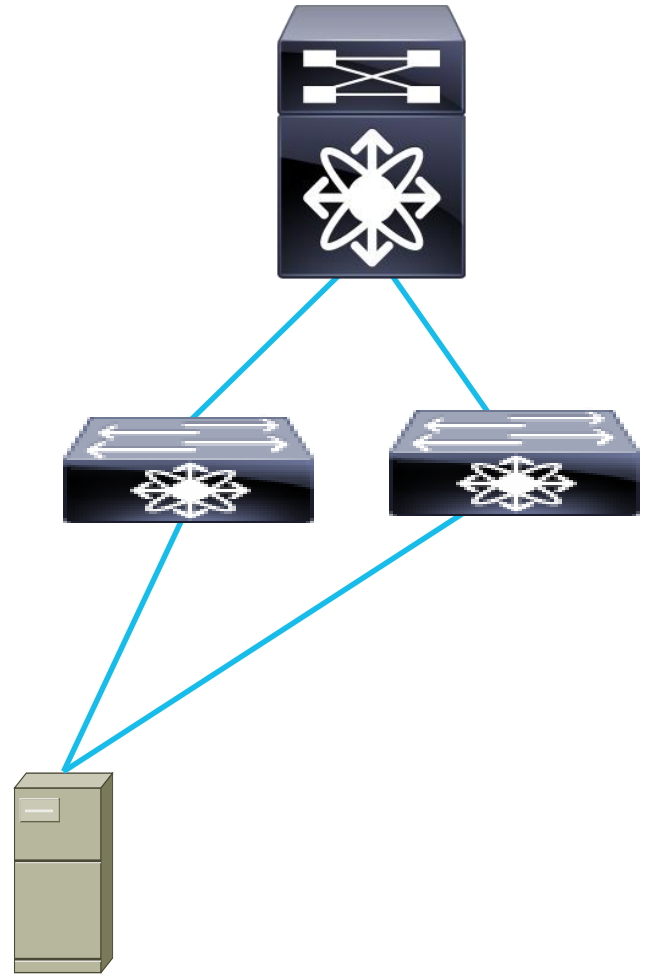

Putting it all Together

What is the recommendation?

	Critical Bug fix & PSIRT	Hardware Upgrades	New Image Version
SMU Patching	Recommended ✓	Not Recommended X	Not Recommended X
ISSU	Possible ✓	Not Recommended X	Recommended ✓
GIR	Not Recommended X	Recommended ✓	Not Recommended X
Box reload (Cold Boot)	Possible ✓	Not Recommended X	Possible ✓

Hitless Upgrade Models

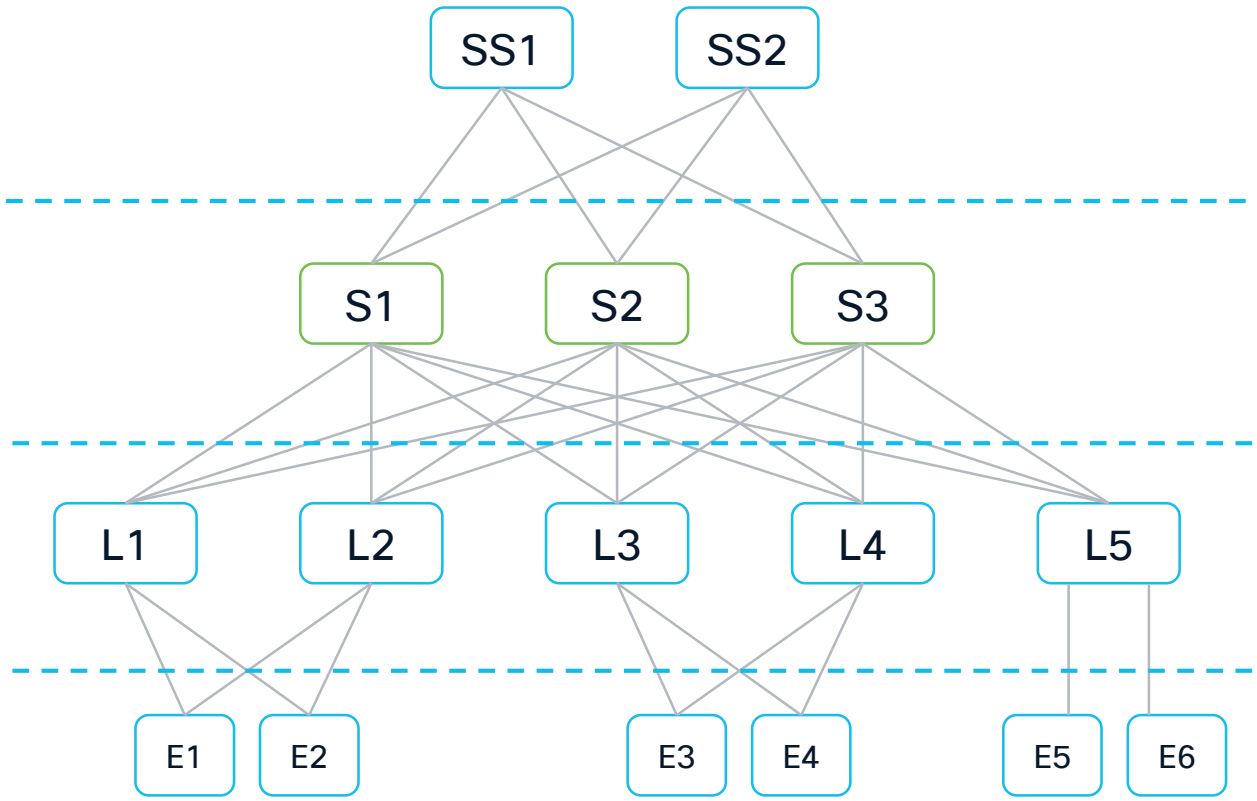

Platform options	Recommended upgrade options
Modular, Fixed	✓ GIR/maintenance mode ✓ SMU (Patch)
Modular, Fixed	✓ GIR/maintenance mode ✓ SMU (Patch)
Fixed	✓ GIR/maintenance mode ✓ SMU (Patch) ✓ ISSU*

GIR, SMU support on all Nexus

ISSU/Enhanced ISSU on NX-OS fixed switches

Image upgrade automation from Controller

Conclusion

Your network is as reliable as its weakest component

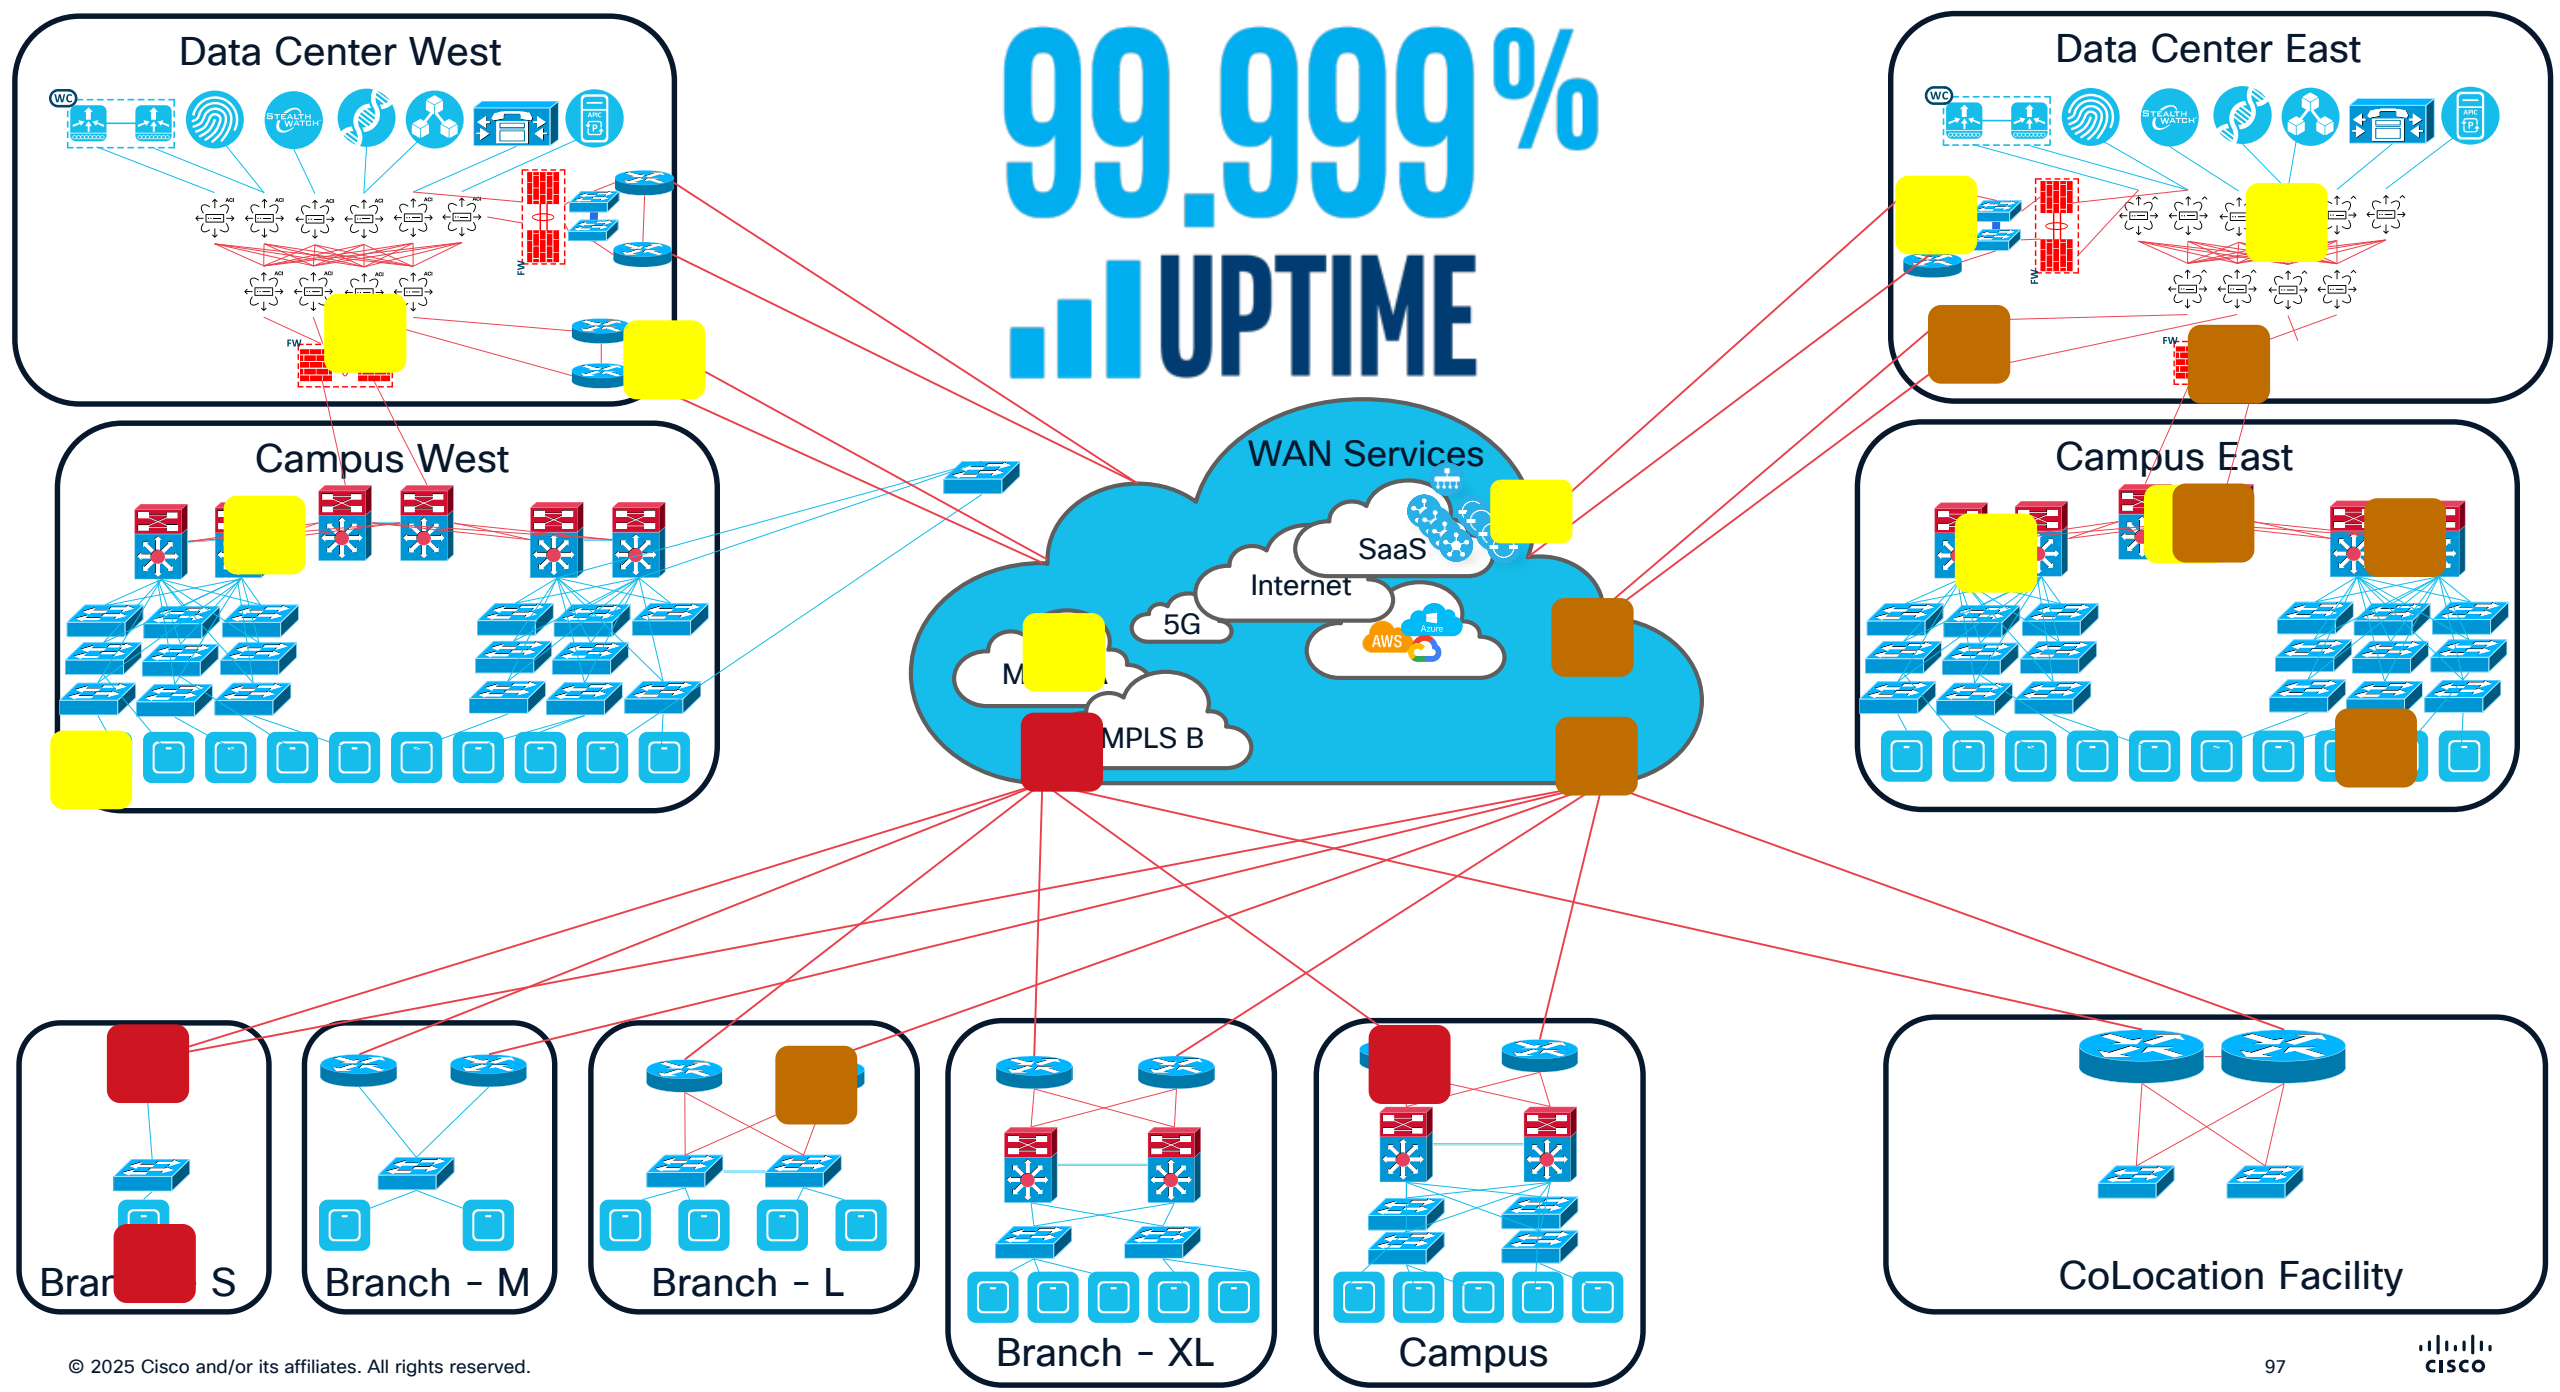

References

For Your Reference

- Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/high_availability/guide/b_Cisco_Nexus_9000_Series_NX-OS_High_Availability_and_Redundancy_Guide_7x.html
- Understanding Network-Level High Availability: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/high_availability/configuration/guide/ha_network.html
- ACI Upgrade Best Practices and Troubleshooting: https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/213618-aci-upgrade-best-practices-and-troublesh.html
- Configuring APIC High Availability: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/cli/nx/cfg/b_APIC_NXOS_CLI_User_Guide/b_APIC_NXOS_CLI_User_Guide_chapter_010101.pdf

Complete your session evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to win 1 of 5 full conference passes to Cisco Live 2026.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn exclusive prizes!

Complete your surveys in the Cisco Live mobile app.

Continue your education

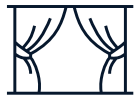

Visit the Cisco Showcase for related demos

Book your one-on-one Meet the Engineer meeting

Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs

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

Thank you

cisco Live !


CISCO