The bridge to possible

Why You Shouldn't Fear Upgrading Your ACI Fabric

The Handbook!

Takuya Kishida – Technical Leader, Technical Marketing Engineering
Joey Ristaino – Technical Leader, Technical Marketing Engineering
BRKDCN-2910

cisco Live!

#CiscoLive

Cisco Webex App

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2910

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 7, 2024.

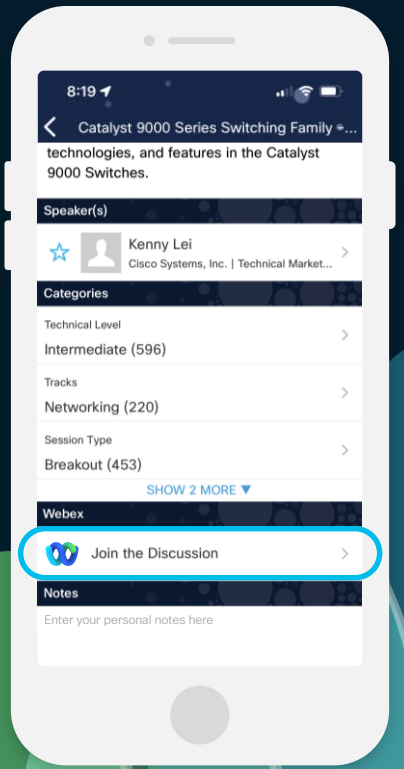

Agenda

- Why Upgrade?
- Upgrade Architecture
  - ACI Firmware Upgrade Types
  - Upgrade Architecture – APIC
  - Upgrade Architecture – Switches
  - (Bonus) Upgrade Enhancements
- Best Practices
  - Best Practices Workflow Review
  - Best Practices Configurations
  - “Pre-Upgrade Checklist” Review and Execution
  - “Do’s and Don’ts”

Why Upgrade?

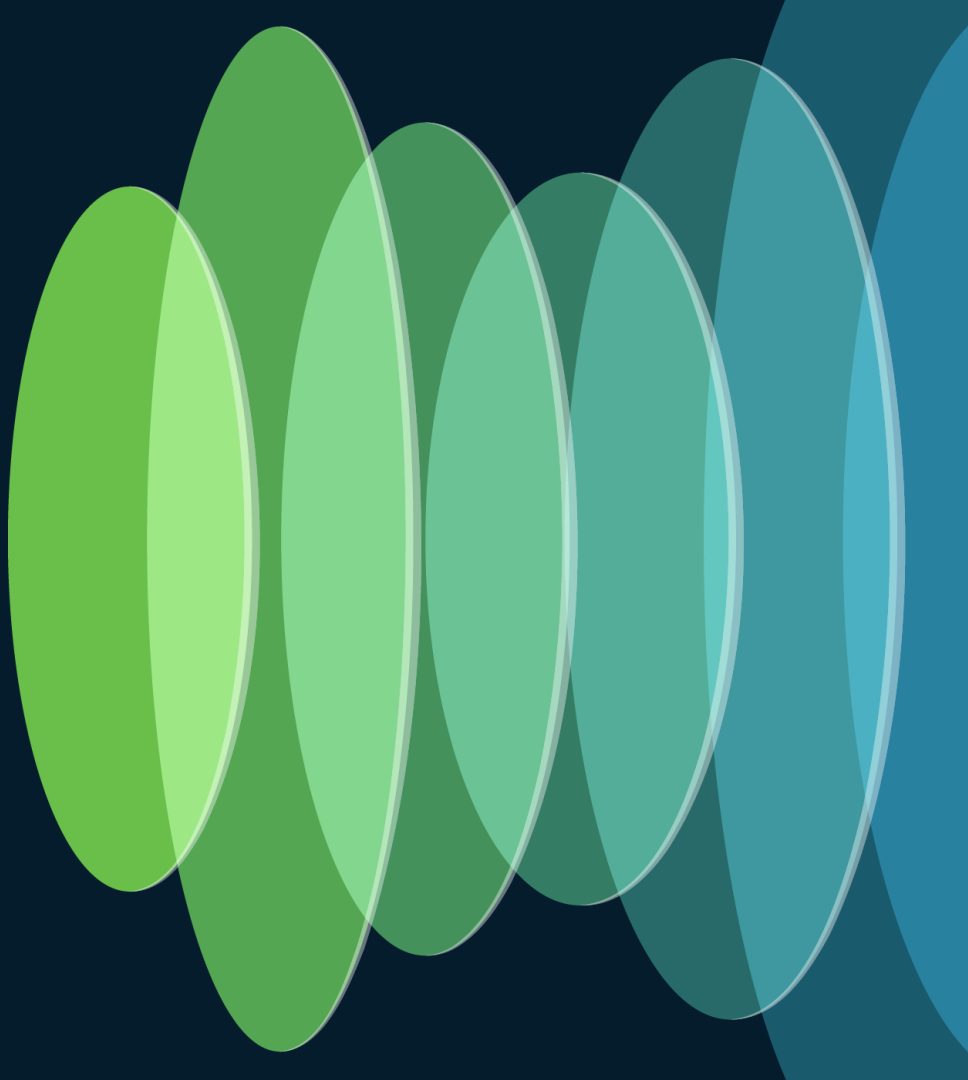

New features, new use cases, new possibilities

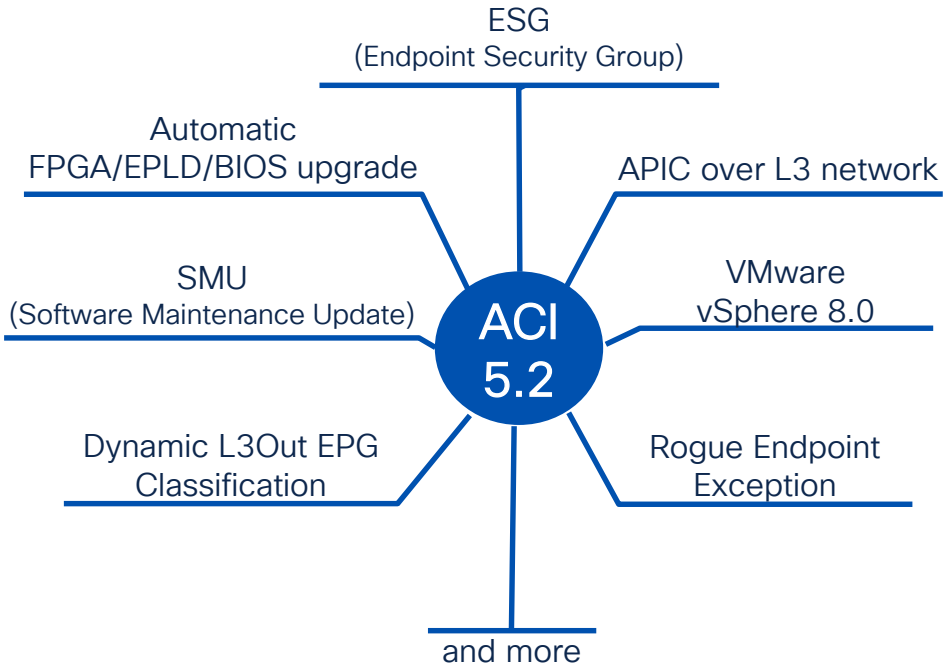

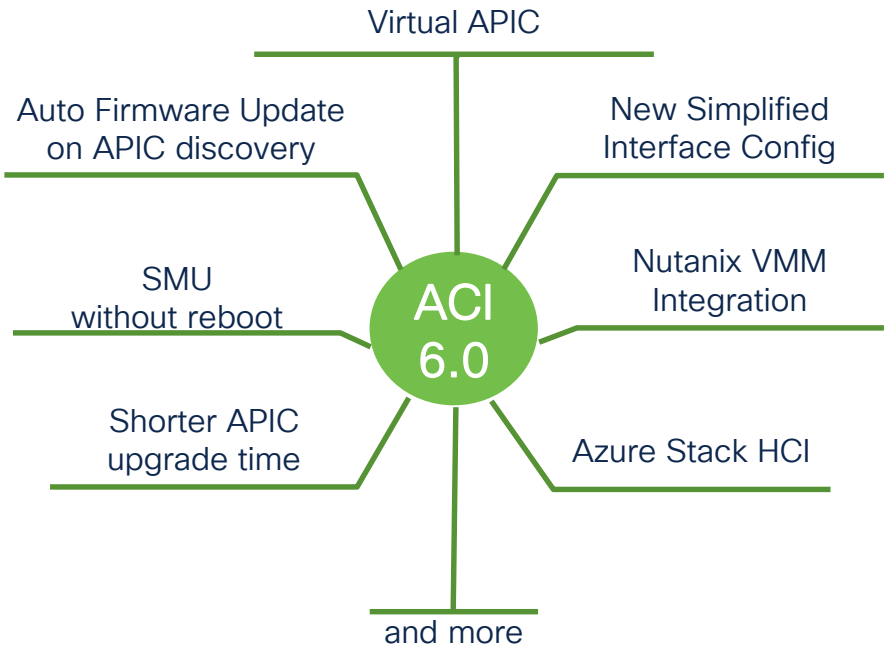

And more with NDI (Nexus Dashboard Insight) for ACI 4.2(5) or newer

Get out of the vicious cycle

By knowing the upgrade architecture and best practices

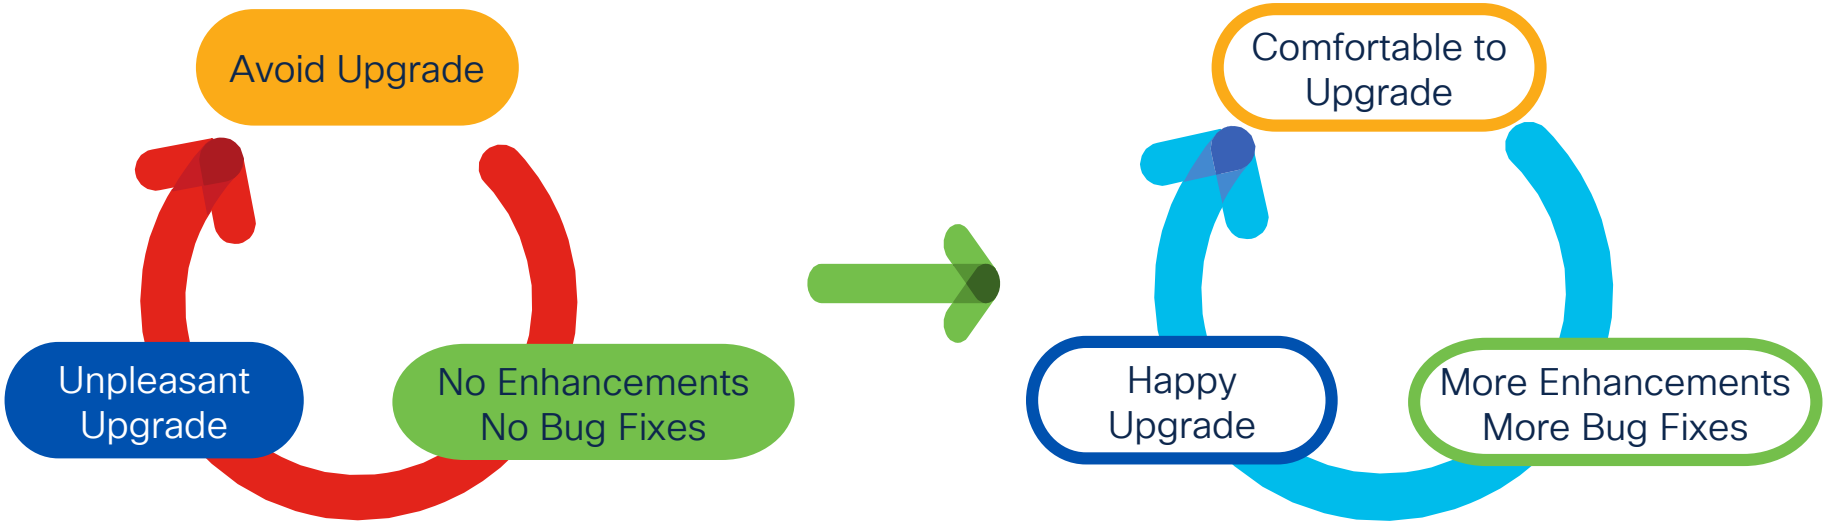

ACI Firmware Upgrade Types

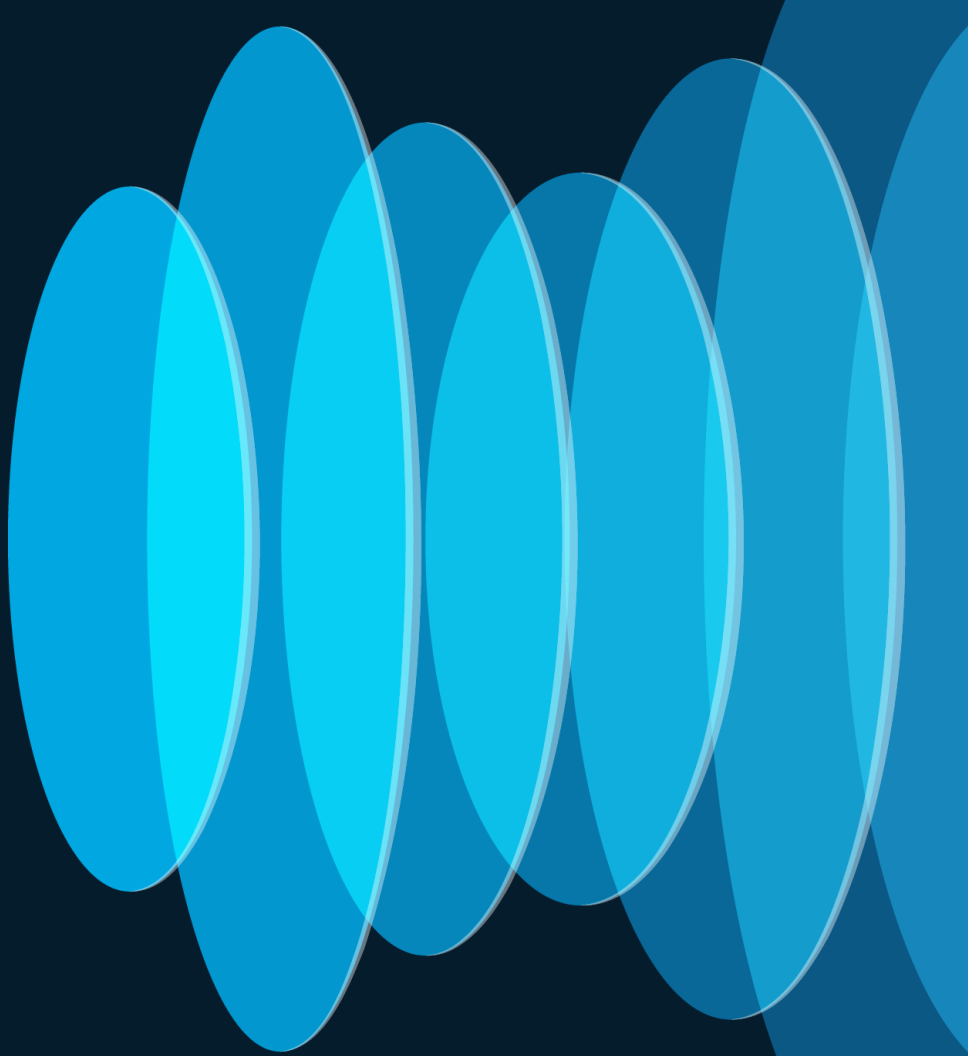

ACI Firmware Upgrade Types

ACI Firmware Upgrade Types (Regular)

- Regular Upgrade
- Software Maintenance Upgrade (SMU)
- EPLD/FPGA Upgrade (Only Switches)

Base OS firmware upgrade

In principle, all APICs and switches should be on the same version

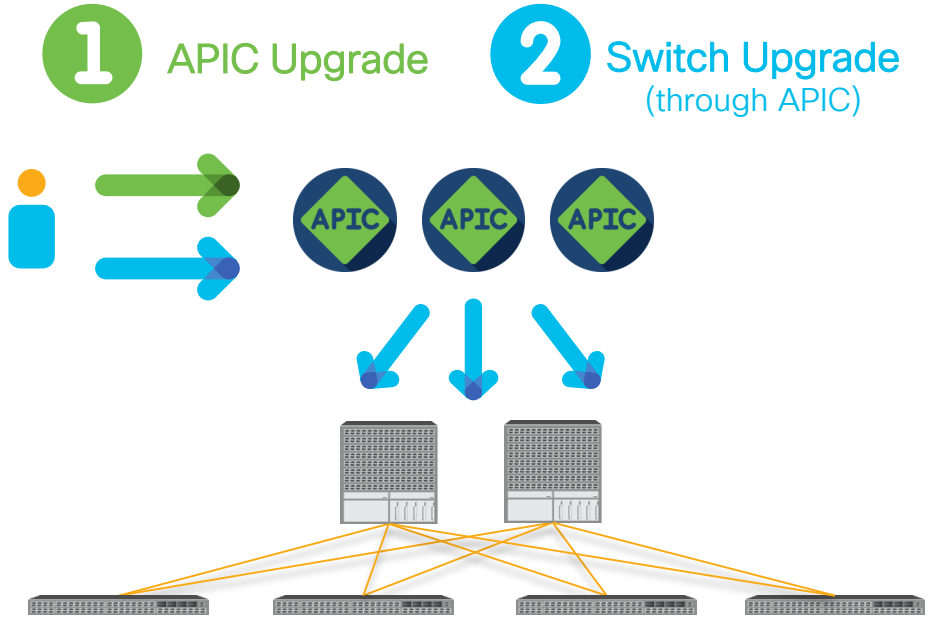

Different versions in the same fabric??

In principle, this should be avoided.

What if I cannot finish upgrades in a single upgrade window?

- Available options

  APIC firmware

  ➢ All APICs must be on the same version

  Switch firmware

  ➢ Switches can be on different versions with limited operations.

Create, update and delete BDs, EPGs, contracts, L3Outs, VMM domains, Access Policies

Collect configuration backups, techsupports, or troubleshoot with SPAN

Physical operations such as enabling disabling interfaces, replacing a node

See Upgrade Guide for the complete list:
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/all/apic-installation-aci-upgrade-downgrade/Cisco-APIC-Installation-ACI-Upgrade-Downgrade-Guide/m-operations-allowed-during-mixed-versions-on-cisco-aci-switches.html

ACI Firmware Upgrade Types (SMU)

5.2(1) NEW

- Regular Upgrade
- Software Maintenance Upgrade (SMU)
- EPLD/FPGA Upgrade (Only Switches)

A patch for a specific defect

No need to upgrade the entire fabric. You can apply it only to APICs or affected switch nodes

1. SMU for all APICs
2. SMU for specific switches (through APIC)

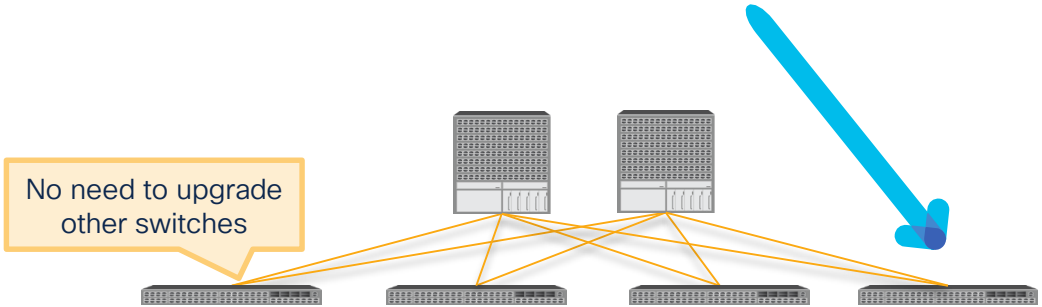

ACI Firmware Upgrade Types (EPLD/FPGA)

- Regular Upgrade
- Software Maintenance Upgrade (SMU)
- EPLD/FPGA Upgrade (Only Switches)

Hardware related firmware

Each ACI switch version has the desired EPLD/FPGA version.
Automatically upgraded via Regular Upgrade through APIC.
- No user configurations

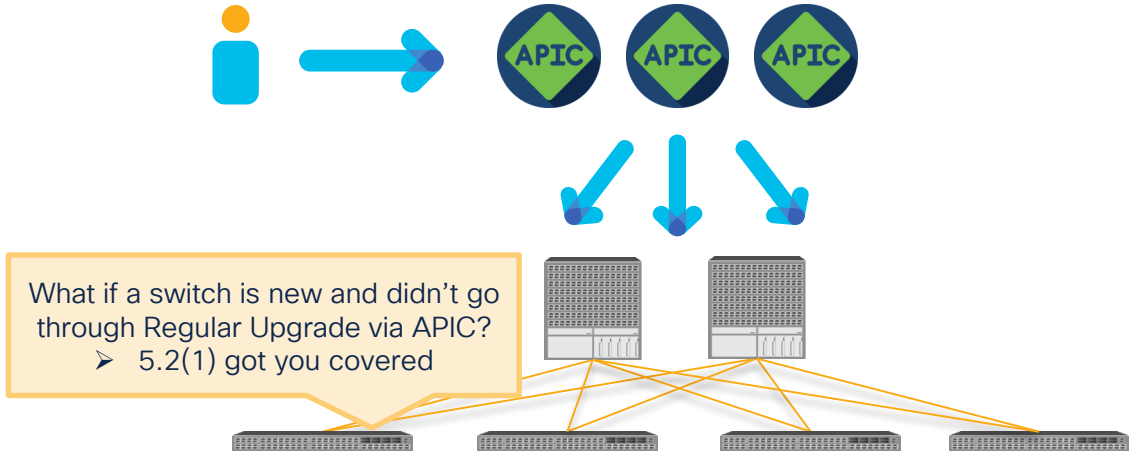

APIC Upgrade Architecture

Note: for 4.0 or newer APICs

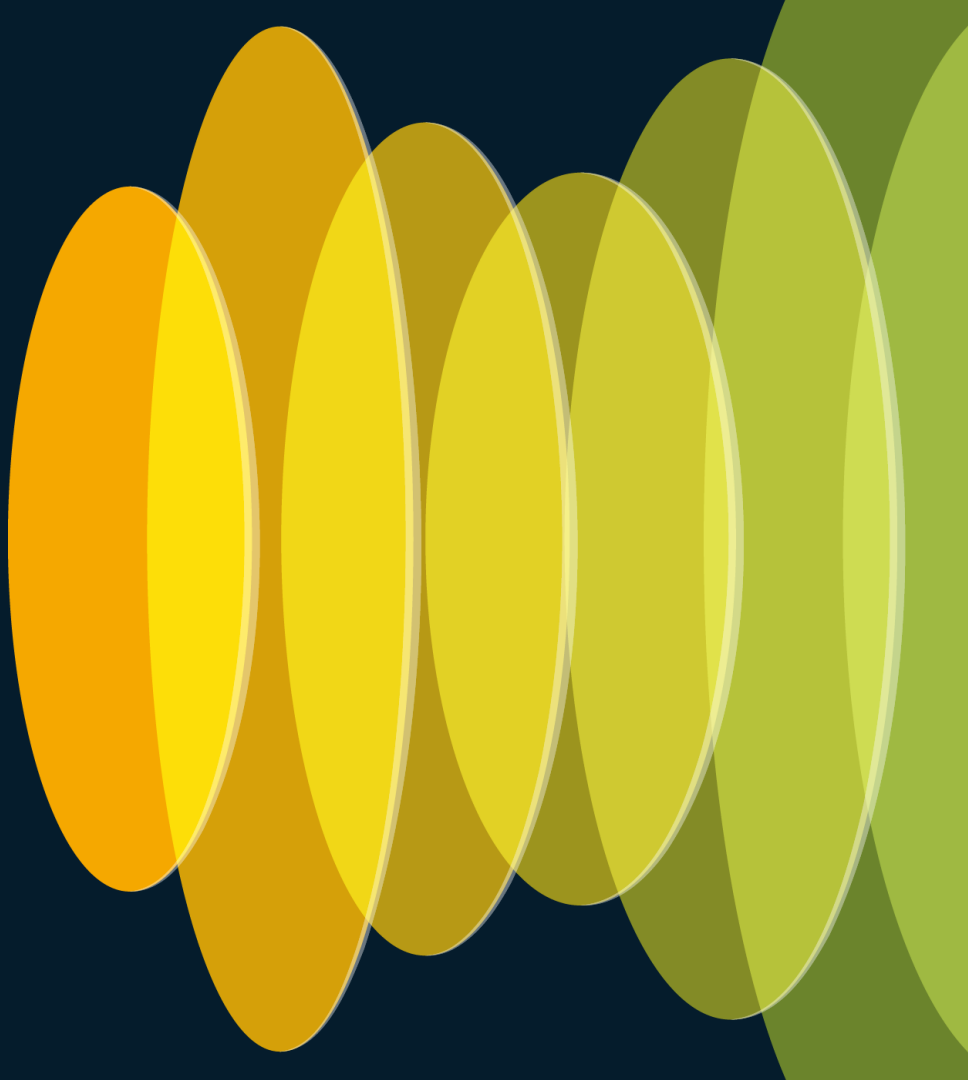

cisco Live!

APIC Upgrade Architecture

Image Upload

- A user uploads the APIC image on one of APICs
- After md5sum check, the image is copied to other APICs

Trigger

Install

Data Conversion & Reboot

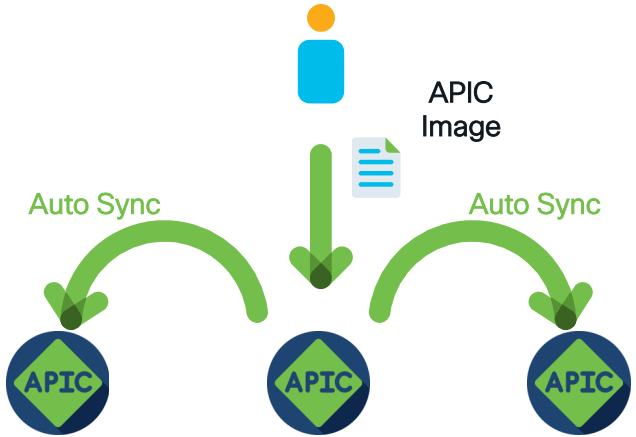

APIC Upgrade Architecture

Estimated Time
A few min.

Image Upload

Trigger
- Set the target version on all APICs
- APIC1 informs shards on all APICs of upgrades

Install

Data Conversion & Reboot

No disruptive operations from this point.
(details in later slides)

TIP Shard – user configurations and data spread across APICs
Replica – back up for each shard

APIC Upgrade Architecture

Image Upload

Trigger

Install

- Install APIC OS in a backup partition
- All APICs perform this in parallel

Data Conversion & Reboot

Estimated Time
A few min.

APIC Upgrade Architecture

Image Upload

Trigger

Install

Data Conversion & Reboot

- Convert user configurations and data to the target version format
- Conversion happens one APIC at a time

NEW

6.0(3) or newer requires **significantly less time** due to an internal enhancement in data conversion

Estimated Time

Depends on the size of data.

A fair estimation would be 40 min per APIC (potentially more or less)

Convert data starting from APIC 1, then reboot.

After reboot, APIC1's upgrade is considered completed.

Wait until lower numbered APICs finish data conversion and reboot.

ACI Switch Upgrade Architecture

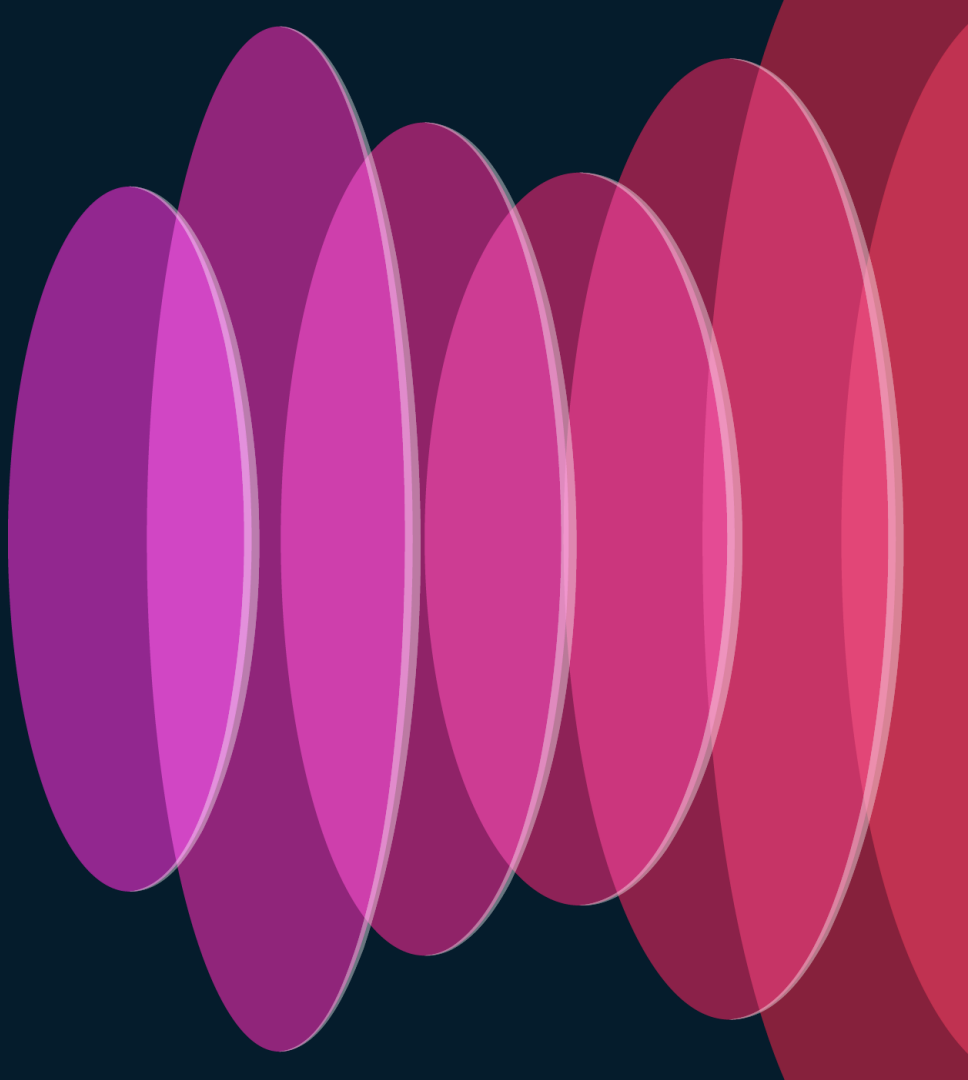

ACI Switch Upgrade Flow

Image Download

- The switch downloads the image from APIC
- The download is via infra TEP

Queuing

Preparation

Reboot

Boot Up

No Traffic Impact

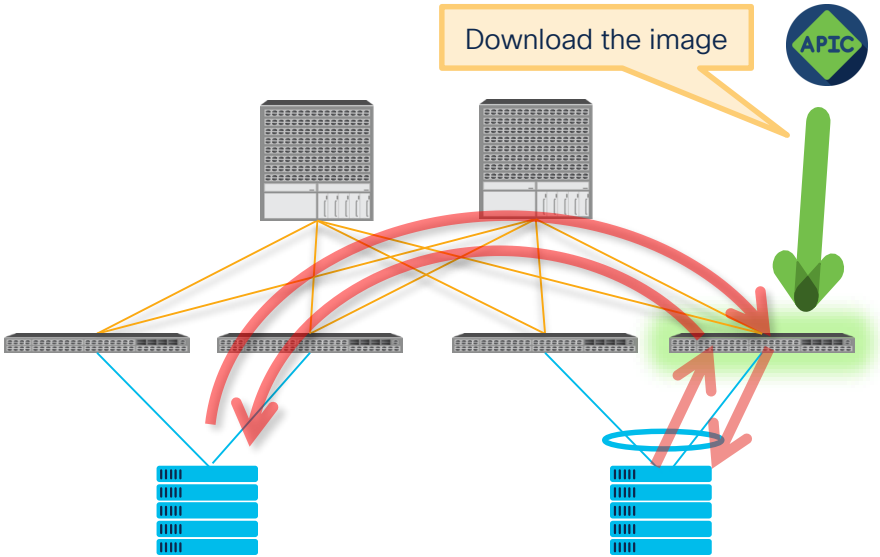

ACI Switch Upgrade Flow

Image Download

Queuing

- The switch receives approval from APIC
- Controls switches that are upgraded in parallel

Preparation

Since APIC 4.1(1)

- One leaf at a time in each vPC pair
- Not all spines in each pod if graceful option is used

Reboot

Boot Up

No Traffic Impact

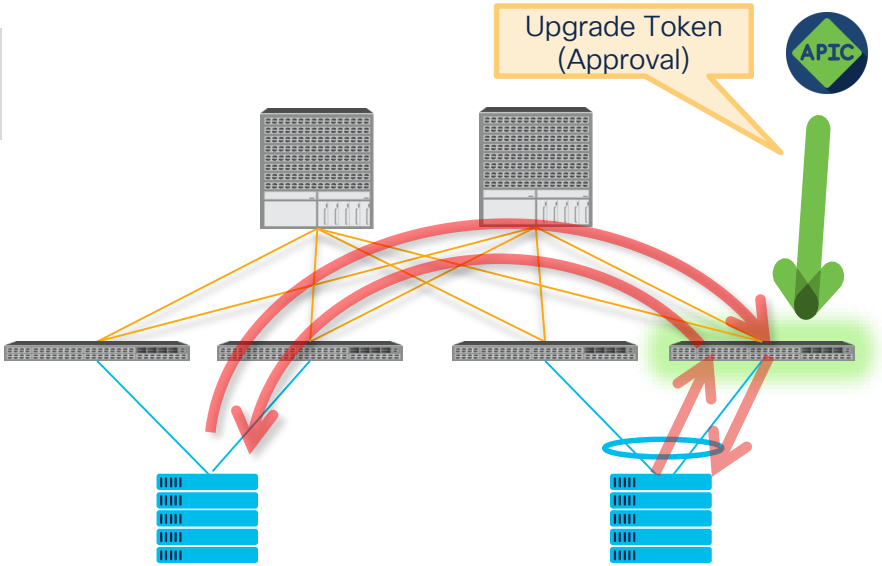

ACI Switch Upgrade Flow

Image Download

Queuing

Preparation

- The switch extracts the image.
- The switch sets the boot var and so on.

Reboot

Boot Up

No Traffic Impact

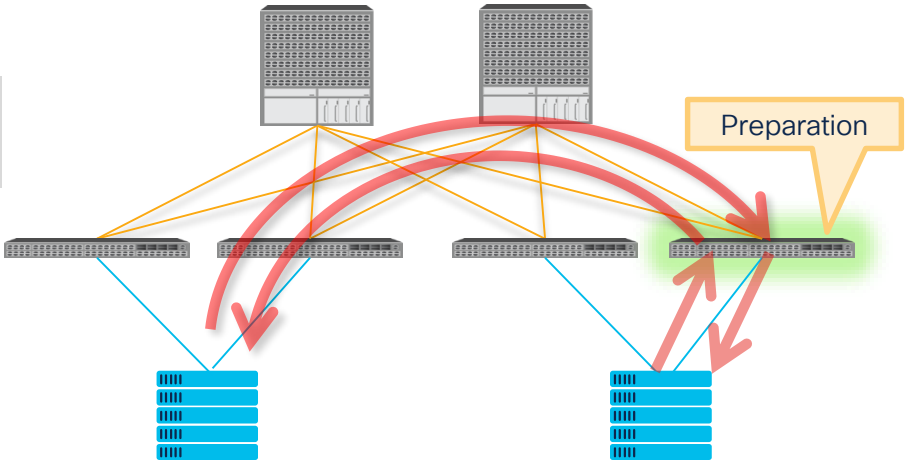

ACI Switch Upgrade Flow

Depends on other conditions such as:
- Link failure detection time on external devices
- Routing protocol and so on

< 100 msec Traffic Impact in the best case

Image Download

Queuing

Preparation

Reboot

- Wipe the config and reboot (i.e. clean reboot)
- Traffic failover relies on link failure

Boot Up

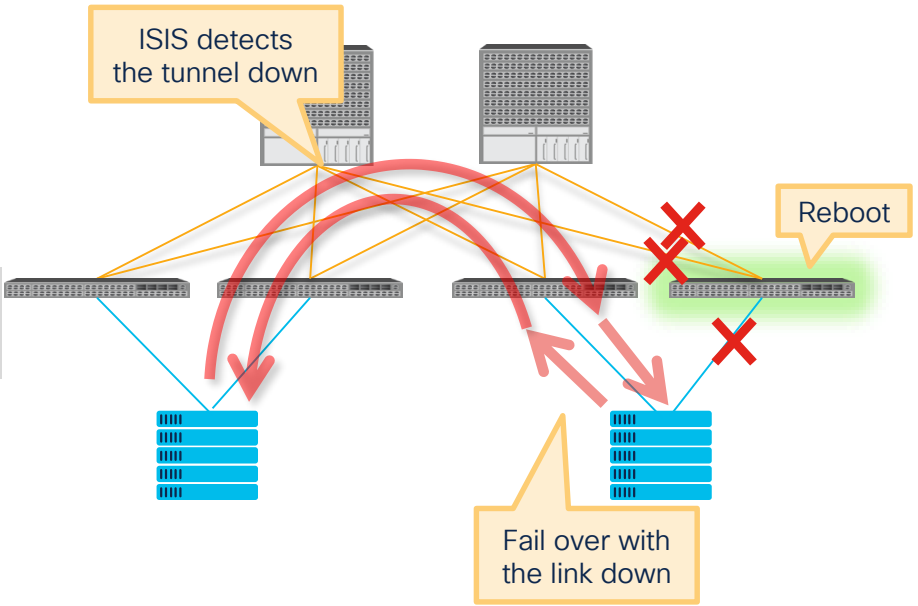

ACI Switch Upgrade Flow

Image Download

Queuing

Preparation

Reboot

Boot Up

- Various traffic flow optimizations
- (Continue to next slides)

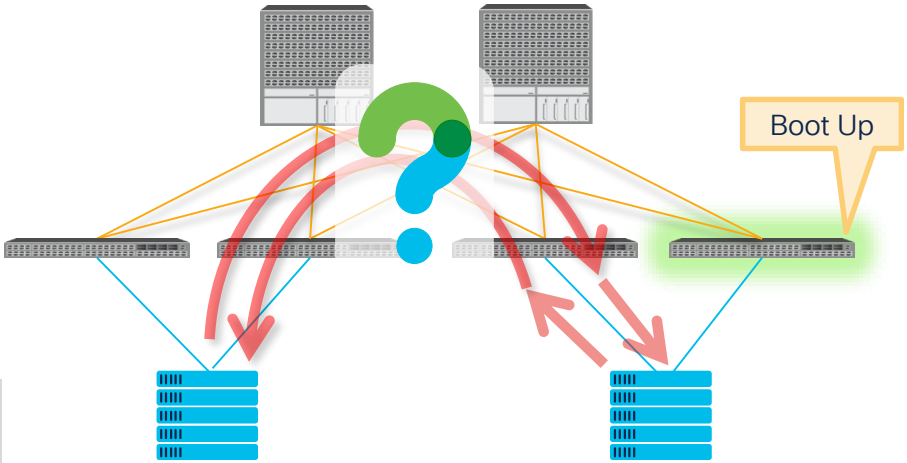

ACI Switch Upgrade Flow (Boot Up Sequence)

Boot Up · Various traffic flow optimizations

01
- Bring up fabric links
- Bring up APIC connected down links
- Admin down other down links

02

03

04

05

06

07

No Traffic Flow Change

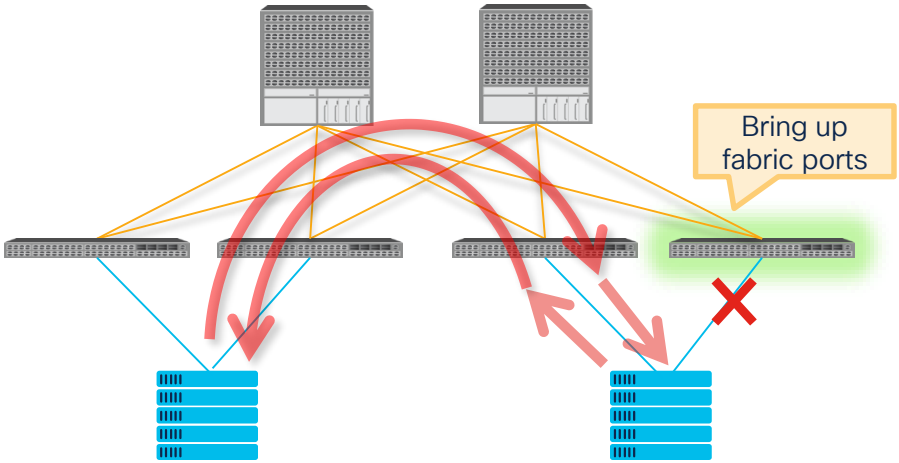

ACI Switch Upgrade Flow (Boot Up Sequence)

Boot Up · Various traffic flow optimizations

01
- Bring up fabric links
- Bring up APIC connected down links
- Admin down other down links

02
- An APIC discovers the switch via DHCP/LLDP
- The same TEP IP is assigned

03

04

05

06

07

No Traffic Flow Change

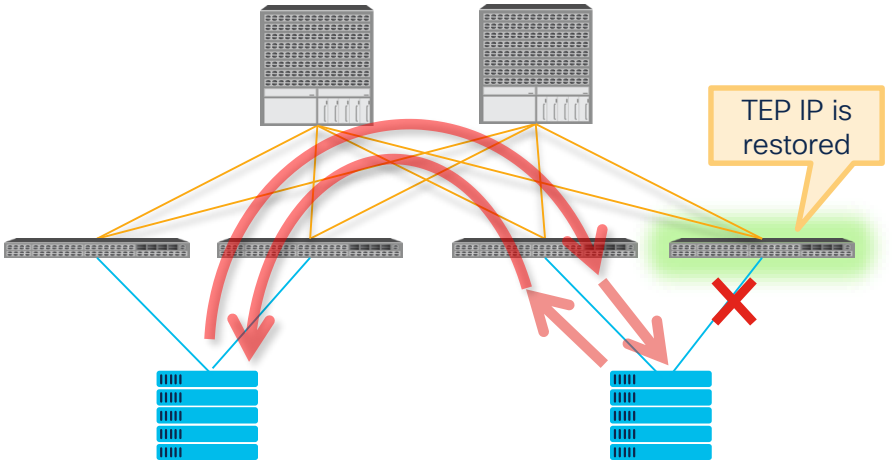

ACI Switch Upgrade Flow (Boot Up Sequence)

Boot Up · Various traffic flow optimizations

01
- Bring up fabric links
- Bring up APIC connected down links
- Admin down other down links

02
- An APIC discovers the switch via DHCP/LLDP
- The same TEP IP is assigned

03
- ISIS overload mode is activated
  - ✓ ISIS advertises the TEP IP with a large metric
  - ✓ ISIS does not advertise BD mcast groups to join

04

05

06

07

No Traffic Flow Change

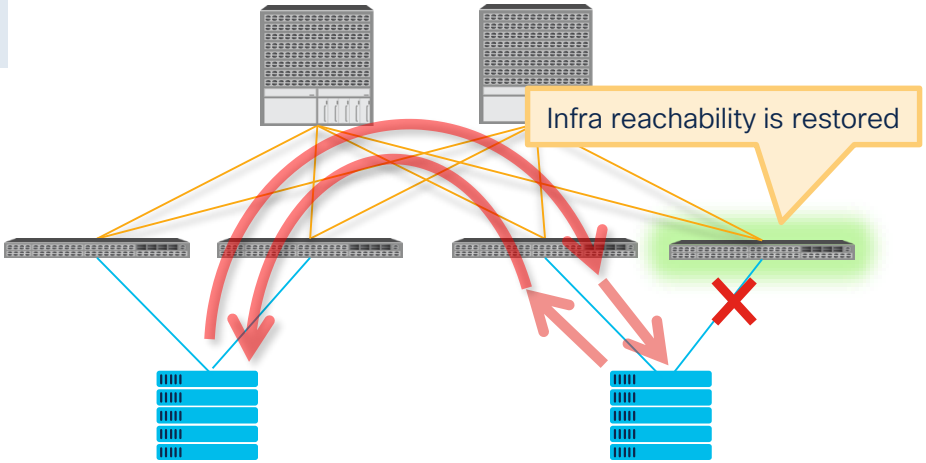

ACI Switch Upgrade Flow (Boot Up Sequence)

Boot Up · Various traffic flow optimizations

01
- Bring up fabric links
- Bring up APIC connected down links
- Admin down other down links

02
- An APIC discovers the switch via DHCP/LLDP
- The same TEP IP is assigned

03
- ISIS overload mode is activated
  - ✓ ISIS advertises the TEP IP with a large metric
  - ✓ ISIS does not advertise BD mcast groups to join

04
- Starts downloading configurations from an APIC

05

06

07

No Traffic Flow Change

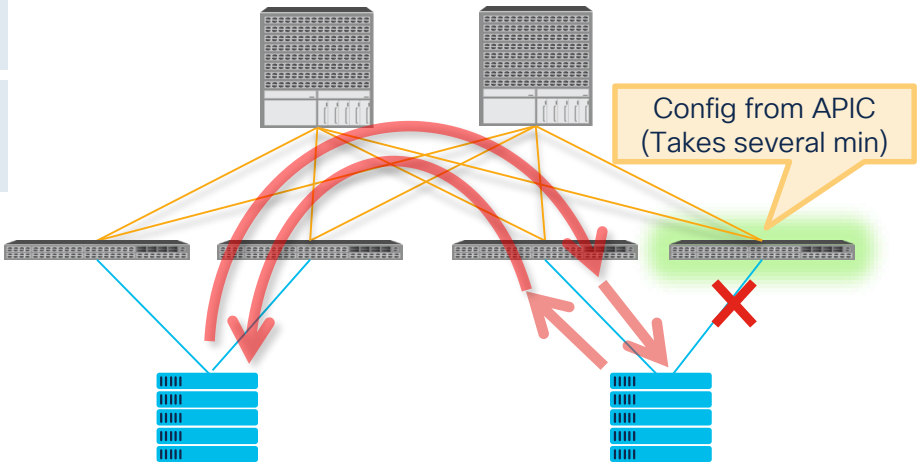

ACI Switch Upgrade Flow (Boot Up Sequence)

Boot Up · Various traffic flow optimizations

01
- Bring up fabric links
- Bring up APIC connected down links
- Admin down other down links

02
- An APIC discovers the switch via DHCP/LLDP
- The same TEP IP is assigned

03
- ISIS overload mode is activated
  - ✓ ISIS advertises the TEP IP with a large metric
  - ✓ ISIS does not advertise BD mcast groups to join

04
- Starts downloading configurations from an APIC

05
- ISIS multicast overload mode completes (i.e. flood)
- vPC peer is established at the same time

06

ISIS multicast overload timer
- Leaf nodes – Fixed 1min
- Spine nodes – When FTAG tree is created (Fixed 1 min prior to Switch 14.2(1))

07

No Traffic Flow Change

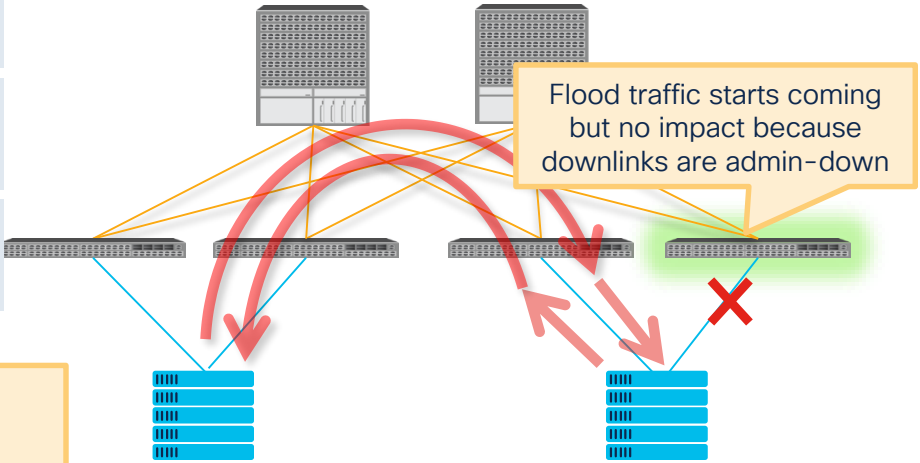

Spine ISIS multicast overload timer (CSCvp79708)

Why not a fixed 1 min?

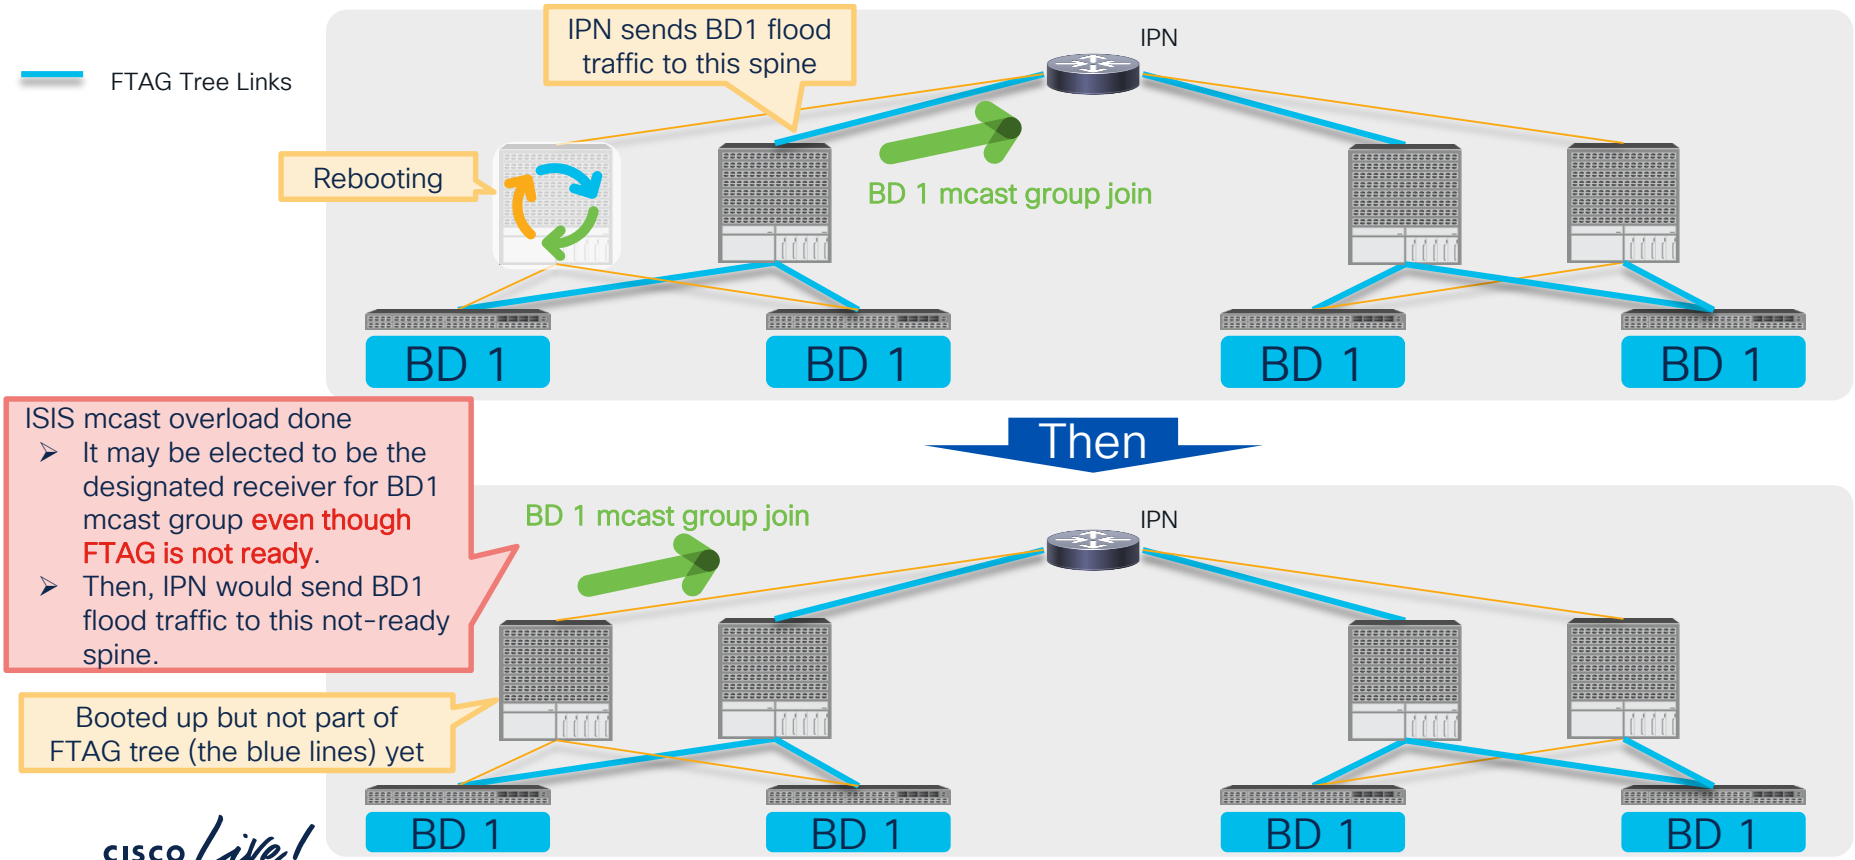

ACI Switch Upgrade Flow (Boot Up Sequence)

Boot Up · Various traffic flow optimizations

01
- Bring up fabric links
- Bring up APIC connected down links
- Admin down other down links

02
- An APIC discovers the switch via DHCP/LLDP
- The same TEP IP is assigned

03
- ISIS overload mode is activated
  - ✓ ISIS advertises the TEP IP with a large metric
  - ✓ ISIS does not advertise BD mcast groups to join

04
- Starts downloading configurations from an APIC

05
- ISIS multicast overload mode completes (i.e. flood)
- vPC peer is established at the same time

06
- Full configuration has been downloaded
  - ✓ Bring up access links (downlinks)
  - ✓ and vPC ports after vPC restore delay timer expires

07

Ready to receive traffic
- VLANs are deployed
  - For VMM, depends on Resolution Immediacy
- Contracts are deployed
  - Depends on Deployment Immediacy
- Spine-Proxy is ready
- Flood handling (FTAG) is ready

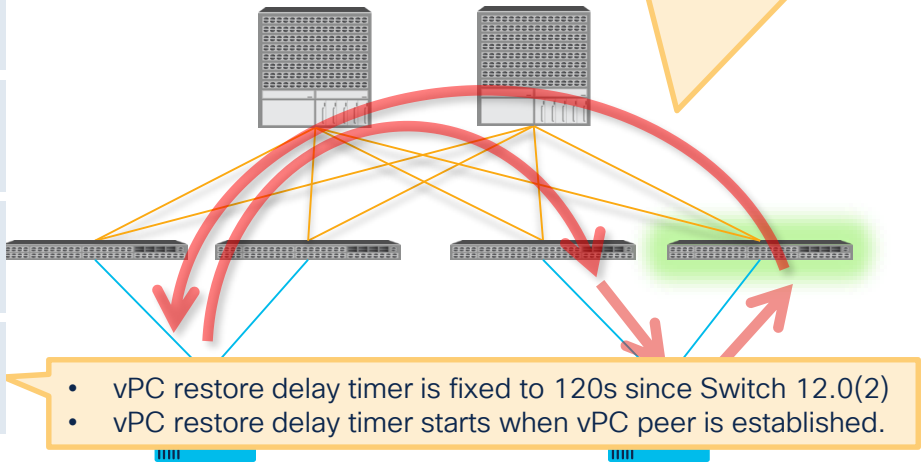

- vPC restore delay timer is fixed to 120s since Switch 12.0(2)
- vPC restore delay timer starts when vPC peer is established.

ACI Switch Upgrade Flow (Boot Up Sequence)

Boot Up · Various traffic flow optimizations

01
- Bring up fabric links
- Bring up APIC connected down links
- Admin down other down links

02
- An APIC discovers the switch via DHCP/LLDP
- The same TEP IP is assigned

03
- ISIS overload mode is activated
  - ✓ ISIS advertises the TEP IP with a large metric
  - ✓ ISIS does not advertise BD mcast groups to join

04
- Starts downloading configurations from an APIC

05
- ISIS multicast overload mode completes (i.e. flood)
- vPC peer is established at the same time

06
- Full configuration has been downloaded
  - ✓ Bring up access links (downlinks)
  - ✓ and vPC ports after vPC restore delay timer expires

07
- ISIS unicast overload mode completes
  - ✓ The TEP IP is advertised with a normal metric

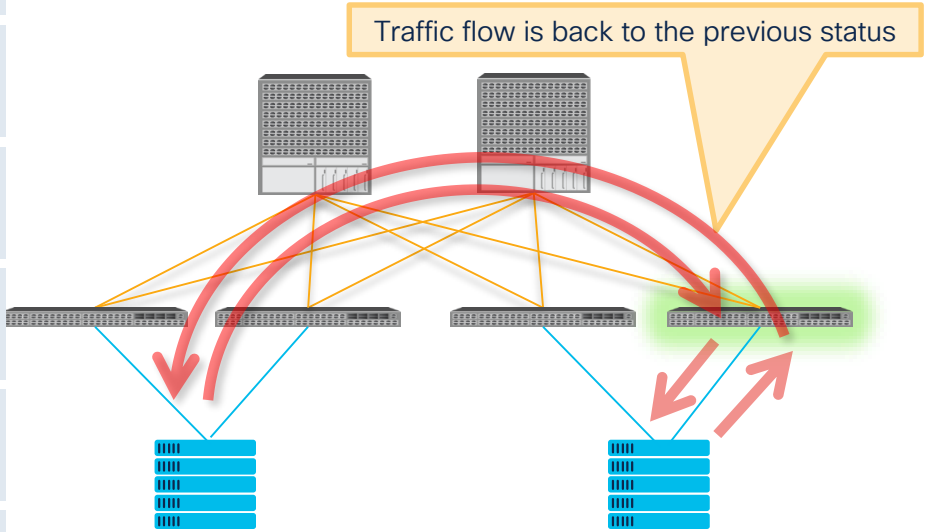

ISIS unicast overload timer – 10 min fixed for all nodes

ACI Switch Upgrade with Graceful Option

(a.k.a. Graceful Upgrade)

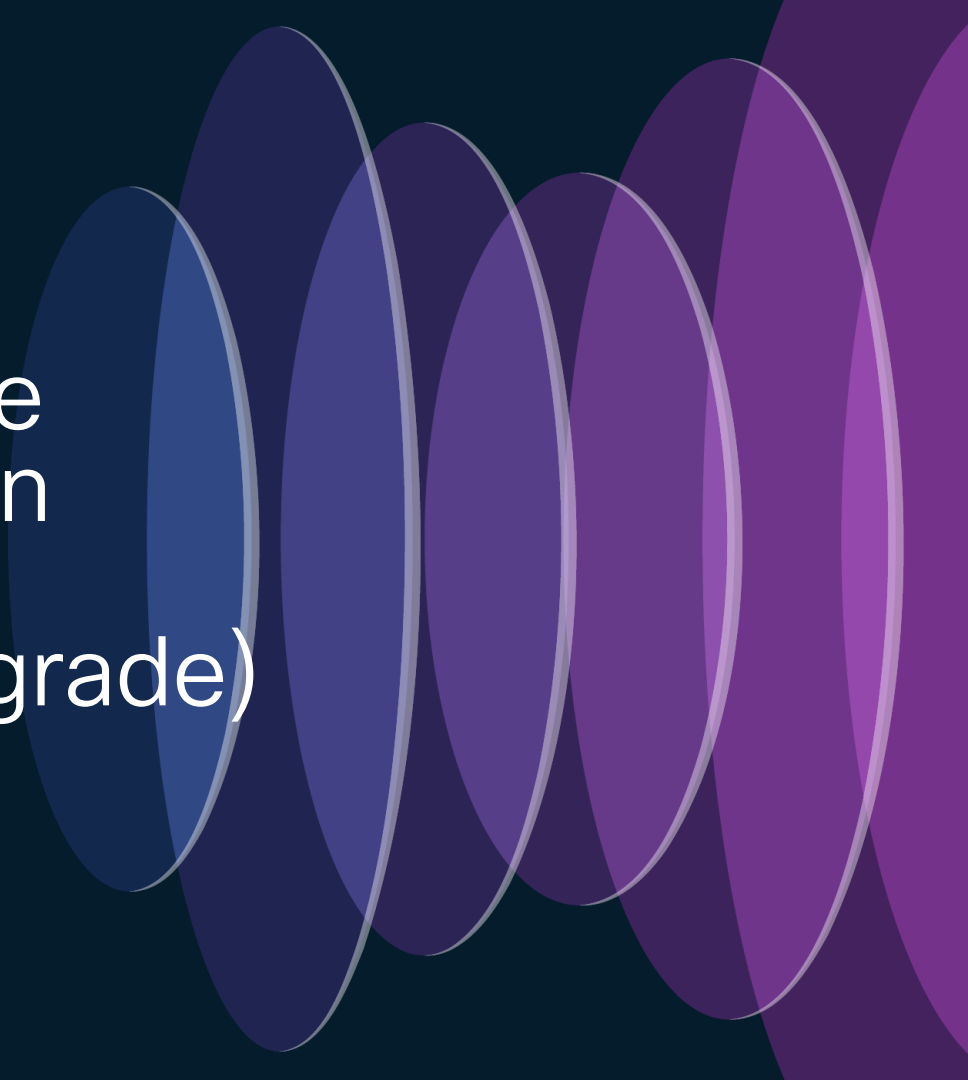

ACI Switch Upgrade with graceful option

Image Download

Scheduler

Preparation

Reboot

- Wipe the config and reboot (i.e. clean reboot)
- ~~Traffic failover relies on link failure~~

Boot Up

The rest is the same as without graceful option.

Graceful Option is to gracefully isolate the switch before the switch goes down for the upgrade

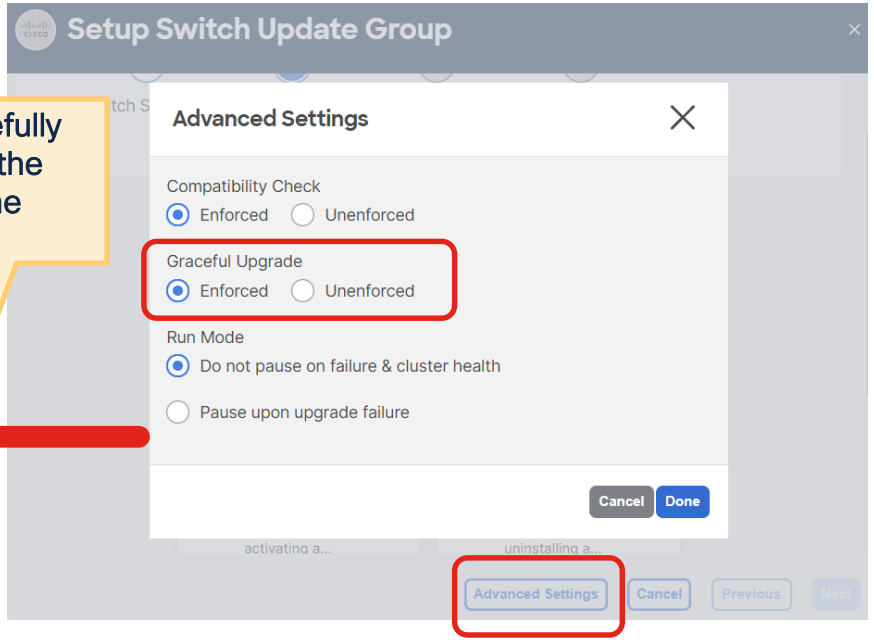

ACI Switch Upgrade with graceful option

Image Download

Scheduler

Preparation

Reboot

- Wipe the config and reboot (i.e. clean reboot)
- ~~Traffic failover relies on link failure~~

Boot Up

The rest is the same as without graceful option.

Graceful Option is to gracefully isolate the switch before the switch goes down for the upgrade

Older APIC GUI

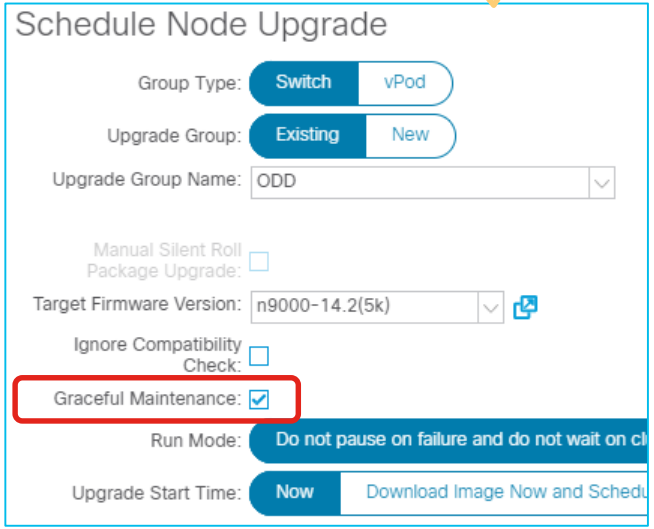

Enhanced reboot sequence with graceful option

- Graceful option **disabled**

Reboot

1. Wipe the config and reboot (i.e. clean reboot)
2. Traffic failover relies on user configured link failure mechanism

- Graceful option **enabled**

Reboot

1. Put the switch into MMode (Maintenance Mode)
   1. ISIS Overload Mode enabled
   2. Graceful Shutdown on Routing Protocols
      - ✓ Leaf – BGP, EIGRP, OSPF for L3Out
      - ✓ Spine – BGP, OSPF for IPN, GOLF
   3. vPC informs its peer that this switch is going down
   4. LACP sends PDUs with aggregation bit zero (starting from 3.1(2))
      - ➢ External devices can exclude the link from the port-channel before the link physically goes down.
   5. Shutdown front panel ports
      - ✓ Leaf – all down links including APIC connected links
2. Wipe the config and reboot (i.e. clean reboot)

Traffic Disruption without Graceful Upgrade

- OSPF DR reboot example

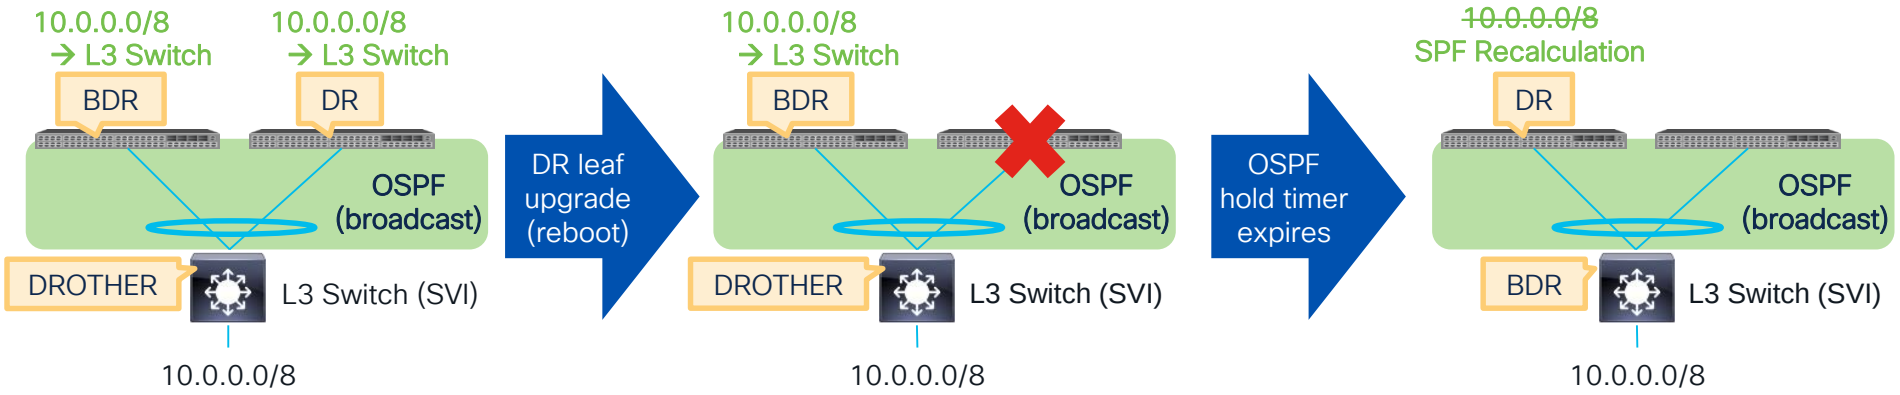

No Traffic Impact

< 100 msec Traffic Impact

Due to failover with link failure

A few seconds loss

Until the external router re-sends OSPF LSA for 10.0.0.0/8

With Graceful Upgrade

- OSPF DR reboot example

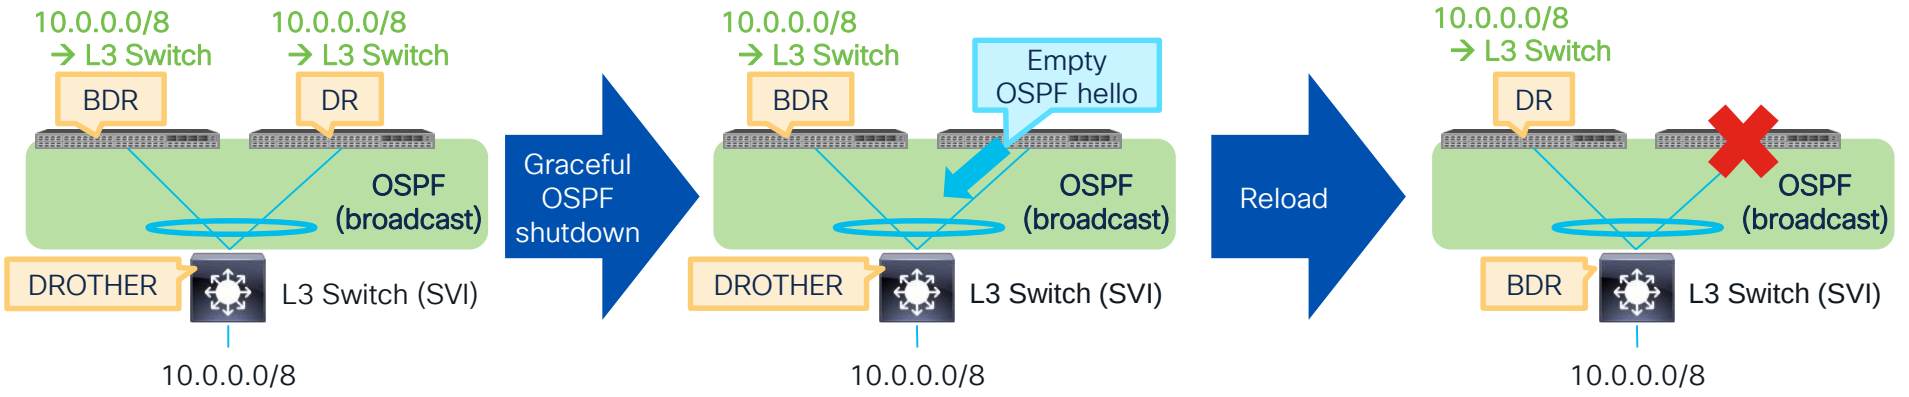

No Traffic Impact

No Traffic Impact

OSPF will _not_ lose existing routes due to SPF recalculation

< 100 msec Traffic Impact

Due to failover with link failure

GIR and Graceful Upgrade in ACI

Both GIR (Graceful Insertion and Removal) and Graceful Upgrade put the switch in MMode (Maintenance Mode) to isolate the switch from the fabric.

However, the use case for these two features are completely different.

GIR (Graceful Insertion and Removal)

Use Case:
- To isolate a switch for further debugging
- To quickly restore service by isolating a malfunctioning switch

Difference:
- It is not supported to upgrade a switch in MMode via GIR

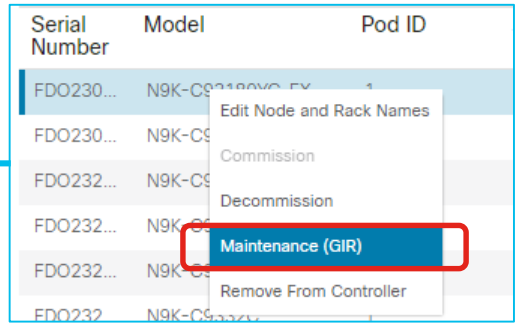

An upgrade with the graceful option

Use Case:
- To upgrade a switch after isolating the switch

Difference:
- The switch will communicate to APIC and perform an upgrade immediately after the switch was put into MMode.

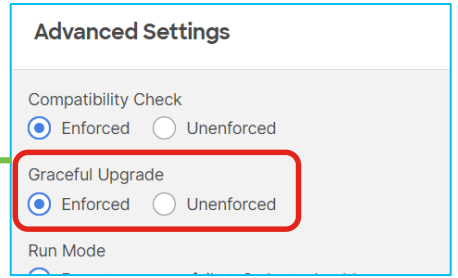

Auto Firmware Update

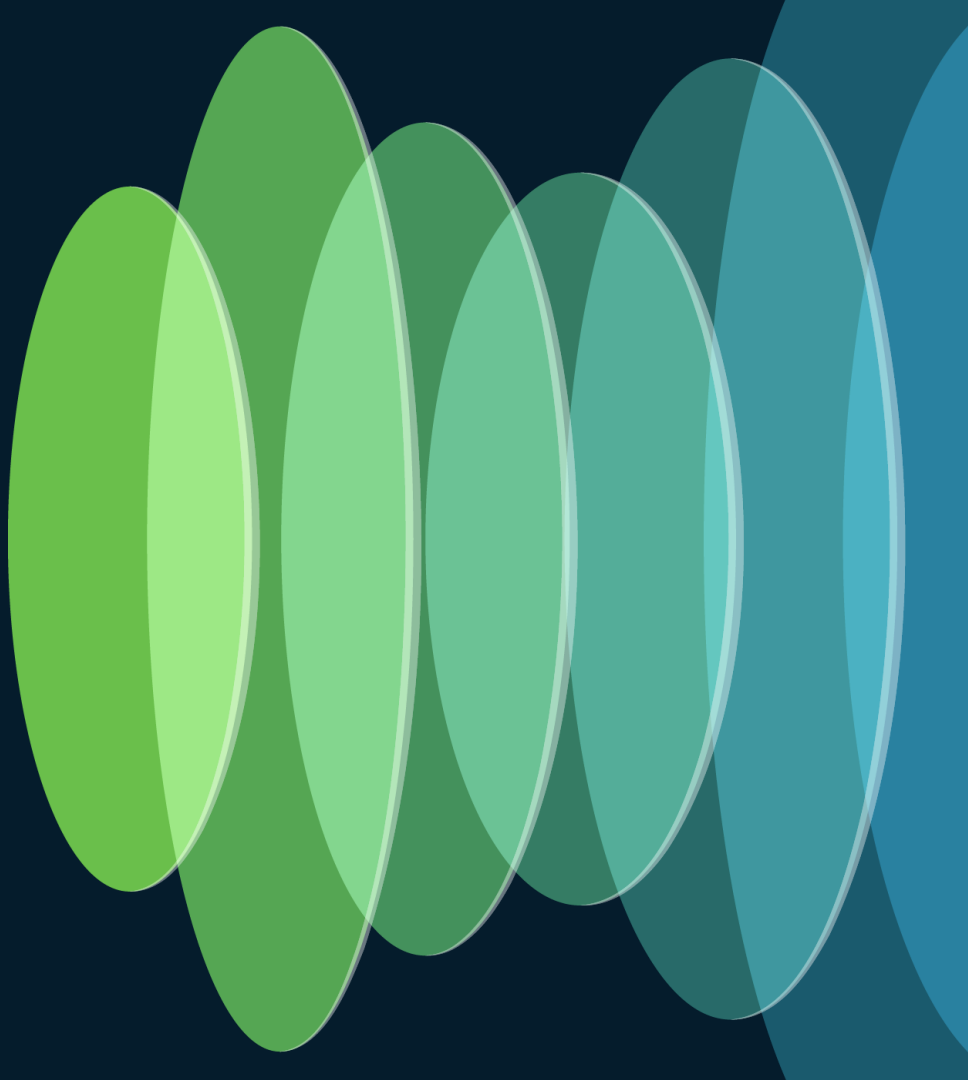

Auto Firmware Update for APIC

6.0(2) NEW

Use Case 1: APIC Replacement

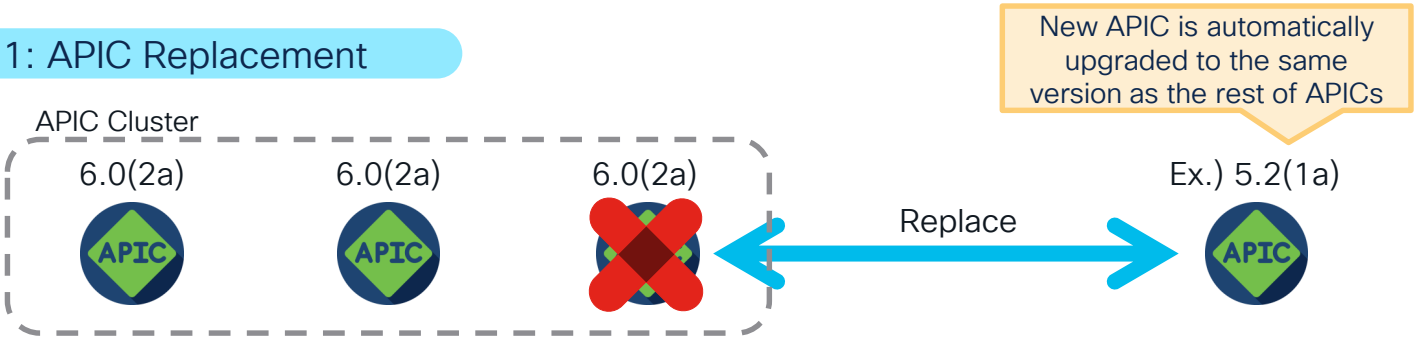

Use Case 2: Cluster Expansion

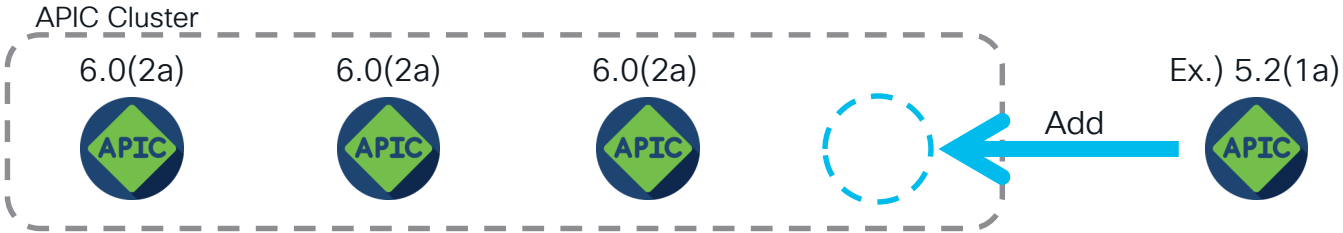

Auto Firmware Update for Switches

Enforcing Version Consistency

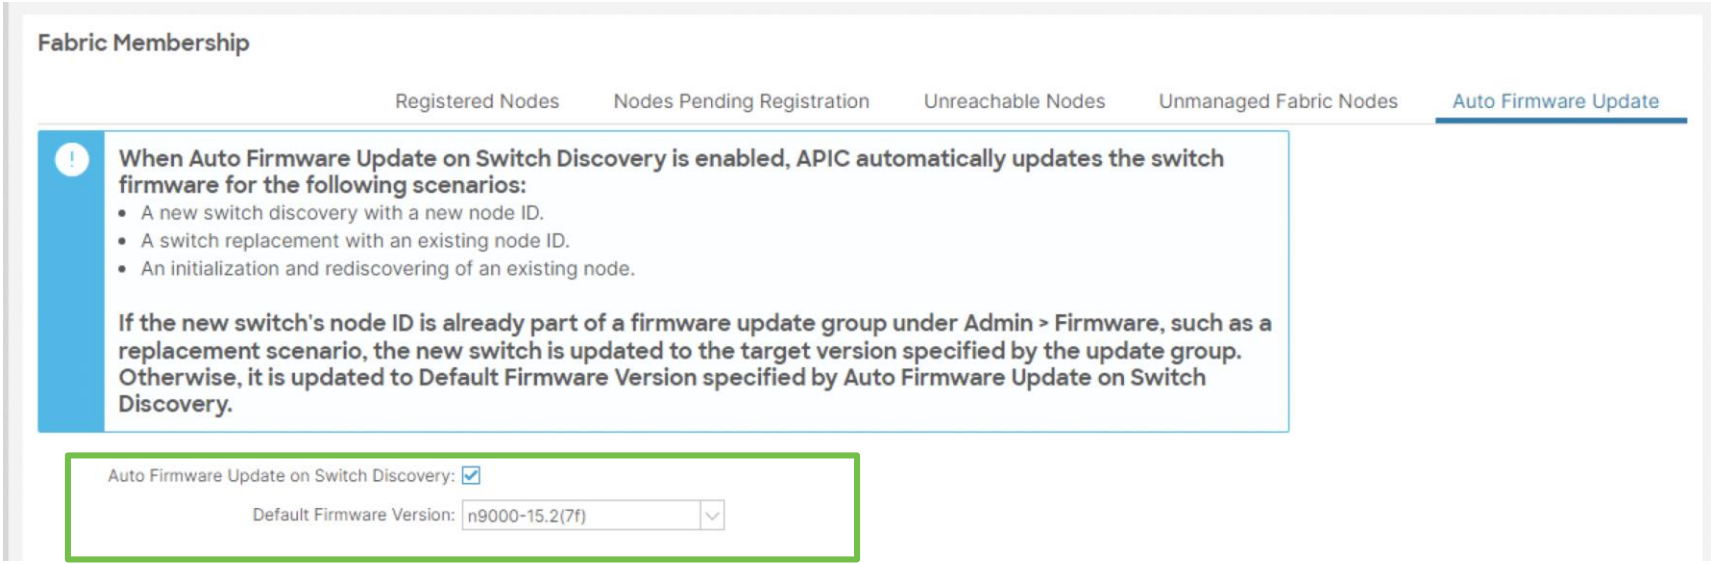

Fabric > Inventory > Fabric Membership > Auto Firmware Update >=5.1(1)

Admin > Firmware > Infrastructure > Nodes > Enforce Bootscript Version Validation < 5.1(1)

32-bit/64-bit Switch Images

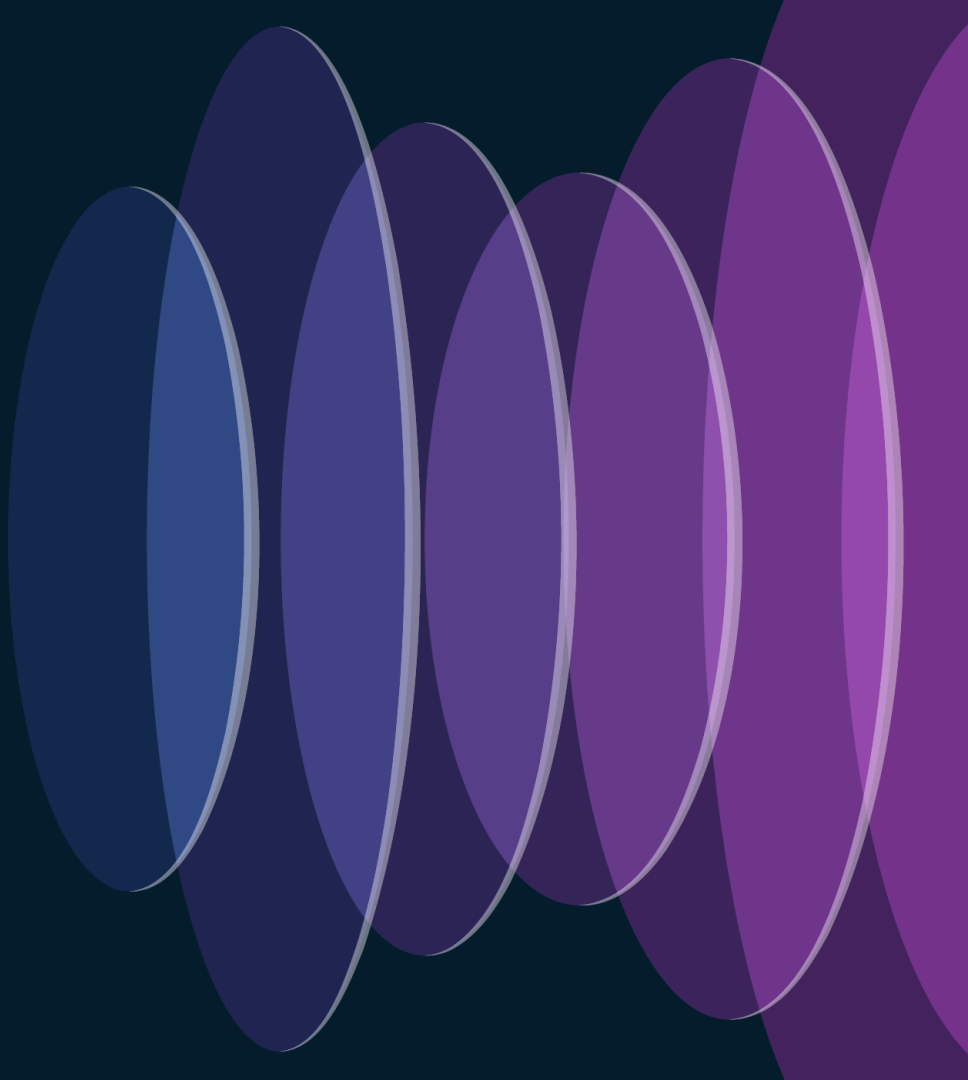

32-bit/64-bit switch images

6.0(2) NEW

https://software.cisco.com

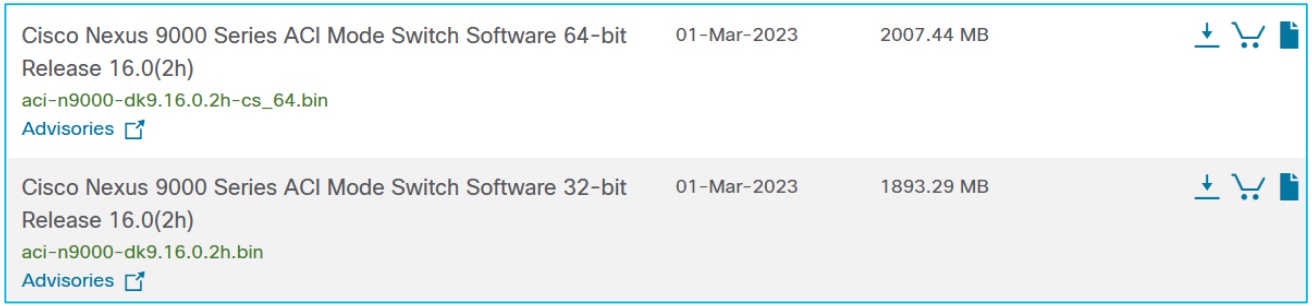

Why?

Scalability

To utilize the most of what the switch hardware has to offer.

Example: Forwarding Scale Profile such as Max LPM profile requires 32 GB memory / 64-bit image.

Which Image?

APIC picks the right one for you

6.0(3) or newer:
  Total memory is 24GB or less: 32 bit
  Otherwise: 64 bit
6.0(2):
  Fixed per switch model

Upgrade Procedure

When do you upload 64-bit switch images to APICs?

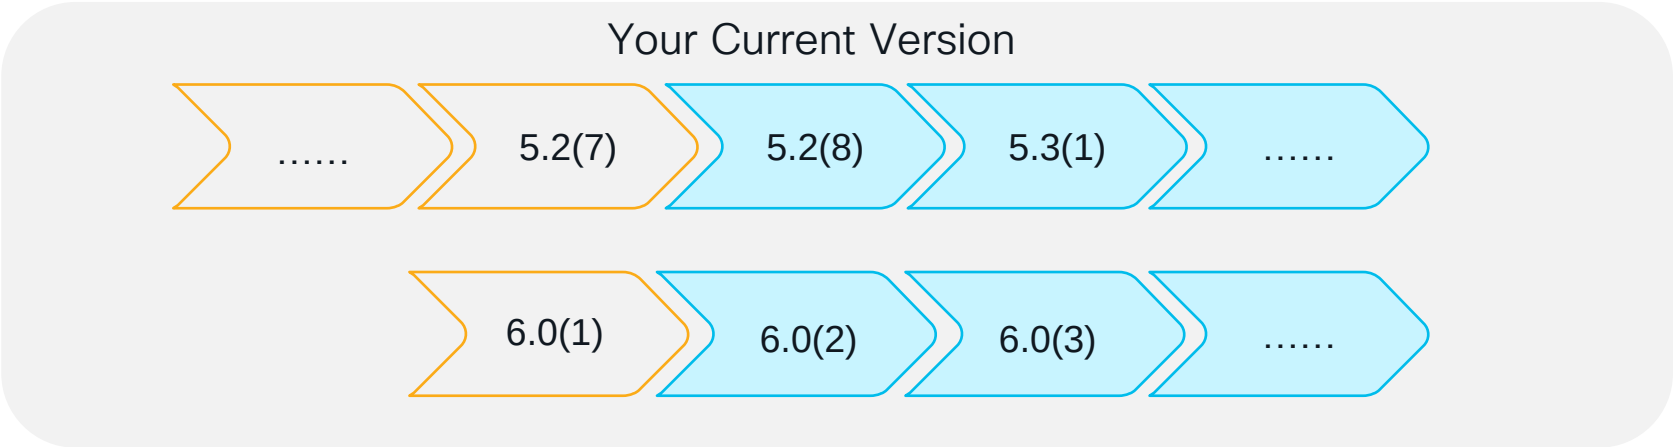

Do not upload the 64-bit image until APICs are upgraded to the blue versions

Upload the 64-bit image before APICs are upgraded

Upgrade Procedure

Always upload both 32-bit and 64-bit switch images on APICs

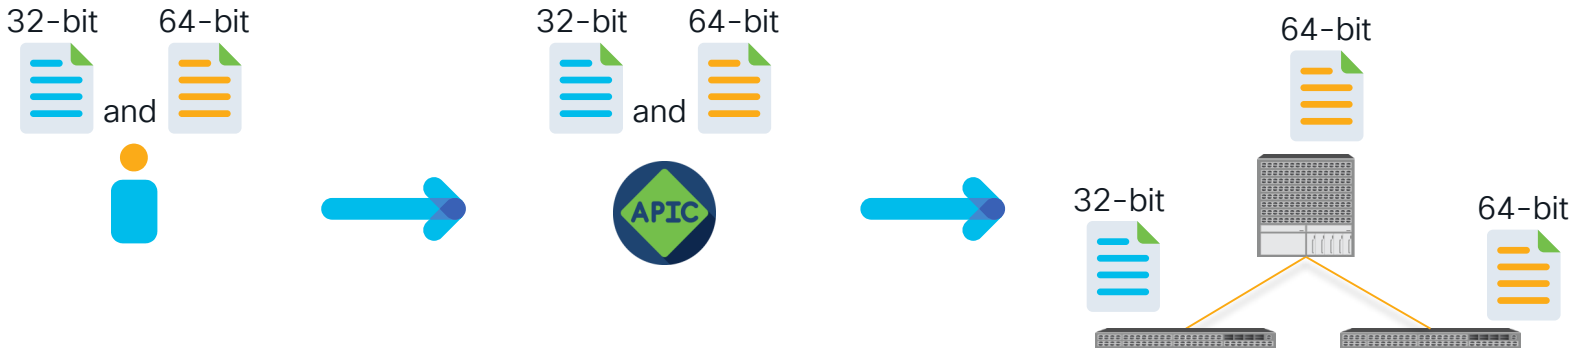

Upgrade and Reload

There may be 2 reloads when a switch uses a 64-bit image for the first time

1 APIC picks 32 or 64

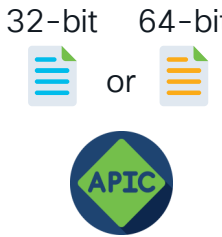

APIC checks its internal catalog for the memory options for each switch model.

If the switch model can have 24GB or less memory, APIC loads 32-bit image first.

2 Switch Upgrade

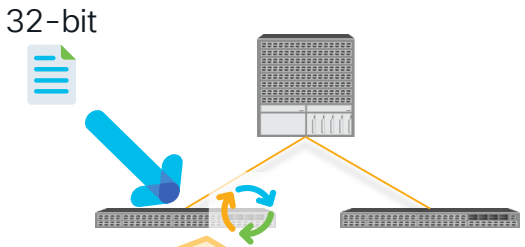

Switches upgrade with whichever the image passed by APICs

3 Switch Upgrade Again

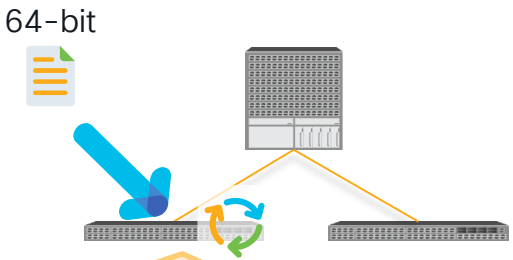

Automatically upgrade again with 64-bit with the conditions below:

- The switch has 32GB or more memory
- The switch is running a 32-bit image of 6.0(3) or later

If the switch is already running 64-bit, the next upgrade will use 64-bit directly. No more 2 reloads.

Memory Options

Switch models that can have 24GB or less memory

Product ID	Possible Memory Size
N9K-C9332C	16GB
N9K-C93180LC-EX	24GB
N9K-C93180YC-EX	24GB
N9K-C93108TC-EX	24GB
N9K-C93180YC-FX	24GB or 32GB
N9K-C93108TC-FX	24GB
N9K-C9348GC-FXP	24GB

Product ID	Possible Memory Size
N9K-C9336C-FX2	24GB
N9K-C9336C-FX2-E	24GB
N9K-C93216TC-FX2	16GB or 24GB
N9K-C93240YC-FX2	16GB or 24GB
N9K-C93360YC-FX2	16GB or 24GB
N9K-C93180YC-FX3	16GB or 32GB
N9K-C93180YC-FX3H	16GB or 32GB
N9K-C93108TC-FX3H	16GB or 32GB
N9K-C93108TC-FX3P	16GB or 32GB
N9K-C9364C-GX	16GB or 32GB

These switches may reload twice

Note:
This list is based on the internal catalog on APIC running 6.0(5). Subject to change.

Upgrade Enhancements

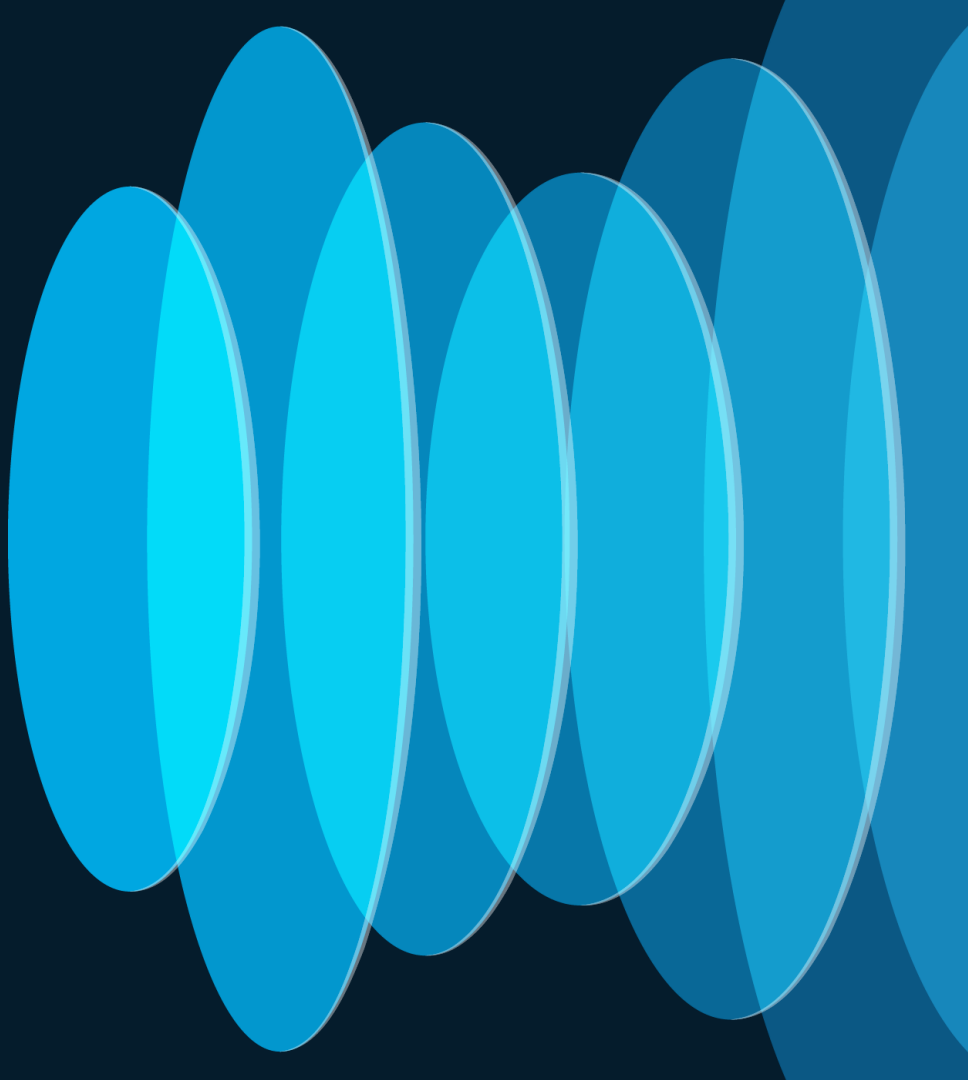

ACI Upgrade Enhancement Quick Summary

	Supported APIC versions →	4.1(1)	4.2(1)	4.2(5)	5.2(1)	5.2(3)	6.0(2)	6.0(3)	Switch version requirements
Upgrade Time Optimization	Switch Image Pre-download	✓	✓	✓	✓	✓	✓	✓	14.1(1) or later
	Multi-Pod Parallel Switch Upgrade			✓	✓	✓	✓	✓	No requirements
	Unlimited Parallel Switch Upgrade By Default			✓	✓	✓	✓	✓	No requirements
	Faster APIC Data Conversion							✓	N/A
Visibility	APIC Detailed Install Stage			✓	✓	✓	✓	✓	N/A
	Switch Image Download Progress			✓	✓	✓	✓	✓	14.5(1) or later
Operation Optimization	Built-in Pre-Upgrade Validation		✓	✓	✓	✓	✓	✓	No requirements
	Pre-Upgrade Validator App				✓	✓	✓	✓	No requirements
	SMU Support				✓	✓	✓	✓	15.2(1) or later
	Auto EPLD/FPGA upgrade				✓	✓	✓	✓	15.2(1) or later
	NXOS to ACI auto conversion via POAP					✓	✓	✓	15.2(3) or later
	Auto Firmware Update for APIC						✓	✓	N/A
	Auto Firmware Update for switches	✓	✓	✓	✓	✓	✓	✓	No requirements

Upgrade Time Reduction

Switch Image Download from APIC to switches

Upgrade multiple pods/switches in parallel

Switch Image Pre-Download with a scheduler

4.1(1)

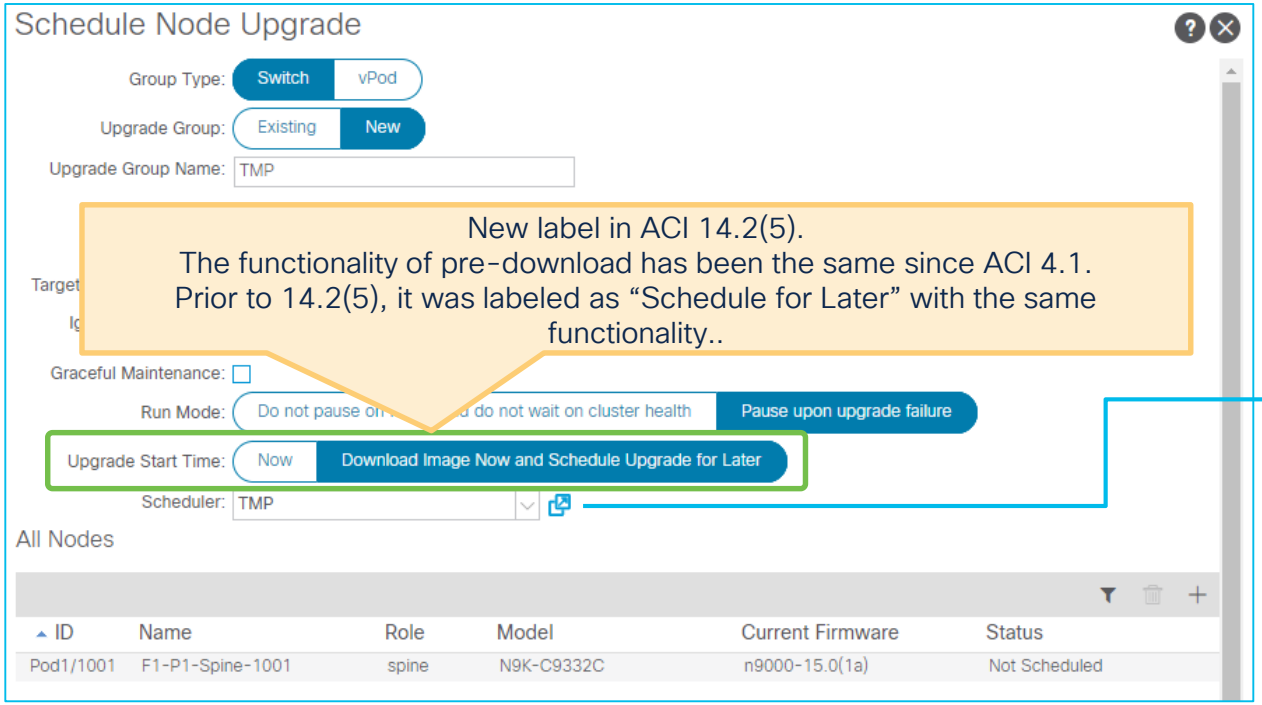

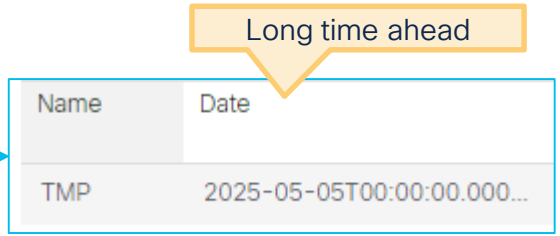

1. Schedule for a long time ahead just to trigger pre-download of a switch image.
2. During the actual maintenance window, come back to this same window (maintenance group) and select “Now” to trigger the upgrade on demand. Switches don’t need to re-download images and can proceed with the upgrade immediately.

Switch Image Download Progress

4.2(5) NEW

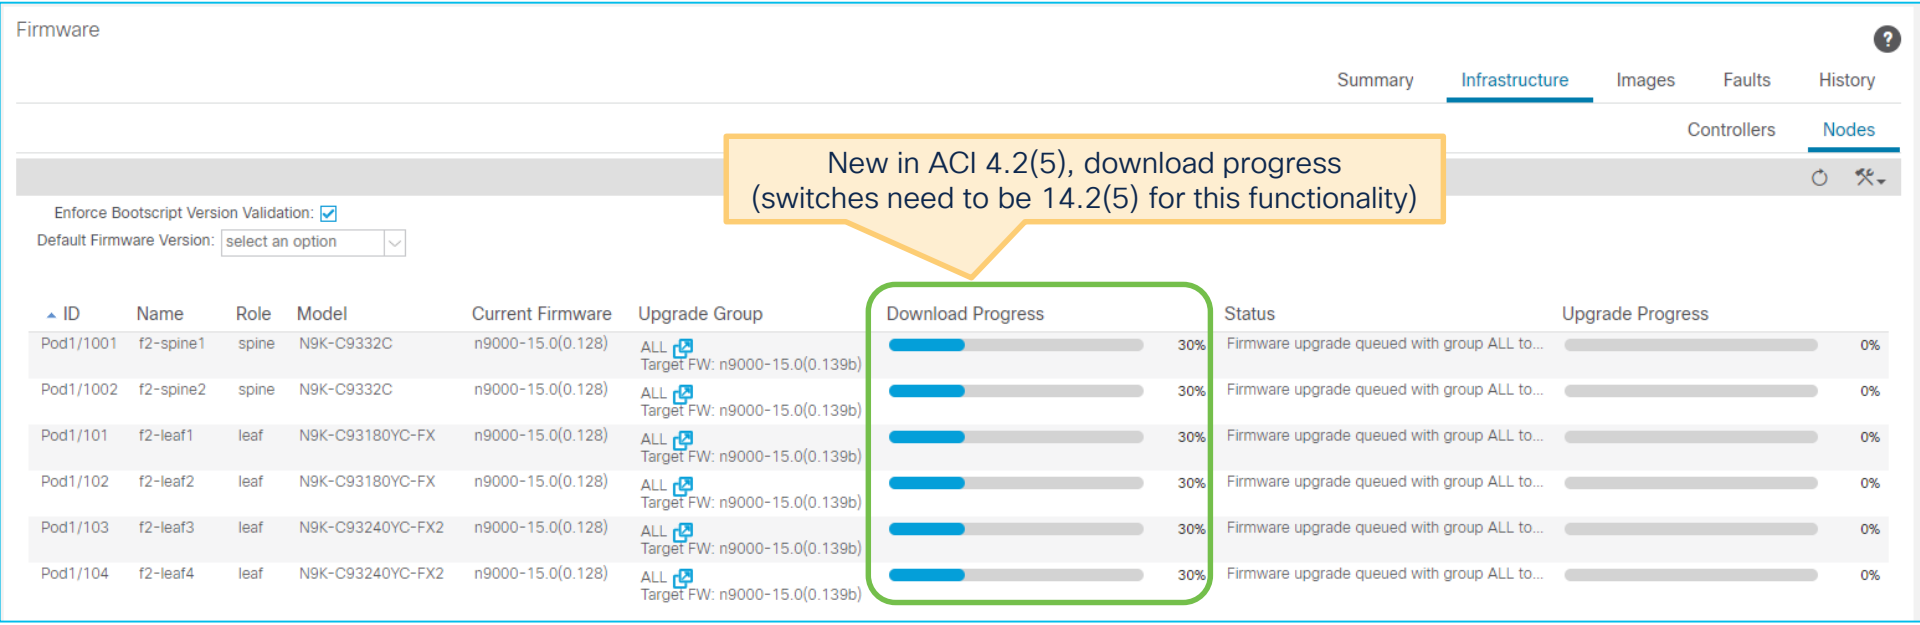

- All switches (regardless of pods or vPC) in the update group download the switch image from APICs in parallel. During this period, the Upgrade Progress remains 0 %.
- With the new Download Progress bar, users can see whether switches finished the download and ready to upgrade.
- If it was triggered with a scheduler, all switches wait after they completed their download.
- If it was triggered with “Upgrade Now”, each switch proceed with the upgrade as soon as it has completed its download.

Switch Image Pre-Download (built-in)

5.1(1) NEW

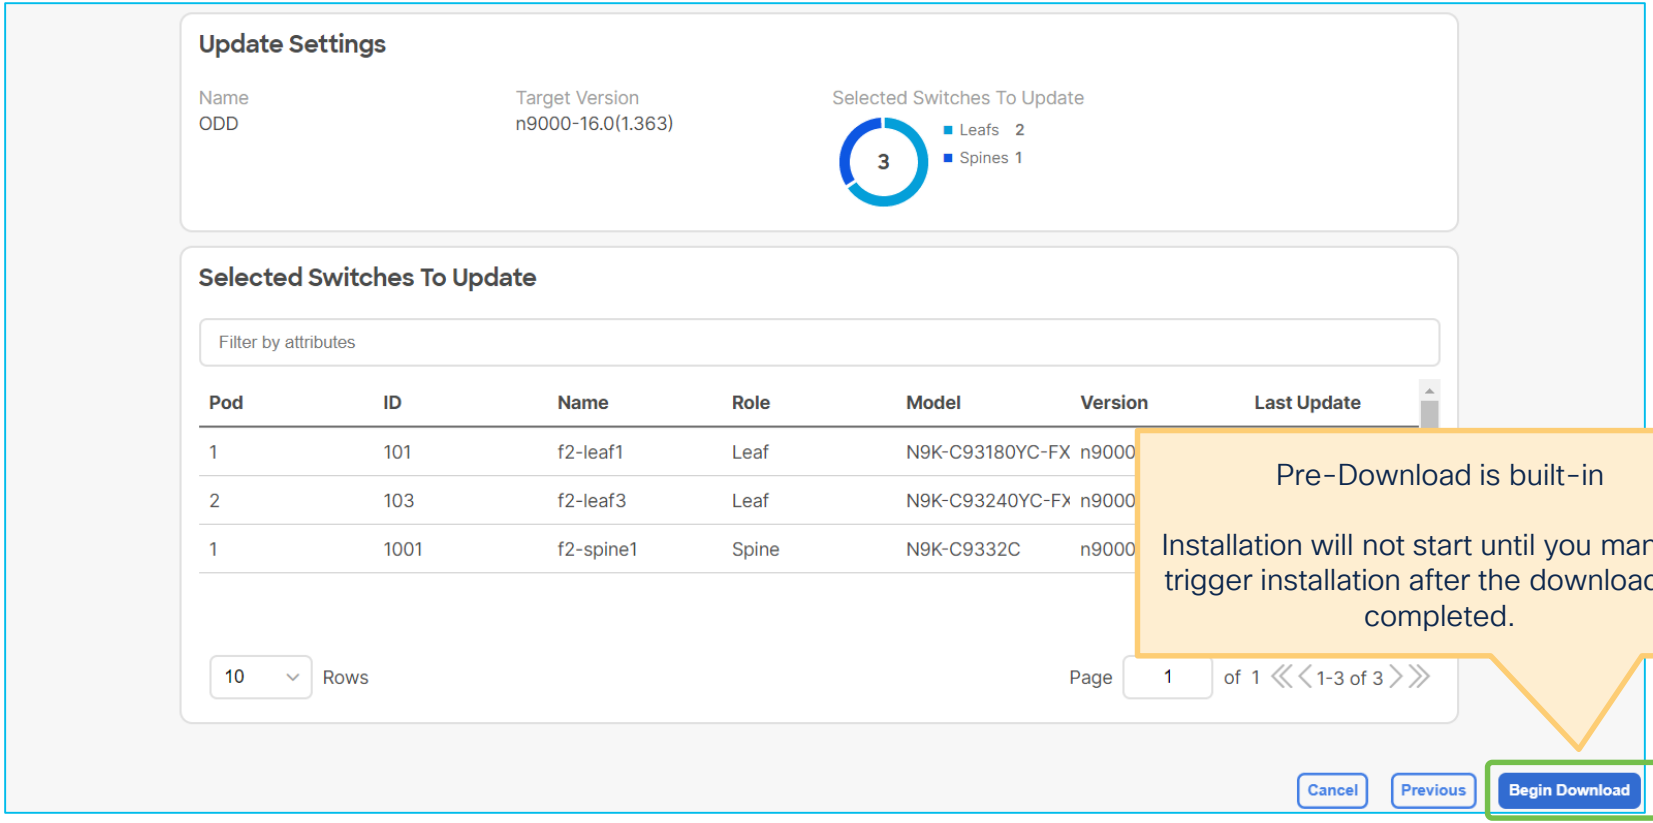

Switch Image Download Progress

5.1(1) NEW

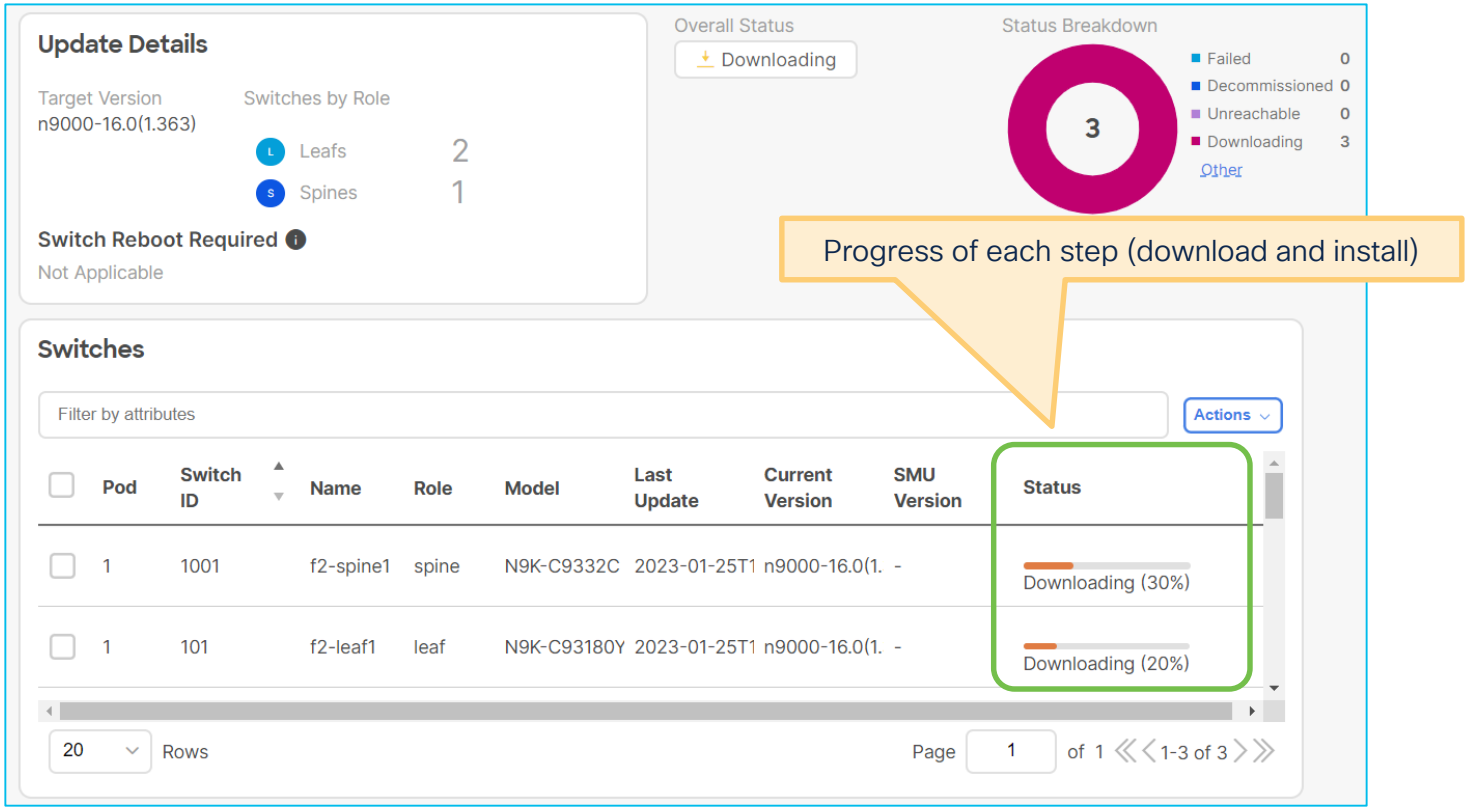

Switch Image Download Progress
(APIC 4.2(5), Switch 14.2(4))

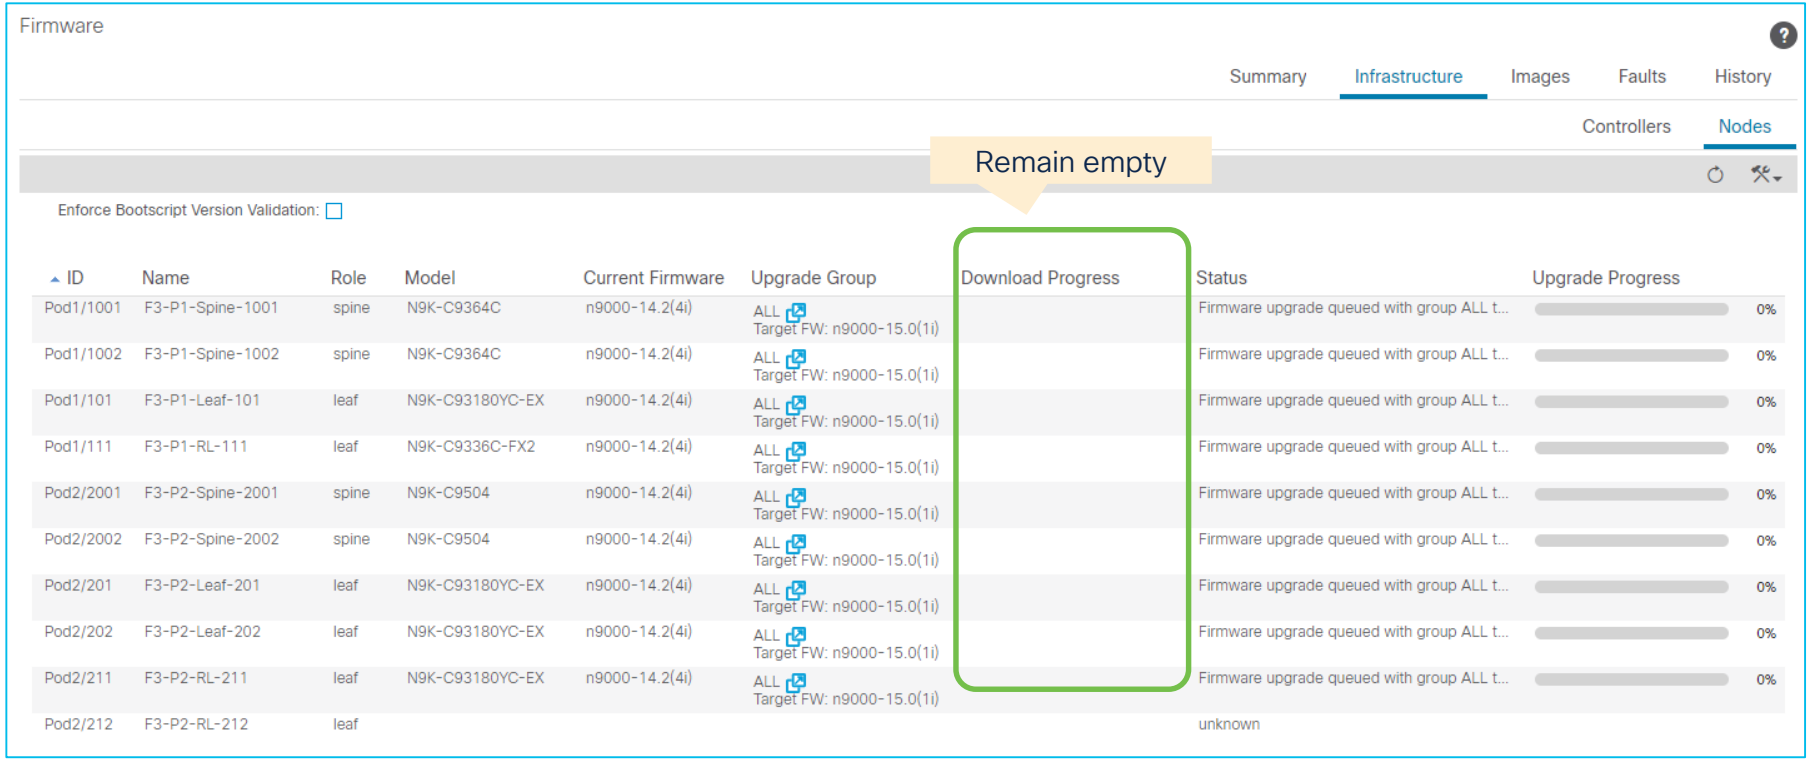

ID	Name	Role	Model	Current Firmware	Upgrade Group	Download Progress	Status	Upgrade Progress
Pod1/1001	F3-P1-Spine-1001	spine	N9K-C9364C	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod1/1002	F3-P1-Spine-1002	spine	N9K-C9364C	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod1/101	F3-P1-Leaf-101	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod1/111	F3-P1-RL-111	leaf	N9K-C9336C-FX2	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod2/2001	F3-P2-Spine-2001	spine	N9K-C9504	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod2/2002	F3-P2-Spine-2002	spine	N9K-C9504	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod2/201	F3-P2-Leaf-201	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod2/202	F3-P2-Leaf-202	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod2/211	F3-P2-RL-211	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL Target FW: n9000-15.0(1i)		Firmware upgrade queued with group ALL t...	0%
Pod2/212	F3-P2-RL-212	leaf					unknown	

Download Progress will not be displayed when switches are older than 14.2(5) even if APIC is 4.2(5) or later

Upgrade multiple pods/switches in parallel

OLD

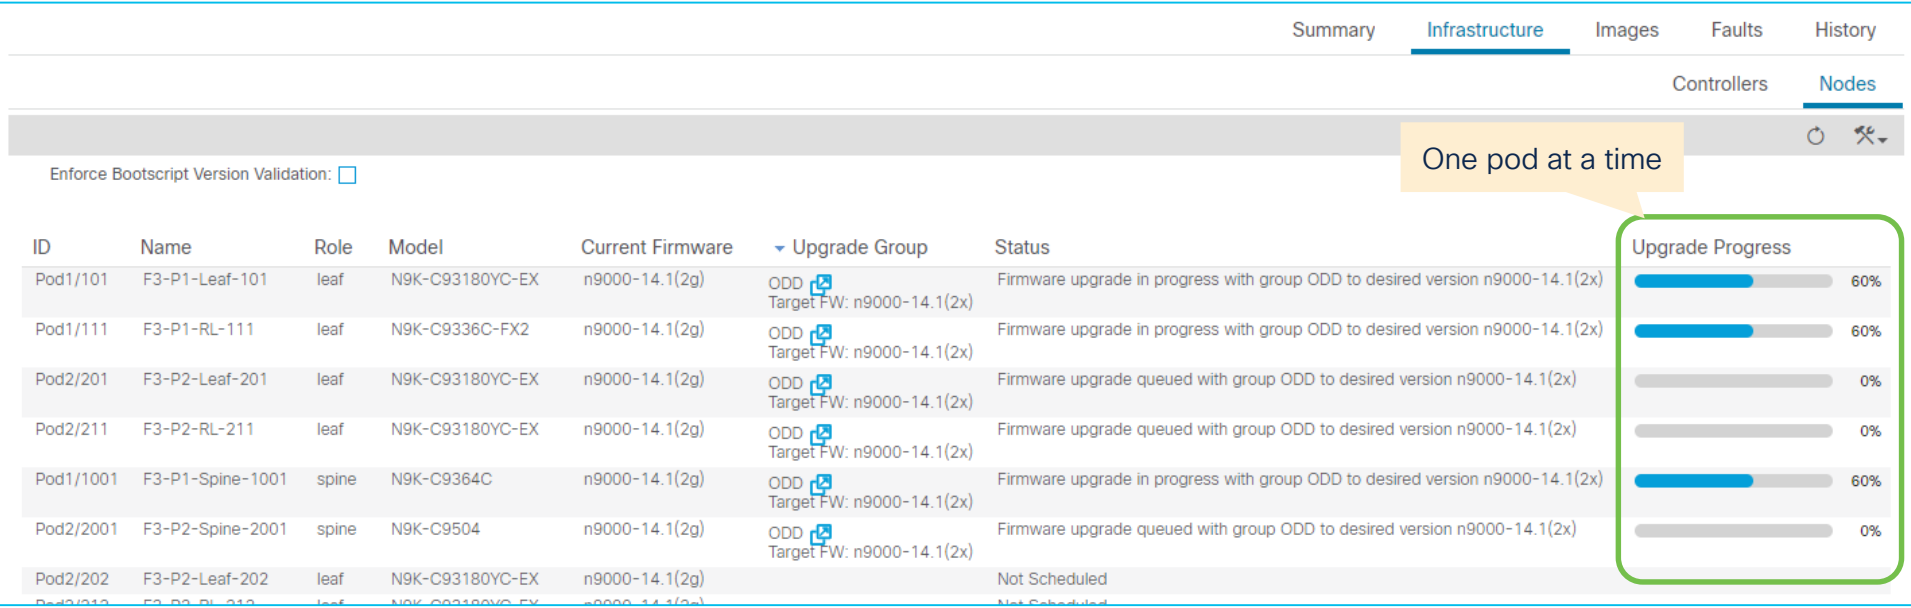

When the actual upgrade starts, APICs allow each switch to upgrade based on the following rules;

- One Pod at a time (14.2(5) has an update)
- When triggered with “Upgrade Now”, 20 switches at a time (14.2(5) has an update)
- When a vPC pair leaf nodes are in the same group, only one of the pair at a time

Unlimited Parallel Upgrade

4.2(5) NEW

All pods at once

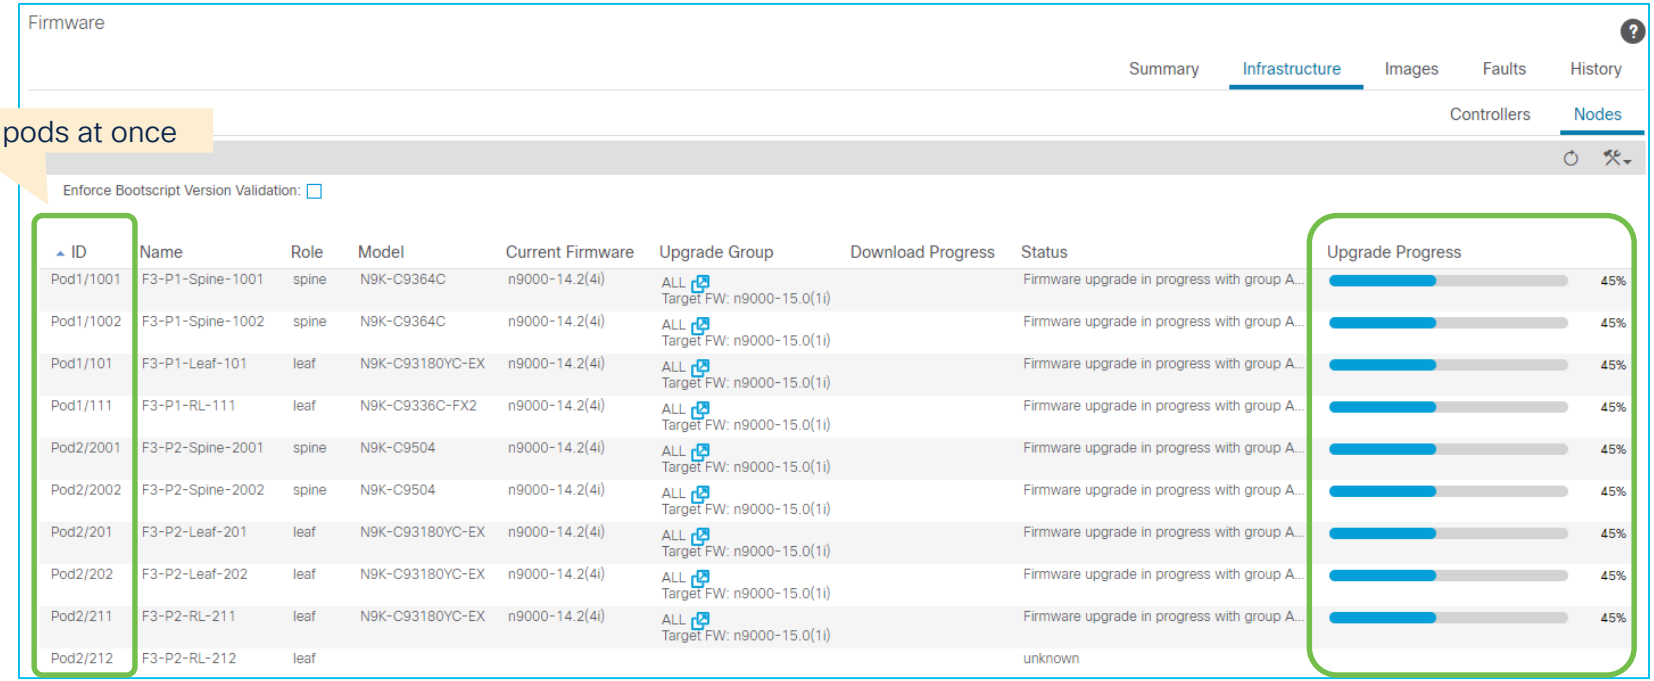

- From APIC 14.2(5) or later, any switches in any pods can be upgraded in parallel
- “Upgrade Now” is no longer limited to 20 switches at a time

Agenda

- Upgrade Architecture
  - ACI Firmware Upgrade Types
  - Upgrade Architecture – APIC
  - Upgrade Architecture – Switches
  - (Bonus) Upgrade Enhancements
- Best Practices
  - Best Practices Workflow Review
  - Best Practices Configurations
  - “Pre-Upgrade Validation” Review and Execution
  - “Do’s and Don’ts”

Recommended Guides

Cisco ACI Upgrade Checklist – Important Starting Point

ACI Upgrade Checklist:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/Cisco-ACI-Upgrade-Checklist.html

Detailed Upgrade Guide (the basis for this presentation)

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/all/apic-installation-aci-upgrade-downgrade/Cisco-APIC-Installation-ACI-Upgrade-Downgrade-Guide.html

Cisco ACI Upgrade Checklist

	Task
☐	Pick your target APIC and ACI switch versions. • Both APICs and ACI switches must be upgraded to the same version. • APIC and ACI switch versions that are compatible to each other are described in the form of x.y(z) and 1x.y(z). For instance, APIC version 5.2(1g) corresponds to ACI switch version 15.2(1g). • Check the *Release Notes* (APIC and ACI switches) for the target version for any open caveats or defects.
☐	See the *APIC Upgrade/Downgrade Support Matrix* for the supported upgrade paths from your current version. • If your current version and the target version are too far apart, you might need to upgrade both your APICs and switches to an intermediate version suggested in the *APIC Upgrade/Downgrade Support Matrix* first. See Multistep Upgrades in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide* for more information. • The *APIC Upgrade/Downgrade Support Matrix* also shows you which CIMC version you need to use for your target APIC version.
☐	Review the ACI upgrade architecture. See ACI Upgrade Architecture in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide* to understand what you should expect along with what you must not perform.
☐	Export your configuration for backup. See the Cisco ACI Configuration Files: Import and Export document for details. Ensure that AES encryption is enabled.
☐	Disable all App Center apps on the APICs. See Guidelines for App Center Apps in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide* for details.
☐	Download both APIC and ACI switch firmware to your APICs. See the *Downloading APIC and Switch Images on APICs* section in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide* for each release for details: • Releases prior to 4.x: Downloading APIC and Switch Images on APICs • Releases 4.x or 5.0: Downloading APIC and Switch Images on APICs • Release 5.1 or later: Downloading APIC and Switch Images on APICs
☐	Download ACI switch firmware from your APICs to each switch. Starting from switch release 14.1(1), switches can download the image from APICs prior to the upgrade. See Rule 5 – Save time by downloading switch images beforehand in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide* for details.
☐	Perform pre-upgrade validations. See Pre-Upgrade Checklists in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide* for details.
☐	Upgrade CIMC on your APICs if suggested so by the Support Matrix. See Upgrading the CIMC Software in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide* for details.
☐	Upgrade APICs. See the *Upgrading the Cisco APIC* section in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide* for each release for details: • Releases prior to 4.x: Upgrading the Cisco APIC from Releases Prior to Release 4.x • Releases 4.x or 5.0: Upgrading the Cisco APIC From Releases 4.x or 5.0 • Release 5.1 or later: Upgrading the Cisco APIC From Releases 5.1x or Later

ACI Firmware Upgrade Best Practice Checklist

- Determine Desired Software and Check Support Matrix
- Review and Implement Best Practice Configurations
- Discover and Clear any issues raised from “pre-upgrade validations”
- Review Upgrade Architecture and “do’s and don’ts”

ACI Software Life Cycle

1. **Cisco Recommended Software Releases**

   https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/recommended-release/b_Recommended_Cisco_ACI_Releases.html

2. **Cisco ACI Release Notes**

   https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

3. **Cisco ACI Upgrade/Downgrade Support Matrix**

   https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/apicmatrix/index.html

4. 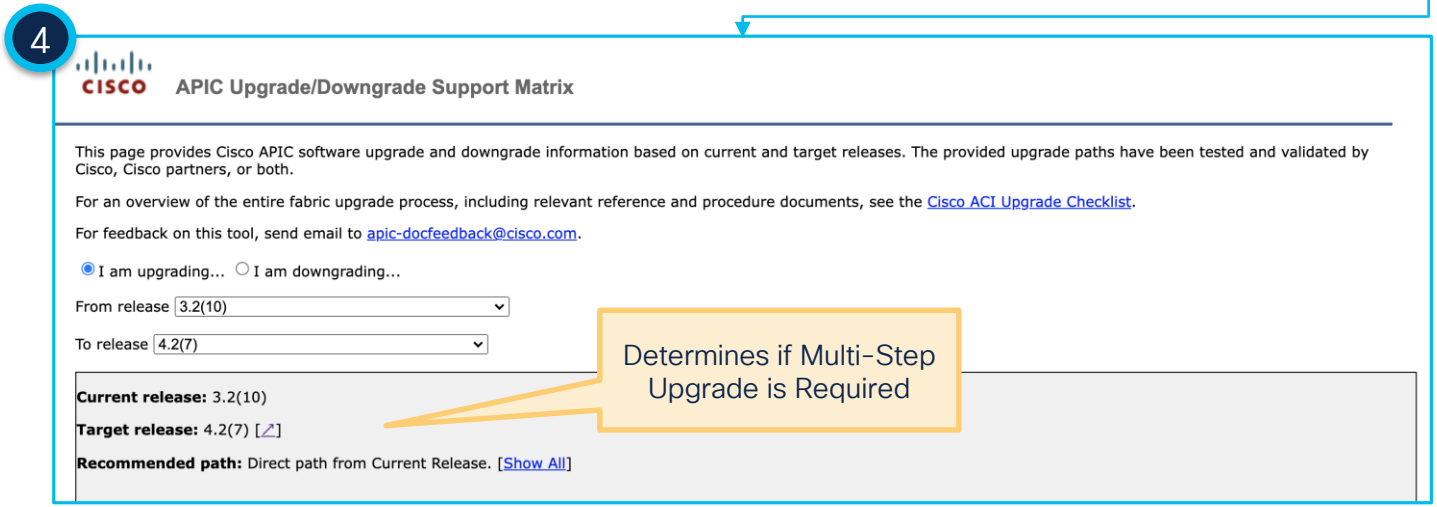

ACI Upgrade Overview

5

Review the ACI Upgrade/Downgrade Guide!

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/all/apic-installation-aci-upgrade-downgrade/Cisco-APIC-Installation-ACI-Upgrade-Downgrade-Guide/m-aci-firmware-upgrade-overview.html#id_48185

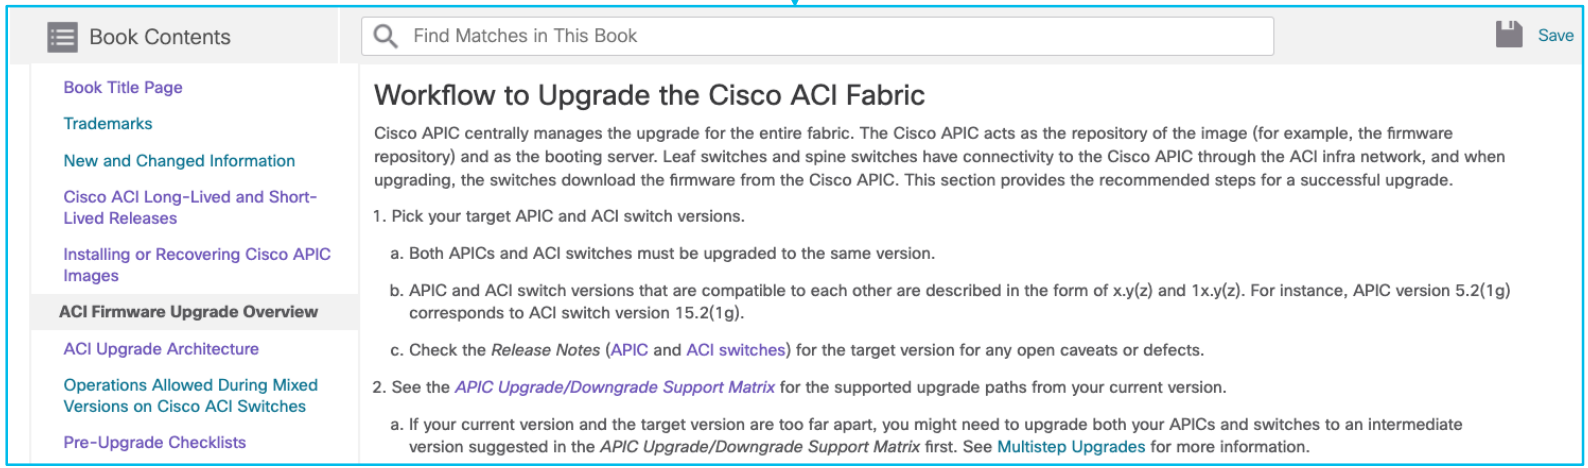

New Release Cadence

Key objectives | Predictable software release cadence | Reach maintenance mode quickly

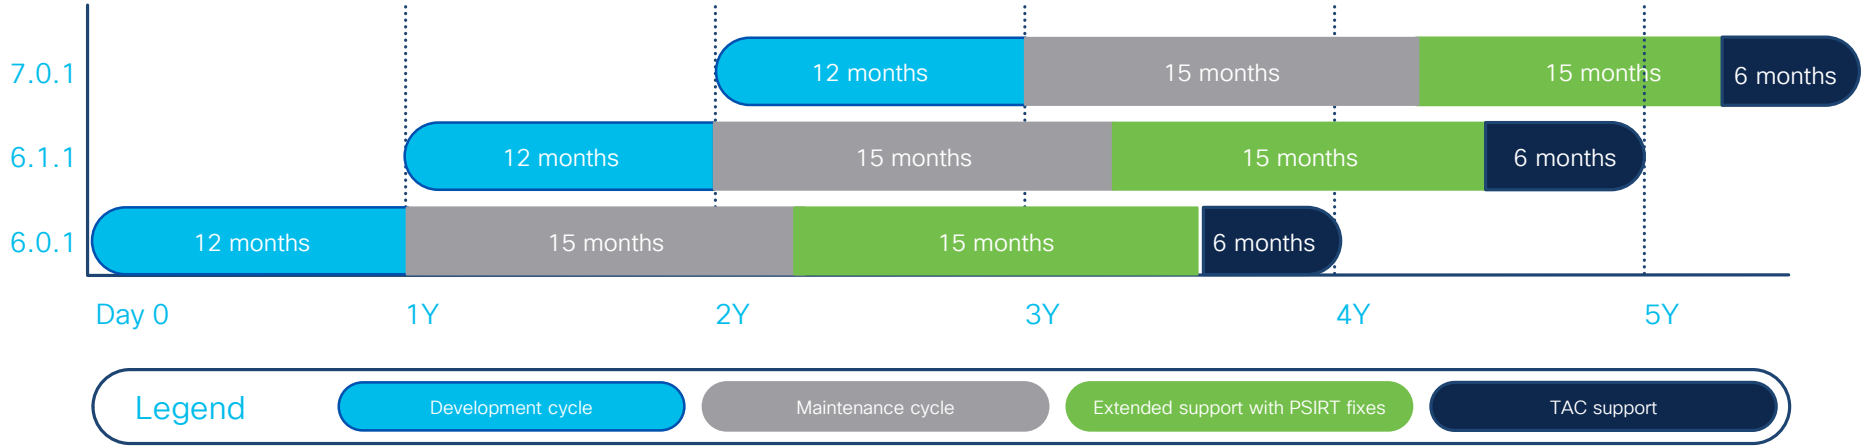

No short-lived and long-lived release tags	Three feature releases from FCS date, including FCS release	Fourth release is a maintenance release (MR), target for golden star	Hardware lifecycle is defined by multiple release and not tied to a single release	Total release lifecycle of four years

CIMC Version Compatibility

6

Option 1: Support Matrix https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/apicmatrix/index.html

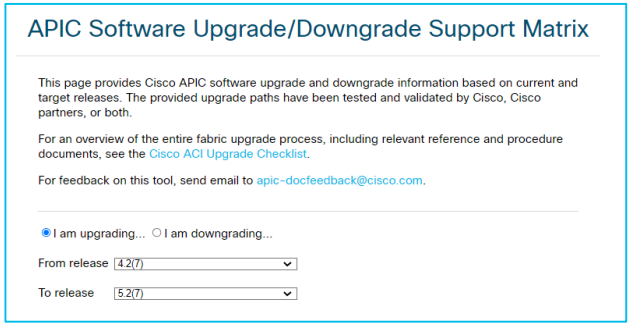

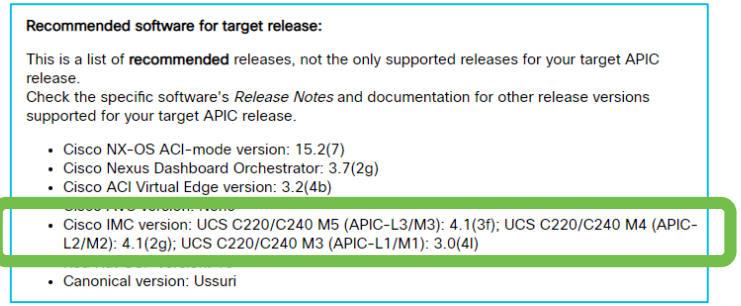

Option 2: APIC Release Note https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

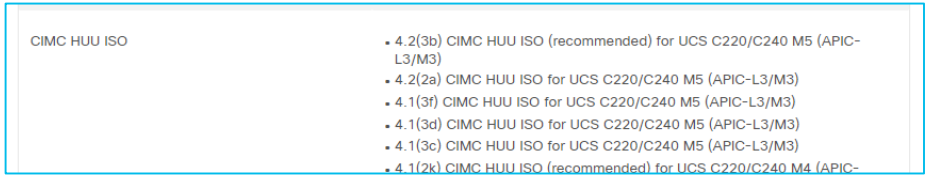

ACI Firmware Upgrade Best Practice Checklist

- [x] Determine Desired Software and Check Support Matrix
- [x] Review and Implement Best Practice Configurations
- [x] Discover and Clear any issues raised from “pre-upgrade validations”
- [x] Review Upgrade Architecture and “do’s and don’ts”

Back Up Configuration

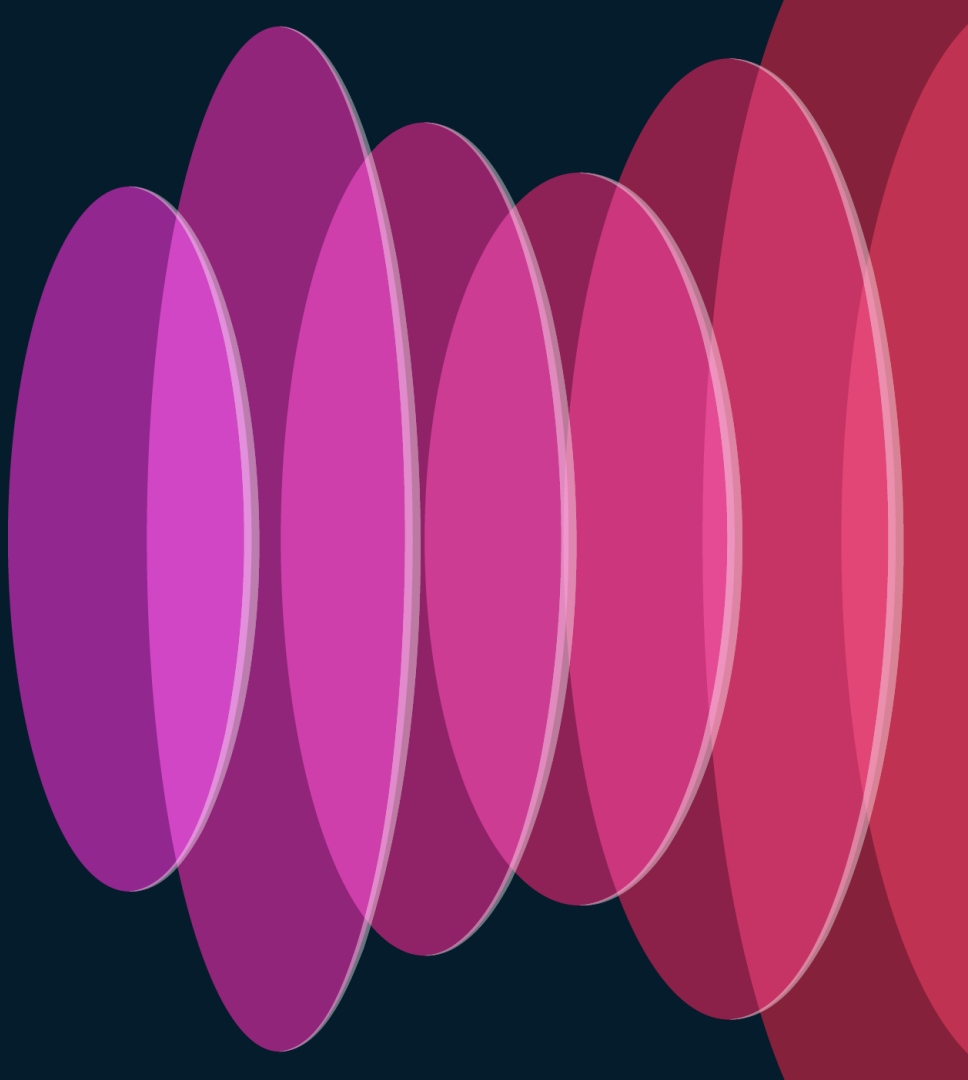

Back Up Configuration with AES File Encryption

- The AES passphrase that generates the encryption keys cannot be recovered or read by an ACI administrator or any other user. The AES passphrase is not stored. Copy your passphrase somewhere safe!
- Setup automatic backups on a scheduler to maintain a consist and up to date backup at all times. Always export it to a remote location.
- In case of upgrade failure, AES backup can be used to recover the system non-disruptively as worst case scenario.

Setting Global AES Encryption allows all the secure properties of the configuration (like credentials) to be successfully imported when restoring the fabric

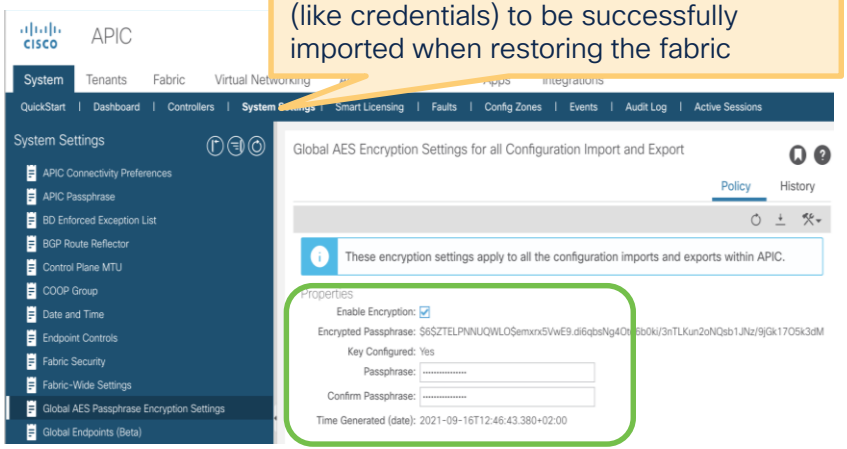

Pre ACI v4.0.1 Setting Location:

Admin > AAA > AES Encryption Passphrase and Keys for Config Export (and Import)

ACI v4.0.1 and later Location:

System > System Settings > Global AES Passphrase Encryption Settings

Technote For Import/Export:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Using_Import_Export_to_Recover_Config_States.html

Switch Upgrade Groups

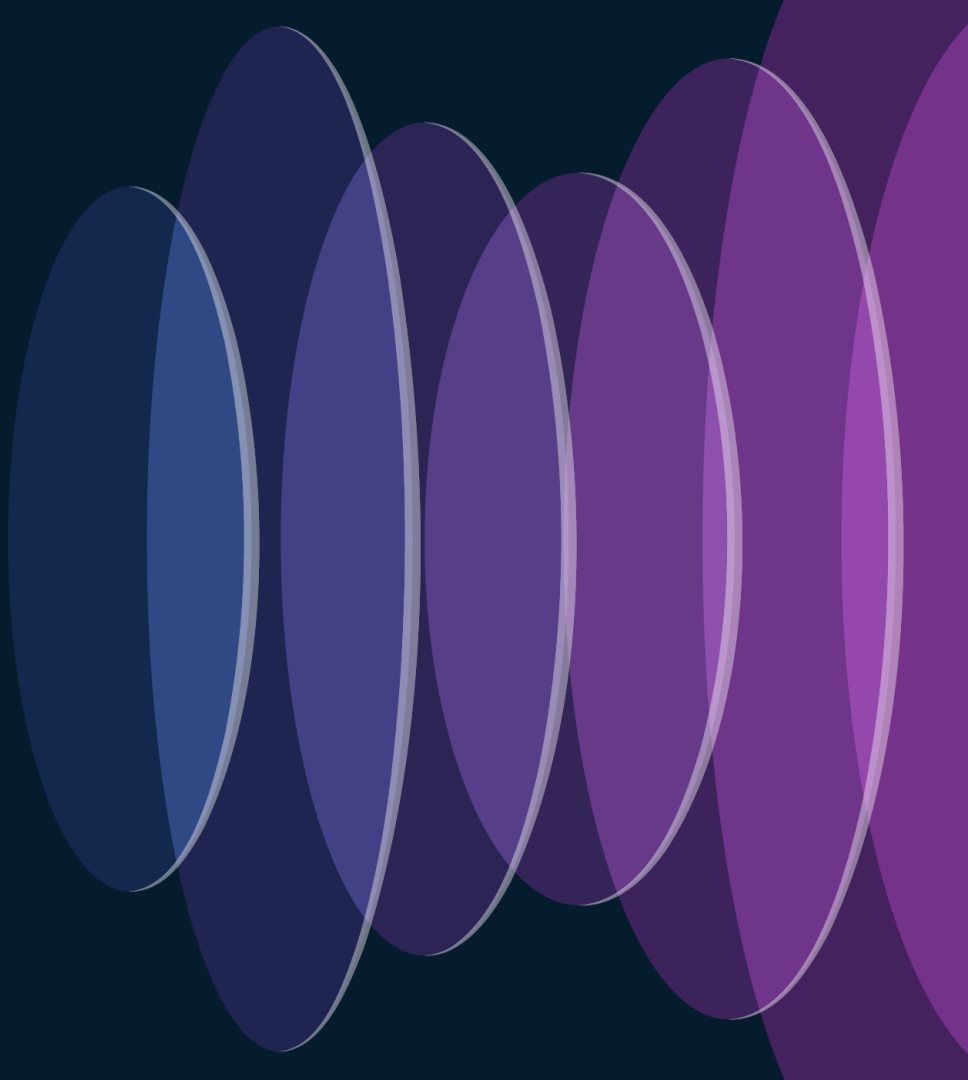

ACI Firmware Upgrade Best Practice 101

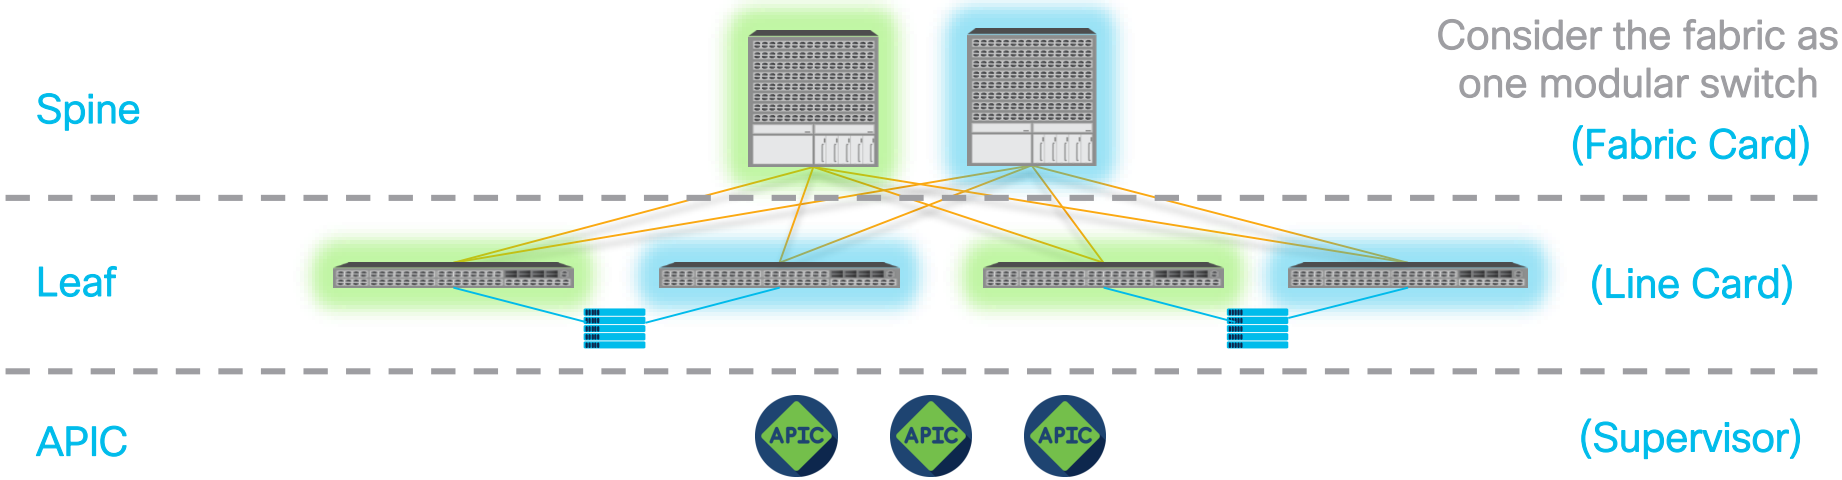

ACI is a solution to manage multiple switches as if it's one huge switch

- APIC (i.e. SUP of the fabric) can be upgraded non-disruptively.
- Each switch (i.e. modules of the fabric) can intelligently choose appropriate switch nodes for non-disruptive traffic flow

Always keep hardware redundancy to achieve zero-to-minimum traffic disruption

1. Upgrade Green switch groups
2. Upgrade Blue switch groups

Switch Upgrade Advanced Options

Upgrade Group

- Name
- Node ID List
- Target Firmware Version
- Scheduler
- **Ignore Compatibility Check**
- **Graceful option**
- **Run Mode**

Later Releases, ie: 5.2
Edit > Version Selection > Advanced Settings

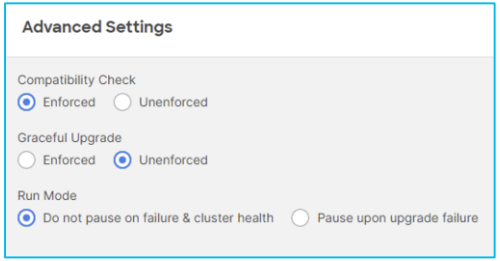

Rule of Thumb
Change defaults only when you must.

- **Compatibility Check** (default: Enforced)
  Only unenforce in a lab where you would like to ignore the supported upgrade path.

- **Graceful Upgrade** (default: Unenforced)
  Only enforce when sub-100ms routing protocol convergence is required.
  Never enforce this when hardware redundancy is lacking. (single spine/leaf pod)

- **Run Mode** (default < 5.1: pause upon upgrade failure
  (default >= 5.1: don't pause upon upgrade failure)
  By default, APIC scheduler will stop putting new switches into queue if
  a) APIC cluster is not fully-fit
  b) The upgrade of previous switches in the same upgrade group failed.
  Ex.) You have 20 leafs in a group. If 1 fails, it will pause all remaining switches that are still queued. If other 19 leafs already started upgrade procedure, those will not be paused.

IS-IS Metric Policy for Multi-Pod and Multi-Site

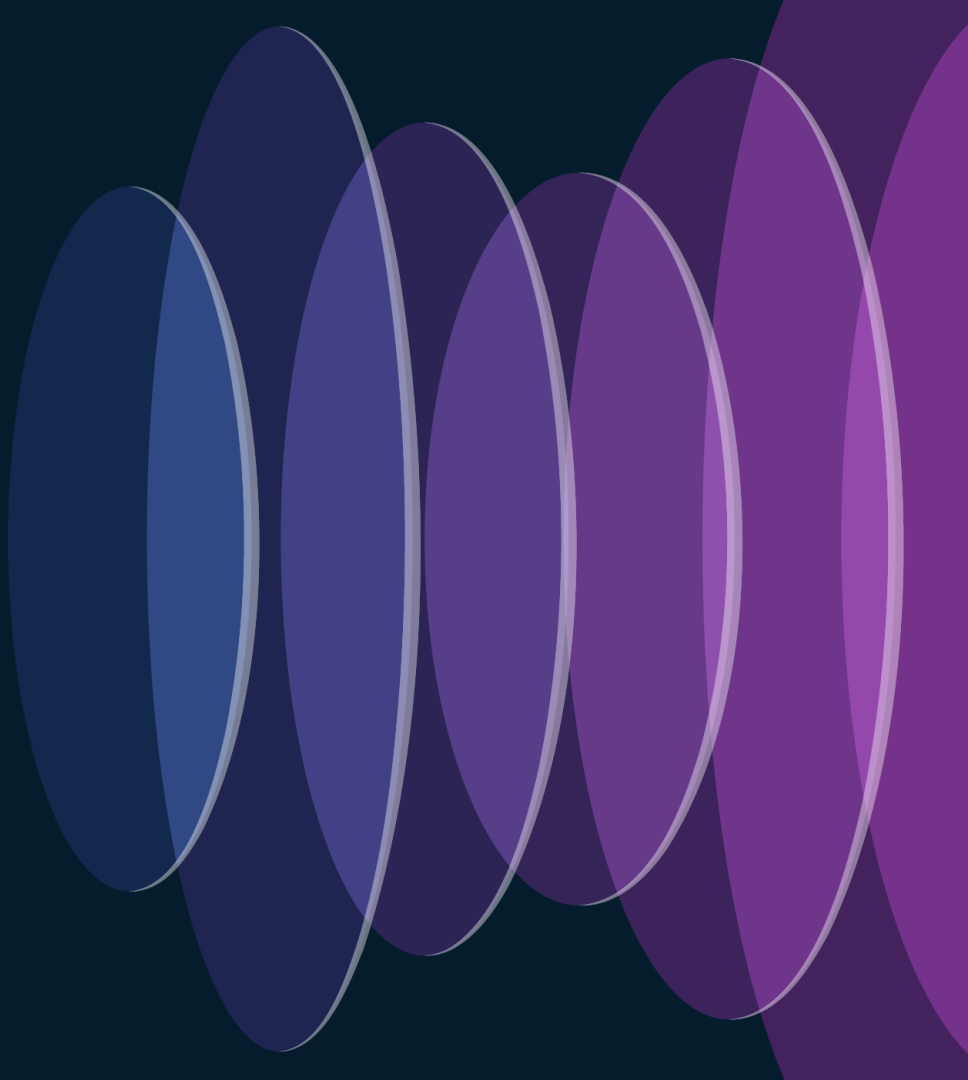

Helpful Tips for Multi-Pod / Multi-Site

ISIS Metric Policy Configuration

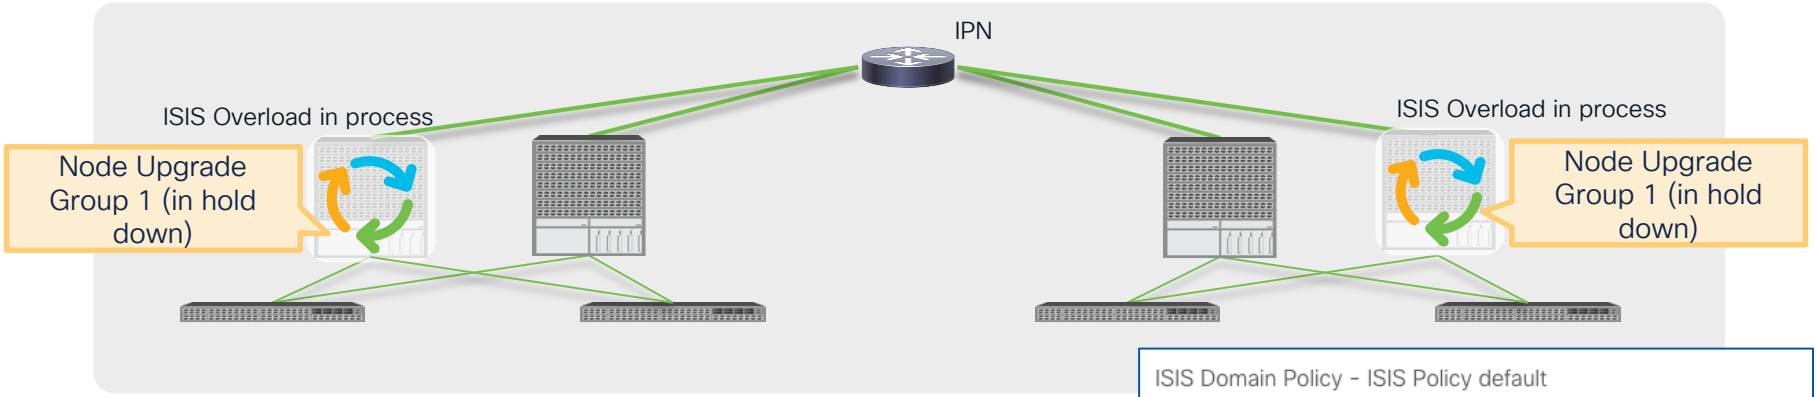

- Default fabric wide IS-IS metric is set at 63 (max value)
- During upgrade, spines set the overload mode while policy is being downloaded.
- If fabric-wide value is already at max, the overload functionality is ineffective.
- This can create unexpected traffic interruption if leaf sends traffic to a spine which is not fully upgraded.

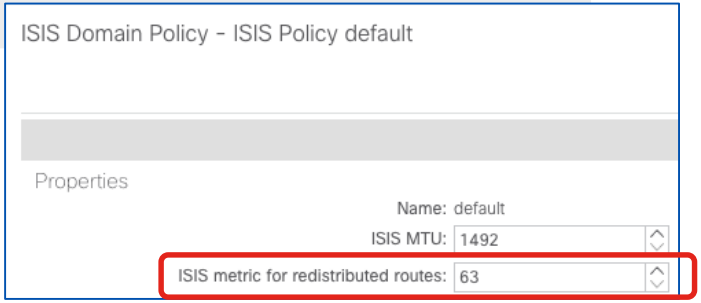

Settings > ISIS Policy (Default Config)

Helpful Tips for Multi-Pod / Multi-Site

ISIS Metric Policy Configuration

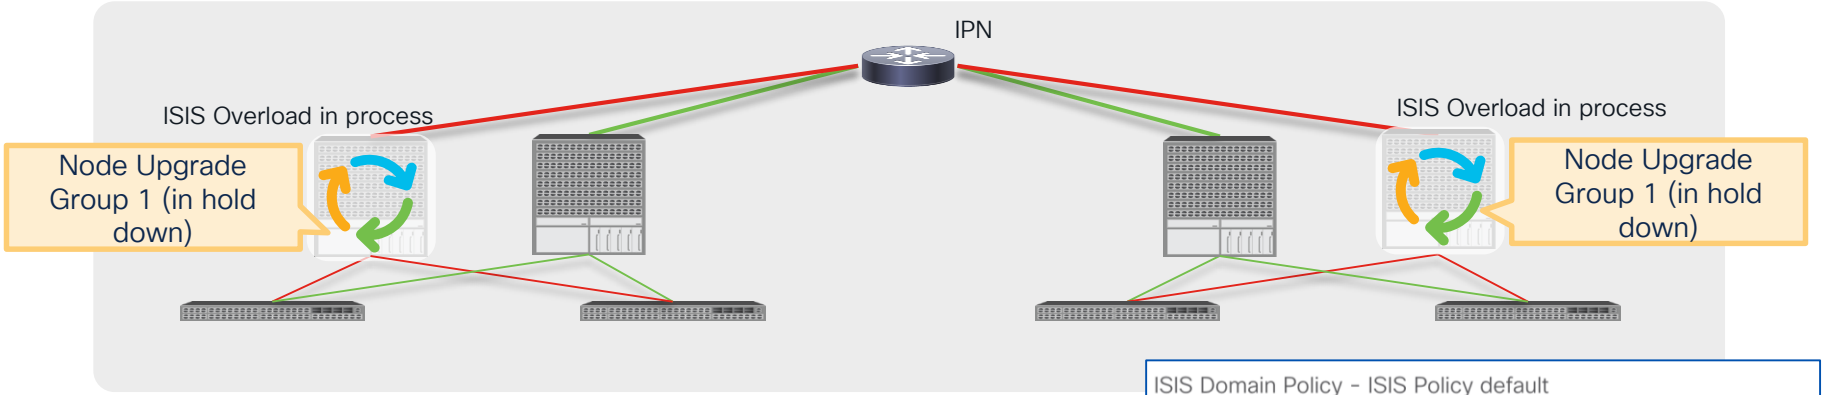

- By Lowering the Value, Remote POD TEP Routes will be preferred through the remaining spines in each POD.
- Once Overload is completed, the spine which was upgraded will advertise these routes using the metric configured.
- This results in ECMP between all spines after the upgrade has completed.

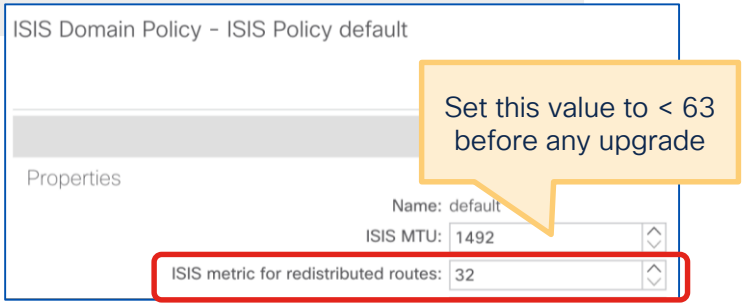

Settings > ISIS Policy

Helpful Tips for Multi-Pod / Multi-Site

Verify Spines are Exchanging Routes to the IPN after upgrade

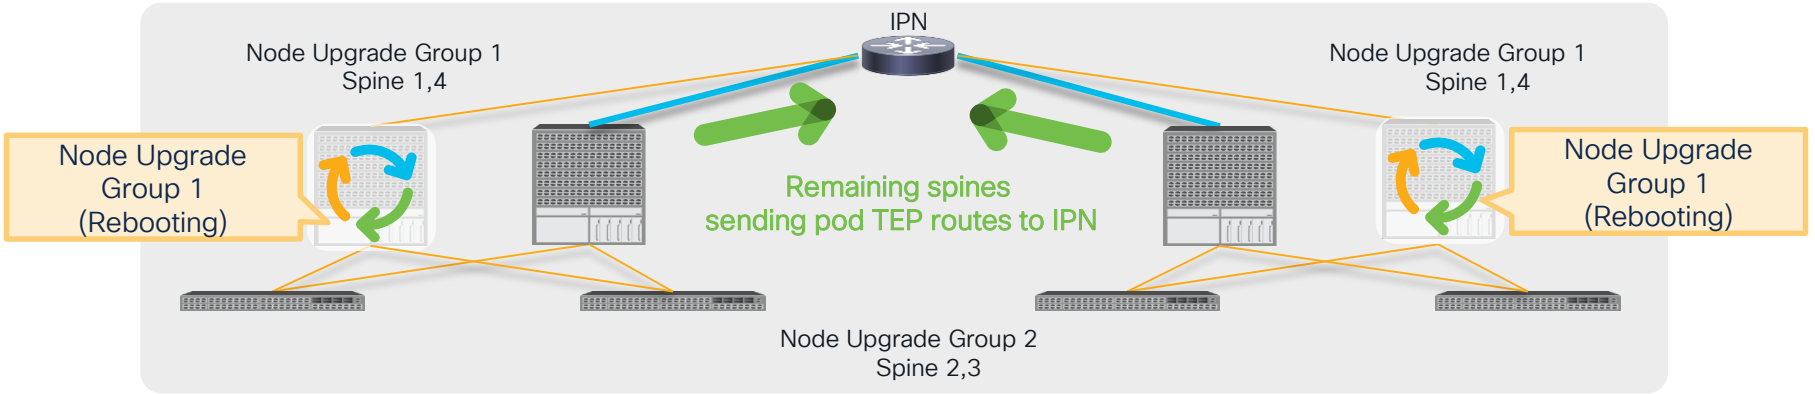

- When Node Upgrade Group 1 finishes, Spines may show as “completed” in upgrade UI but routes towards IPN/ISN may still be in hold down period (up to 10 min)
- Before starting Spine Node Upgrade Group 2, verify that TEP routes of pods / sites are being sent / received from newly upgraded spines in Group 1

Helpful Tips for Multi-Pod / Multi-Site

Verify Spines are Exchanging Routes to the IPN after upgrade

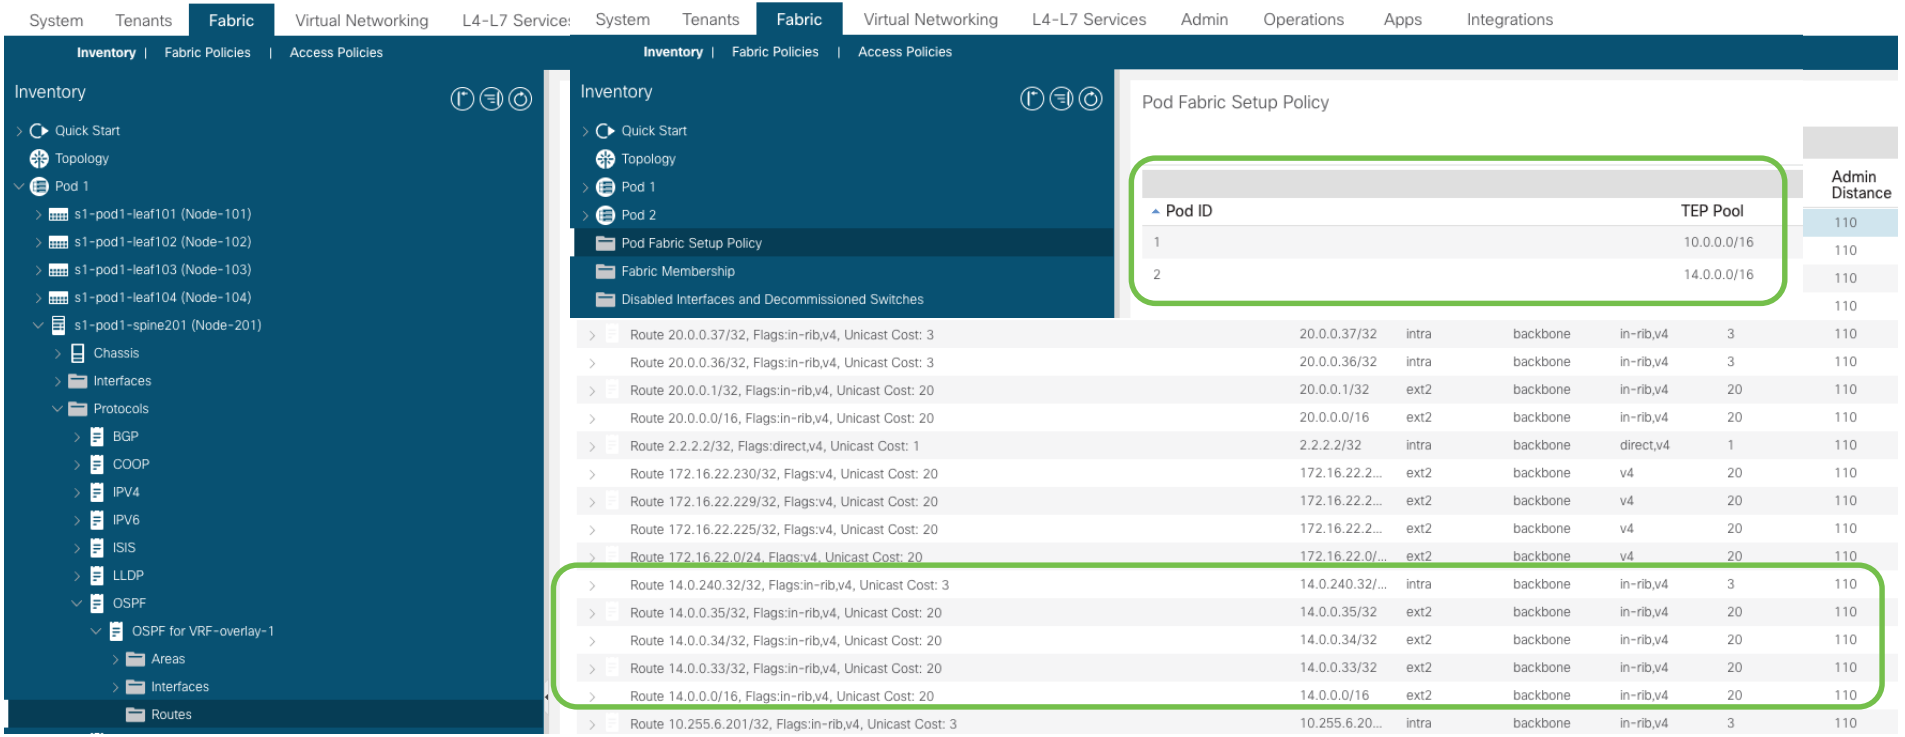

ACI Firmware Upgrade Best Practice Checklist

- Determine Desired Software and Check Support Matrix
- Review and Implement Best Practice Configurations
- Discover and Clear any issues raised from “pre-upgrade validations”
- Review Upgrade Architecture and “do’s and don’ts”

Does your fabric look like this?

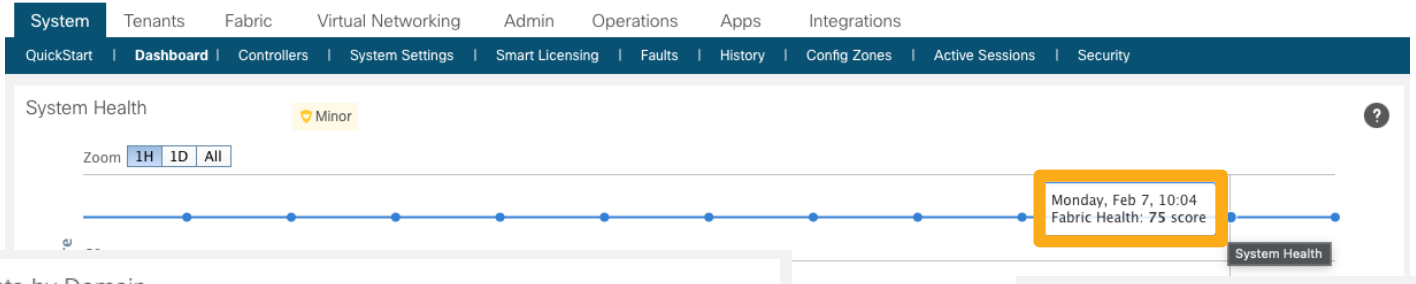

Fault Counts by Domain

Hide Acknowledged Faults Hide Delegated Faults

SYSTEM WIDE	**164**	**1858**	**31173**	**6205**
Access	145	1789	5	82
External	0	0	0	4995
Framework	0	2	0	0
Infra	19	62	1270	1126
Management	0	4	5	0
Security	0	0	214	0
Tenant	0	1	29679	2
Apps	0	0	0	0

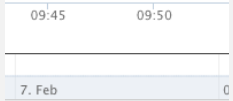

Tenants with Health ≤ 99

Name	Health Score
scale_vpc20	Warning
scale_vpc30	Warning
scale_vpc40	Warning
scale_vpc50	Warning
Tenant_1	Critical
Tenant_10	Critical
Tenant_100	Critical
Tenant_101	Critical

Faults, and the Impact on Upgrades

- Faults can be raised if there is an overlap, or invalid config.
- After an upgrade the switch requests it's configuration "fresh" from APIC. This is the "stateless" behavior of ACI.
- If Logical Config (APIC) has conflicts, the "faulted" config can get pushed before the previously working config.

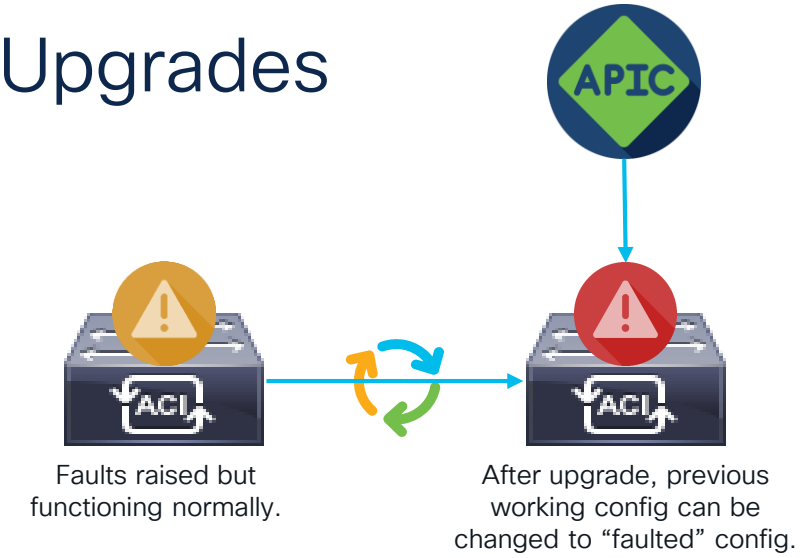


```
L2 Port Config (F0467 port-configured-as-l3)
L3 Port Config (F0467 port-configured-as-l2)
Config On APIC Connected Port (F0467 port-
configured-for-apic)
etc . . .
```

Pre-Upgrade Validation

3.2 – continuing NEW

APIC 3.2, 4.0, 4.1

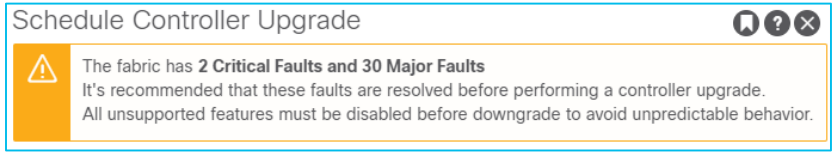

APIC 4.2(1) – 4.2(3)

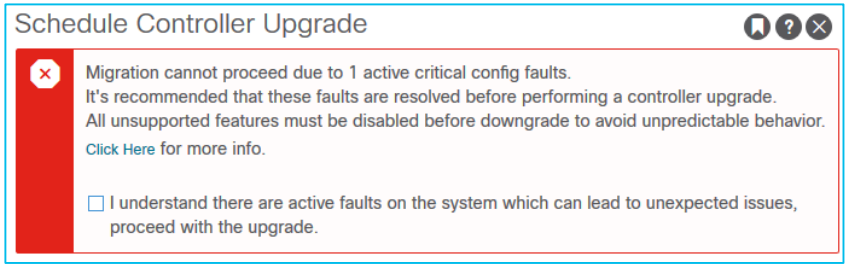

APIC 4.2(4)

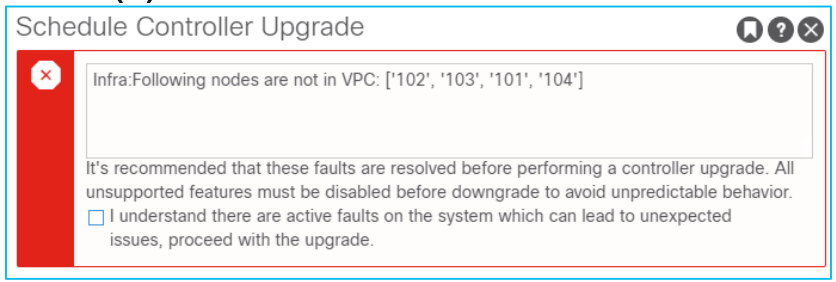

- Prior to 4.2, the APIC upgrade simply warned about the number of all critical and major faults
- On 4.2(1) – 4.2(3), the APIC upgrade warned about
  - ✓ config related critical faults
  - ✓ some specific faults that are known to cause issues during upgrades.
- On 4.2(4), the APIC upgrade warns about
  - ✓ config related critical faults
  - ✓ some specific faults that are known to cause issues during upgrades
  - ✓ A few nonoptimal configurations that may disrupt traffic during the upgrade.
- Additional validation items are being added on each release.

Pre-Upgrade Validation (AppCenter App)

What happens if Cisco adds additional checks?

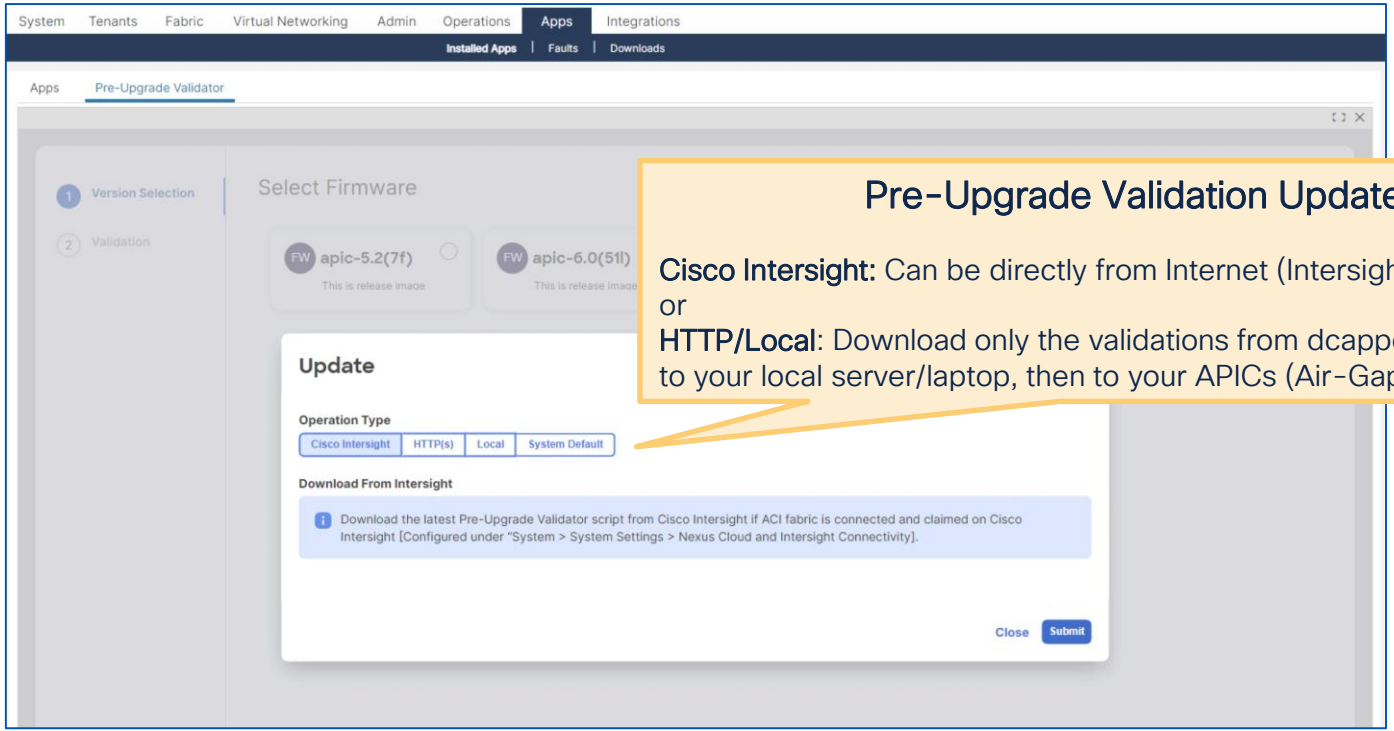

Pre-Upgrade Validation Updates

Cisco Intersight: Can be directly from Internet (Intersight)
or
HTTP/Local: Download only the validations from dcappcenter.cisco.com to your local server/laptop, then to your APICs (Air-Gapped)

https://dcappcenter.cisco.com/pre-upgrade-validator.html
Supported on 5.2. Pre-packaged in APIC in 6.0(2) and later

Pre-Upgrade Validation – Script

https://github.com/datacenter/ACI-Pre-Upgrade-Validation-Script

```
[Check  1/36] APIC Target version image and MD5 hash...
              Checking f2-apic1......

[Check  2/36] Target version compatibility...
[Check  3/36] Gen 1 switch compatibility...
[Check  4/36] Remote Leaf Compatibility... No Remote Leaf Found
[Check  5/36] APIC CIMC Compatibility...
[Check  6/36] APIC Cluster is Fully-Fit...
[Check  7/36] Switches are all in Active state...
[Check  8/36] NTP Status...
[Check  9/36] Firmware/Maintenance Groups when crossing 4.0 Release... Versions not applicable
[Check 10/36] Features that need to be Disabled prior to Upgrade...
  Feature      Name           Status  Recommended Action
  -------      ----           ------  ------------------
  App Center   Policy Viewer  active  Disable the app
  Config Zone  test           Locked  Change the status to "Open" or remove the zone



[Check 11/36] Switch Upgrade Group Guidelines... No upgrade groups found!
[Check 12/36] APIC Disk Space Usage (F1527, F1528, F1529 equipment-full)...
[Check 13/36] Switch Node /bootflash usage... all below 50%
[Check 14/36] Standby APIC Disk Space Usage... No standby APIC found
[Check 15/36] APIC SSD Health (F2731 equipment-wearout)...
[Check 16/36] Switch SSD Health (F3073, F3074 equipment-flash-warning)...
[Check 17/36] Config On APIC Connected Port (F0467 port-configured-for-apic)...
[Check 18/36] L3 Port Config (F0467 port-configured-as-l2)...
[Check 19/36] L2 Port Config (F0467 port-configured-as-l3)...
[Check 20/36] L3Out Subnets (F0467 prefix-entry-already-in-use)...
[Check 21/36] BD Subnets (F1425 subnet-overlap)...
[Check 22/36] BD Subnets (F0469 duplicate-subnets-within-ctx)...
[Check 23/36] VMM Domain Controller Status...
[Check 24/36] VMM Domain LLDP/CDP Adjacency Status... No LLDP/CDP Adjacency Failed Faults Found
```

The goal of the script

To be able to apply the latest validations on any APIC versions via a script

Why the script may be a better choice?:

- Supports older versions – available for everyone!
- Always has the latest checks
- With Github account, you can submit issues or features directly

Both app and script are fully supported by TAC

Pre-Upgrade Validation Script

```
admin@apic1:techsupport> python aci-preupgrade-validation-script.py
    ==== 2024-05-28T08-45-58-0500 ====

Enter username for APIC login          : admin
Enter password for corresponding User  :

Checking current APIC version (switch nodes are assumed to be on the same version)...3.2(10e)

Gathering APIC Versions from Firmware Repository...

[1]: aci-apic-dk9.5.2.7g.bin

What is the Target Version?     : 1

You have chosen version "aci-apic-dk9.5.2.7g.bin"
[Check  1/64] APIC Target version image and MD5 hash...
              Checking fab3-apic1......                                                    DONE
                                                                                           PASS
[Check  2/64] Target version compatibility...                                              PASS
[Check  3/64] Gen 1 switch compatibility...                                                PASS
. . .
. . .
. . .
. . .
. . .
[Check 19/64] L2 Port Config (F0467 port-configured-as-l3)...                              FAIL - OUTAGE WARNING!!
  Fault  Pod    Node      Tenant  AP   EPG   Port    Recommended Action
  -----  ---    ----      ------  --   ---   ----    ------------------
  F0467  pod-1  node-101  jr      ap1  epg1  eth1/6  Resolve the conflict by removing this config or other configs using this port as L3
```

User Enters Credentials

Checks that require login leverage this input

User Selects Target Version

Checks that require target version leverage this input.

Failure Details are Provided

Issue should be corrected (Script Re-Run to validate) before performing upgrade.

Pre-Upgrade Validation Script

```
[Check 59/64] fvUplinkOrderCont with blank active uplinks definition...                    PASS
[Check 60/64] CoS 3 with Dynamic Packet Prioritization...                                  PASS
[Check 61/64] N9K-C93108TC-FX3P/FX3H Interface Down...                                     PASS
[Check 62/64] Invalid FEX Relation Source...                                               PASS
[Check 63/64] LLDP Custom Interface Description Defect...                                  PASS
[Check 64/64] Route-map Community Match Defect...                                          PASS

=== Summary Result ===

PASS                       : 55
FAIL - OUTAGE WARNING!!    :  4
FAIL - UPGRADE FAILURE!!   :  2
MANUAL CHECK REQUIRED      :  1
N/A                        :  2
ERROR !!                   :  0
TOTAL                      : 64

    Pre-Upgrade Check Complete.
    Next Steps: Address all checks flagged as FAIL, ERROR or MANUAL CHECK REQUIRED

    Result output and debug info saved to below bundle for later reference.
    Attach this bundle to Cisco TAC SRs opened to address the flagged checks.

      Result Bundle: /data/techsupport/Scripts/pre-upgrade/preupgrade_validator_2024-05-28T08-45-58-0500.tgz
```

Summary is Provided

All “FAIL” Categories need remediation. Detailed Recommendations to Remediate are in the Upgrade Guide!

Log Bundle is Created

Upload this to any TAC Case if Necessary.

Pre-Upgrade Validation – Script Documentation

[Read the Docs!]

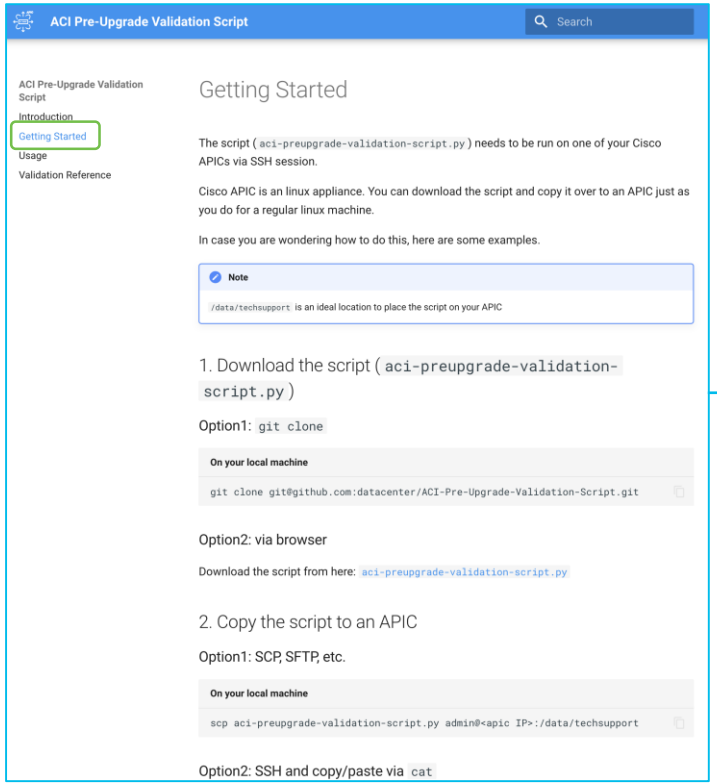

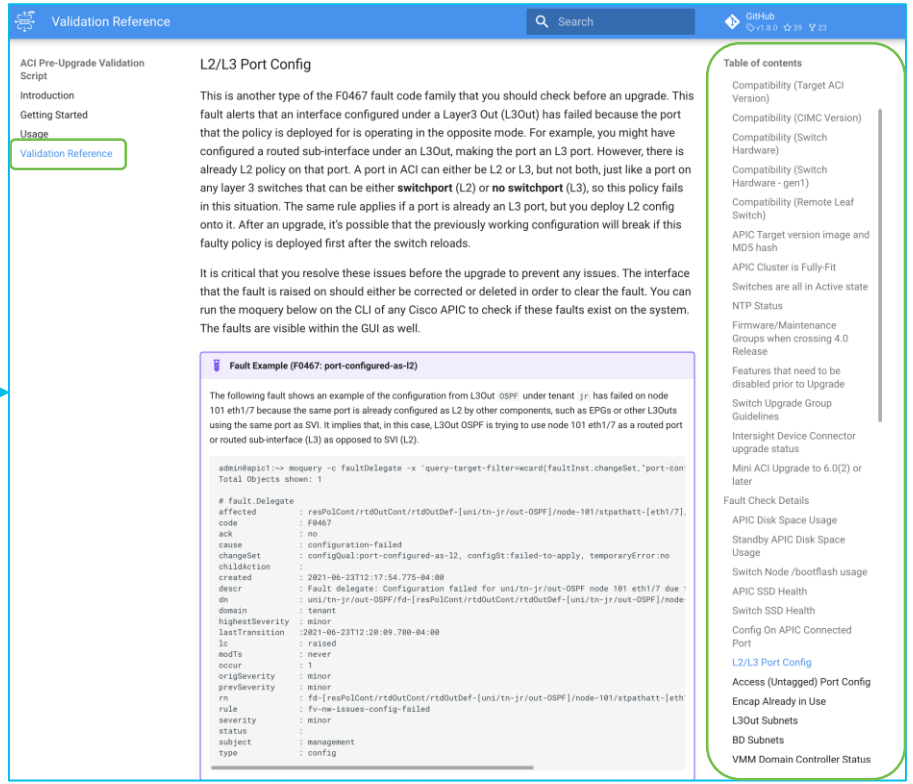

Nexus Dashboard Insights (Optional)

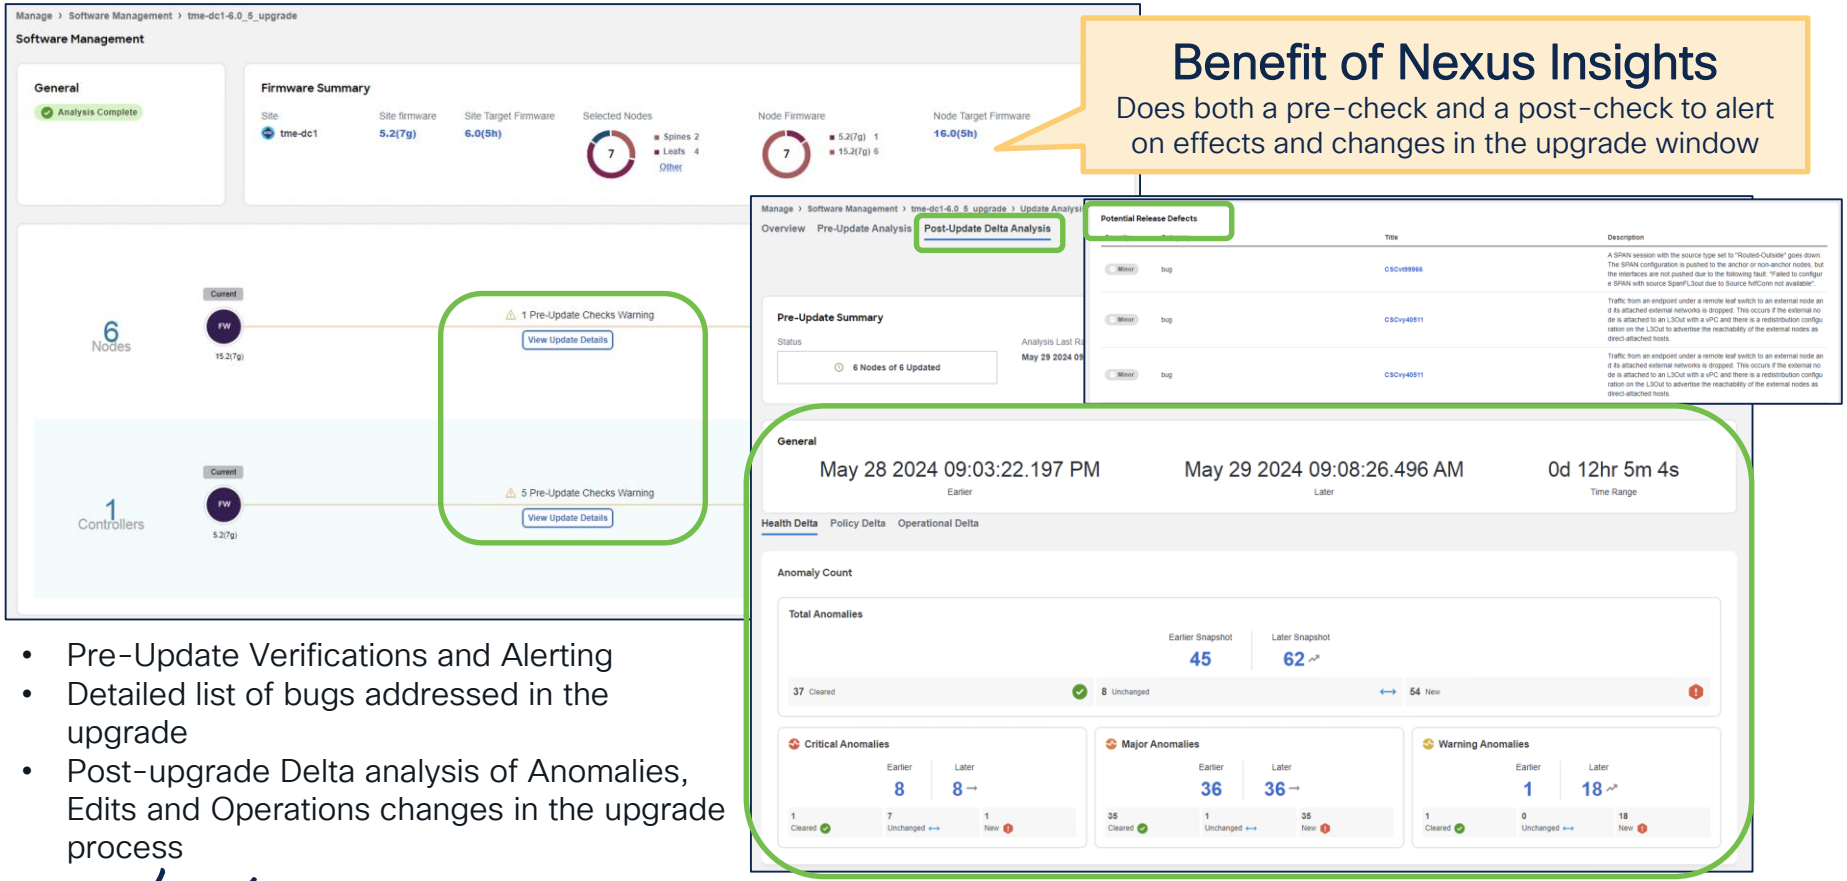

- Pre-Update Verifications and Alerting
- Detailed list of bugs addressed in the upgrade
- Post-upgrade Delta analysis of Anomalies, Edits and Operations changes in the upgrade process

ACI Firmware Upgrade Best Practice Checklist

- Determine Desired Software and Check Support Matrix
- Review and Implement Best Practice Configurations
- Discover and Clear any issues raised from “pre-upgrade validations”
- Review Upgrade Architecture and “Do’s and Don’ts”

Do's and Don'ts

If at any point in time you believe the upgrade/downgrade has either stalled or failed, follow the guidelines below:

Do View the APIC Faults and Installer Logs.

Do Collect the Tech Support Files.

Do Contact Cisco TAC if Needed.

```
admin@apic1:logs> pwd
/firmware/logs
admin@apic1:logs> ls -l
2021-04-15T07:42:57-50
2021-05-28T10:18:33-50
admin@apic1:logs> ls -l ./2021-05-28T10:18:33-50
atom_installer.log
insieme_4x_installer.log

leaf101# pwd
/mnt/pss
leaf102# ls installer_detail.log
installer_detail.log
```

```
admin@apic1:~> techsupport local
```

Do’s and Don’ts

If at any point in time you believe the upgrade/downgrade has either stalled or failed, it is critical that you do not take any of the actions listed below:

Don't reload any APIC in the cluster manually.

Don't decommission any APIC in the cluster.

Don't change the firmware target version back to the original version.

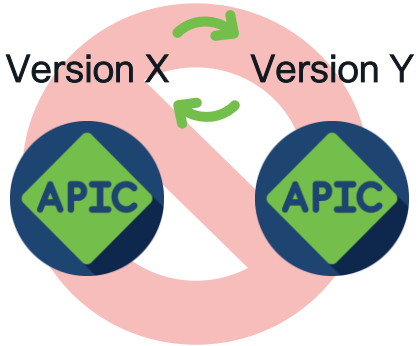

Final Tip

You've read the "Do's and Don'ts"...

When in Doubt, Contact Cisco Support

With Proper Backups, Recovery is Always an Option

ACI Firmware Upgrade Best Practice Checklist

- Determine Desired Software and Check Support Matrix
- Review and Implement Best Practice Configurations
- Discover and Clear any issues raised from “pre-upgrade validations”
- Review Upgrade Architecture and “do’s and don’ts”

Key points to remember

- Always make sure you are performing a supported upgrade.
- Best Practice Configuration and Backups are Critical to Success
- ACI Pre-Upgrade Validations will prevent known issues from impacting the upgrade.
- Never perform a disruptive procedure during an upgrade without help from Cisco.

Reference

- Cisco APIC Installation and ACI Upgrade and Downgrade Guide
  https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/all/apic-installation-aci-upgrade-downgrade/Cisco-APIC-Installation-ACI-Upgrade-Downgrade-Guide.html
- Cisco ACI Upgrade Checklist
  https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/Cisco-ACI-Upgrade-Checklist.html
- Cisco APIC Release Notes
  https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html
- Release Notes for Cisco Nexus 9000 Series Switches in ACI Mode
  https://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-release-notes-list.html
- Getting Started Guide (NX-OS to ACI POAP Auto-conversion)
  https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/getting-started/cisco-apic-getting-started-guide-52x/fabric-initialization-52x.html#d5018e3247a1635

Reference

- Cisco APIC Installation and ACI Upgrade / Downgrade Guide
  https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/all/apic-installation-aci-upgrade-downgrade/Cisco-APIC-Installation-ACI-Upgrade-Downgrade-Guide.html
- Cisco ACI Upgrade Checklist
  https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/Cisco-ACI-Upgrade-Checklist.html
- Cisco APIC Release Notes
  https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html
- Release Notes for Cisco Nexus 9000 Series Switches in ACI Mode
  https://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-release-notes-list.html
- Cisco ACI Upgrade Matrix
  https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/apicmatrix/index.html
- Pre-Upgrade Validation Script
  https://github.com/datacenter/ACI-Pre-Upgrade-Validation-Script
- Pre-Upgrade Validation Script Documentation
  https://datacenter.github.io/ACI-Pre-Upgrade-Validation-Script/

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 2025.

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

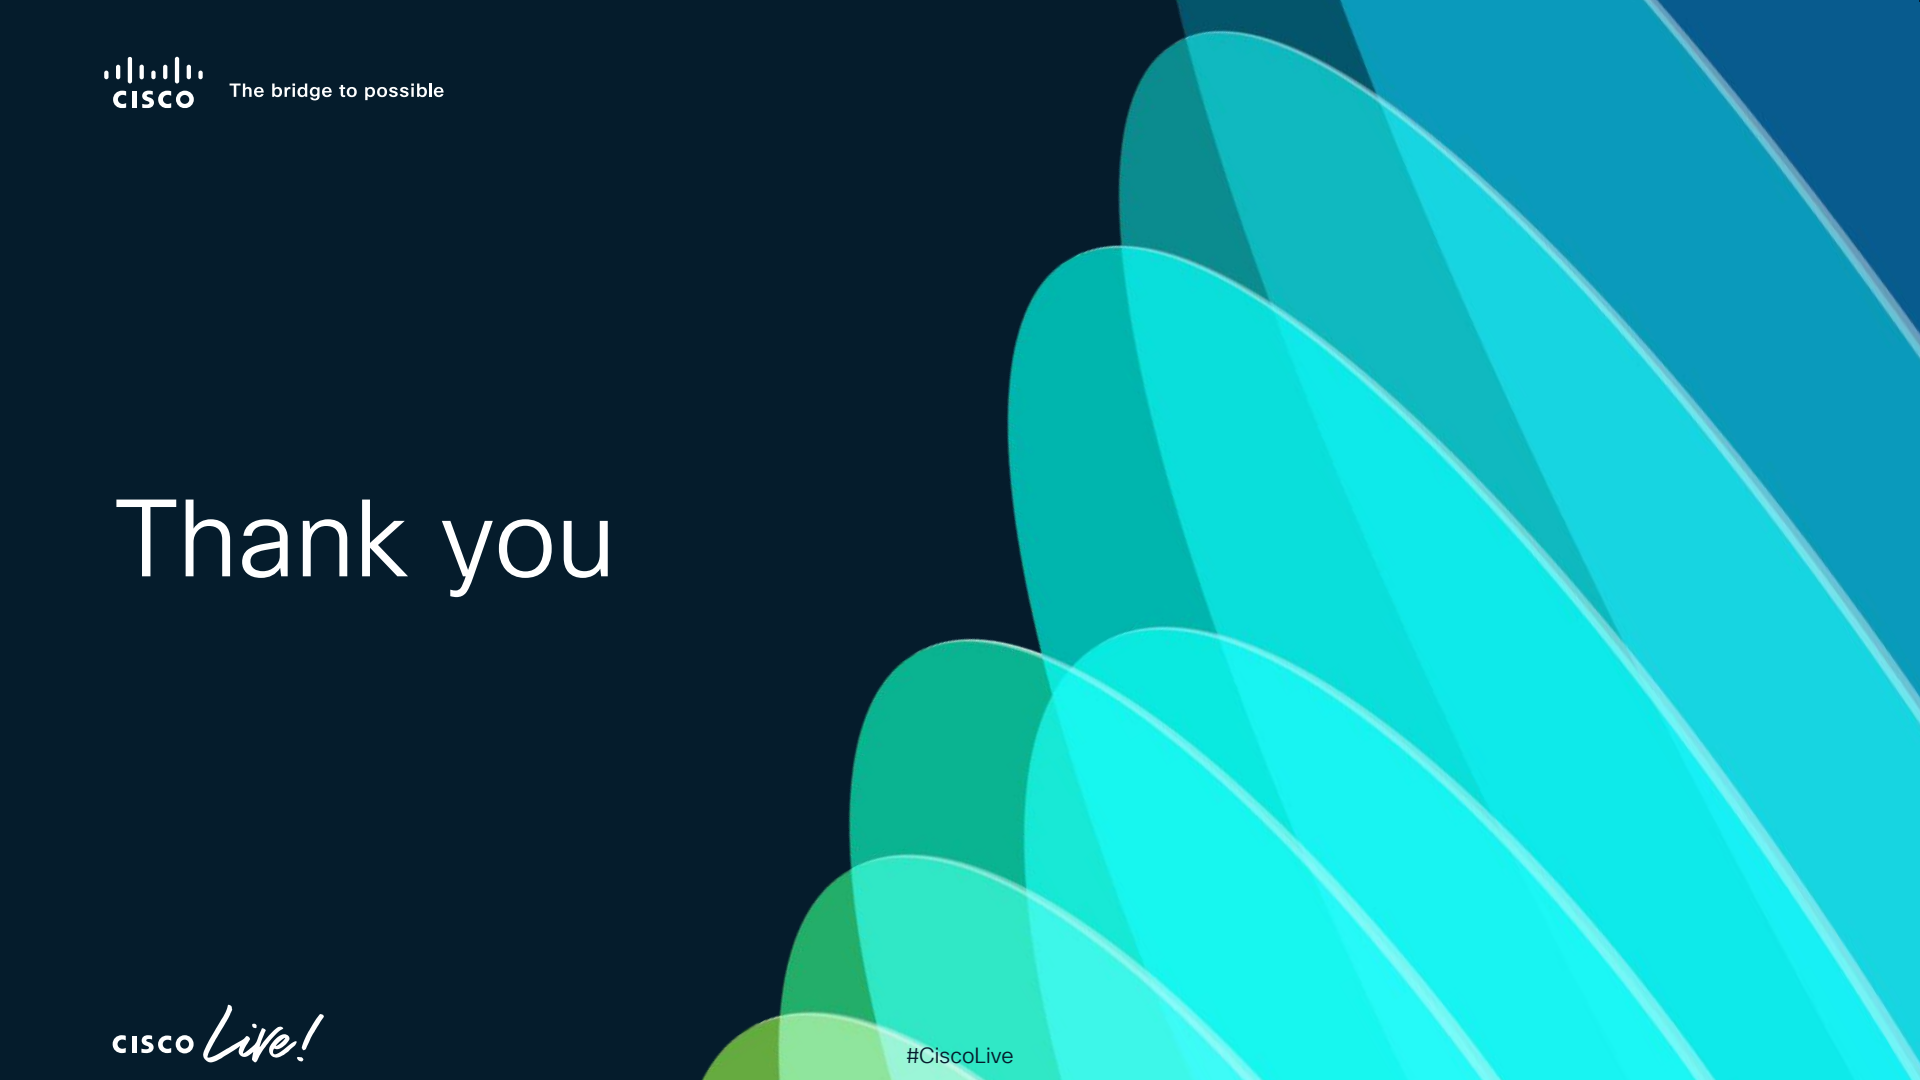
The bridge to possible
Thank you
cisco Live!
#CiscoLive