cisco Live!

Let's go

The bridge to possible

Why You Shouldn’t Fear Upgrading Your ACI Fabric

The Handbook!

Takuya Kishida, Technical Leader, Technical Market Engineering
Joey Ristaino, Technical Leader, Technical Market Engineering

BRKDCN-2910

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?

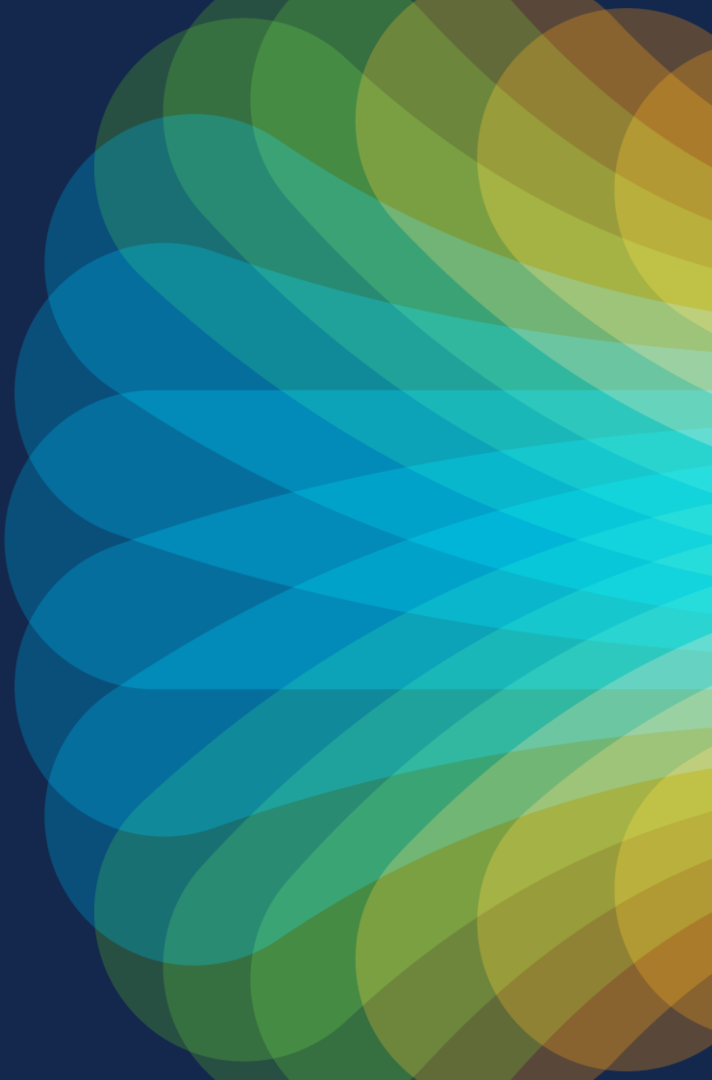

New features, new use cases, new possibilities

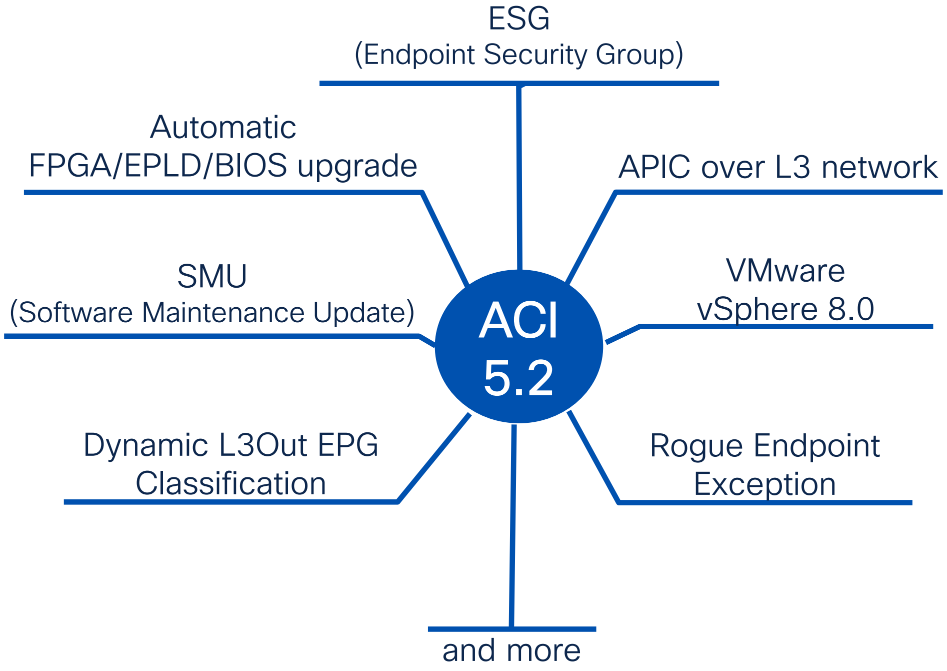

ACI 5.2

ESG (Endpoint Security Group)

Automatic FPGA/EPLD/BIOS upgrade

APIC over L3 network

SMU (Software Maintenance Update)

VMware vSphere 8.0

Dynamic L3Out EPG Classification

Rogue Endpoint Exception

and more

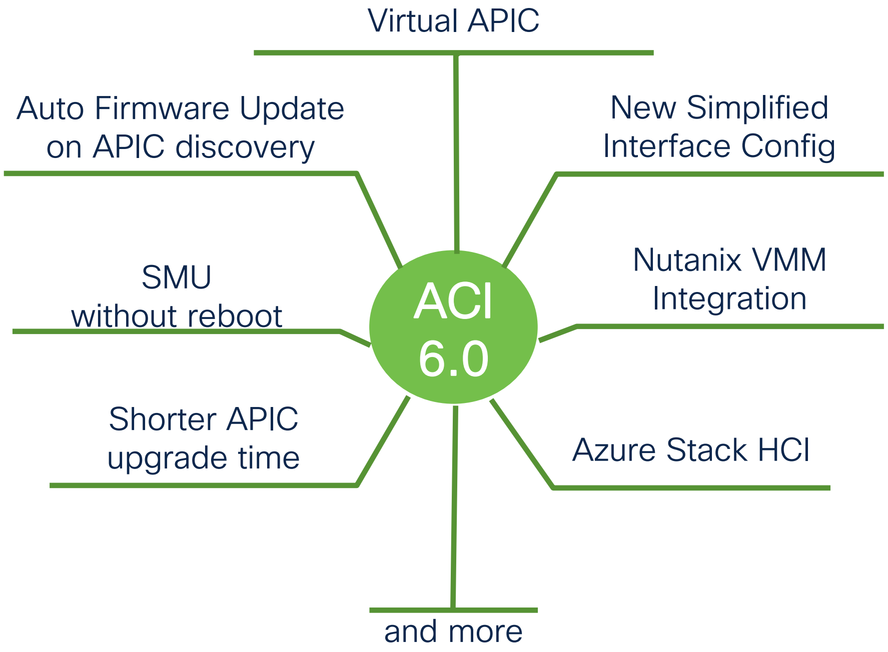

ACI 6.0

Virtual APIC

Auto Firmware Update on APIC discovery

New Simplified Interface Config

SMU without reboot

Nutanix VMM Integration

Shorter APIC upgrade time

Azure Stack HCI

and more

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

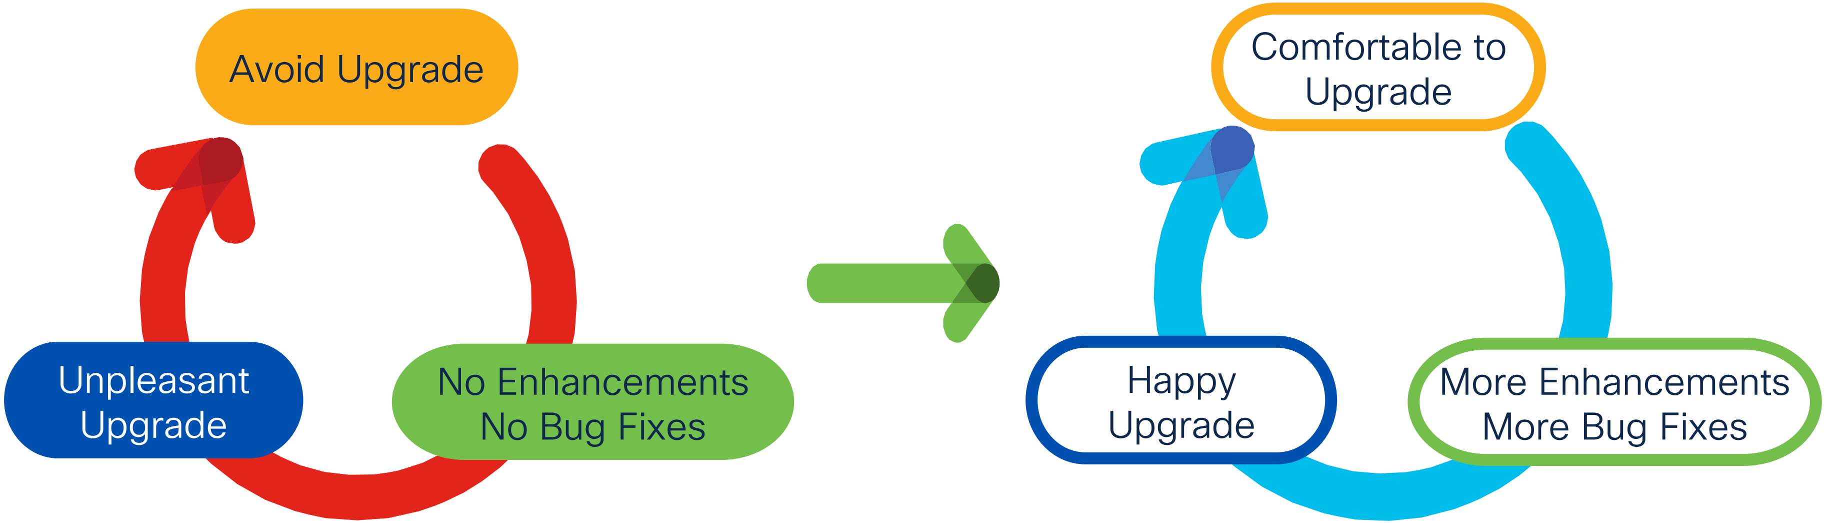

ACI Firmware Upgrade Types

ACI Firmware Upgrade Types

Regular Upgrade

Software Maintenance Upgrade (SMU)

EPLD/FPGA Upgrade (Only Switches)

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

1 APIC Upgrade

2 Switch Upgrade (through APIC)

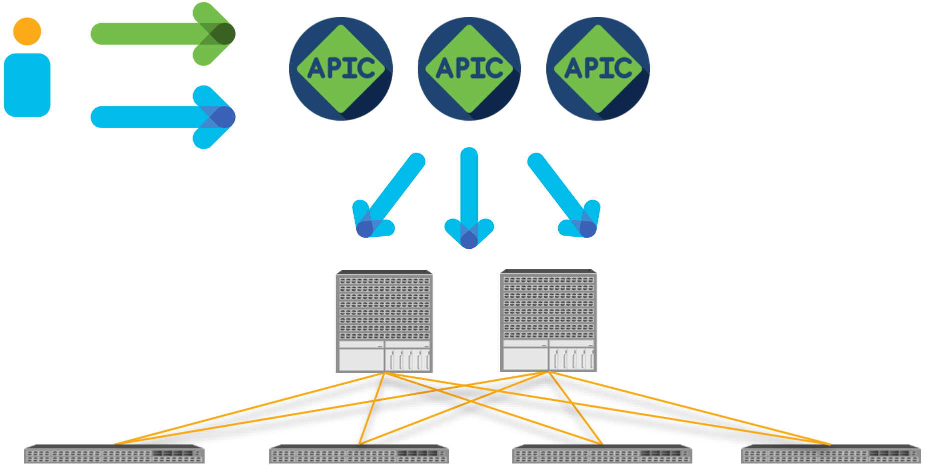

Different versions in the same fabric??

TIP

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.

Supported Operations with different switch versions

✓ 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)

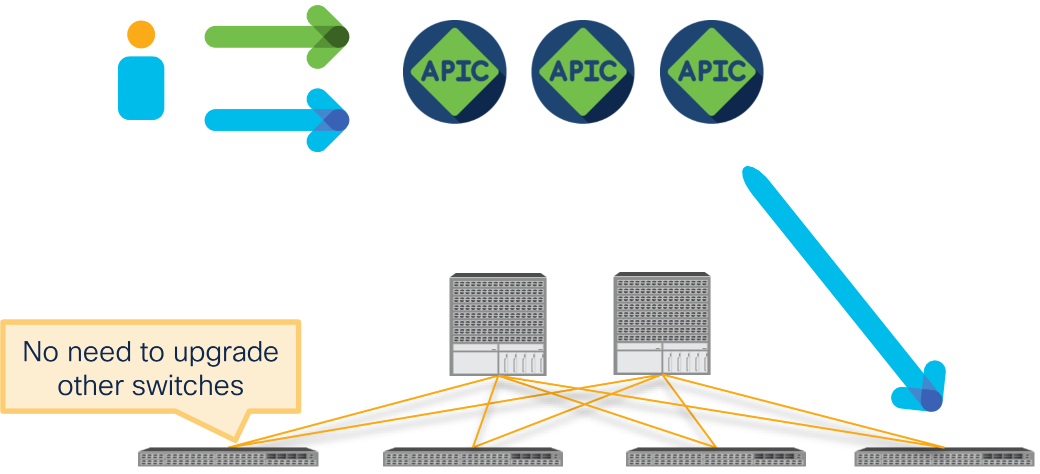

No need to upgrade other switches

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

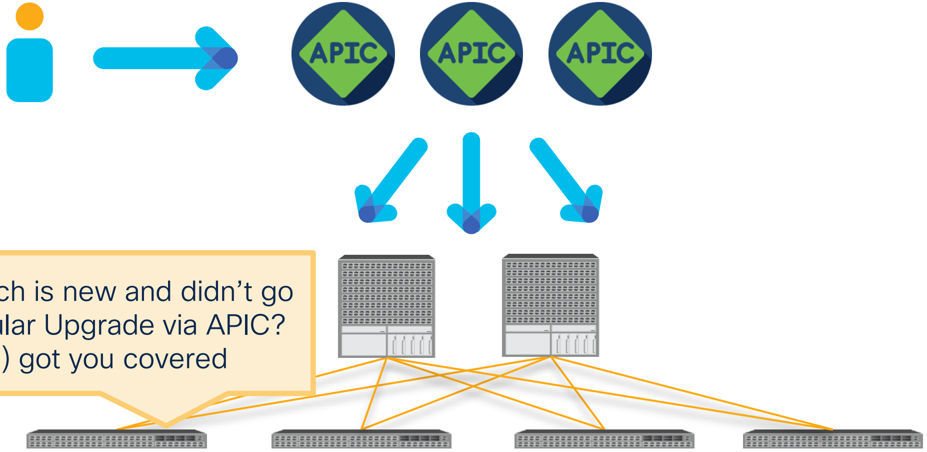

What if a switch is new and didn't go through Regular Upgrade via APIC?

➢ 5.2(1) got you covered

APIC
Upgrade Architecture

Note: for 4.0 or newer APICs

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

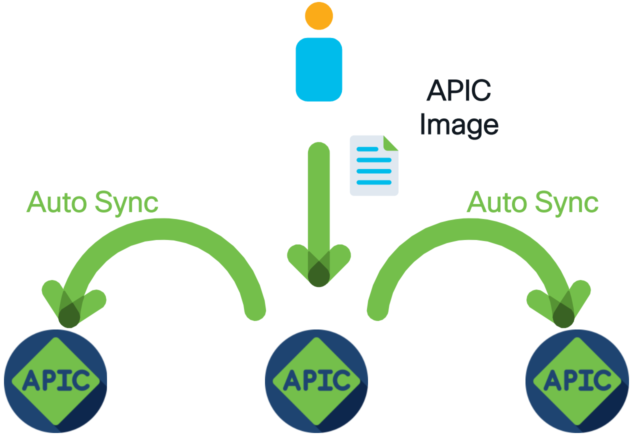

APIC Upgrade Architecture

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)

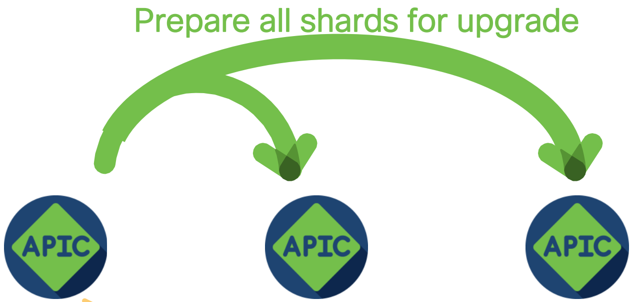

Each shard has 3 replicas across APICs.
Prepare all replicas for upgrade.

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.

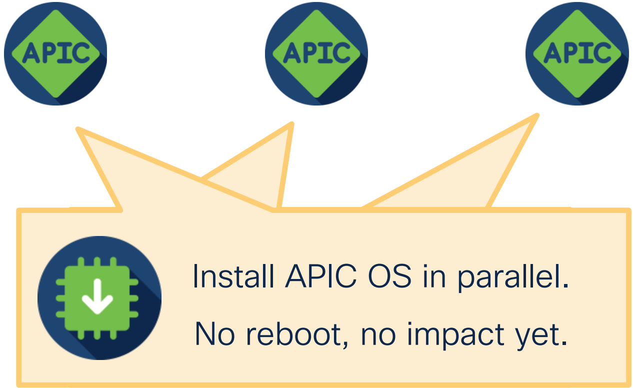

Install APIC OS in parallel.

No reboot, no impact yet.

APIC Upgrade Architecture

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)

Image Upload

Trigger

Install

Data Conversion & Reboot

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

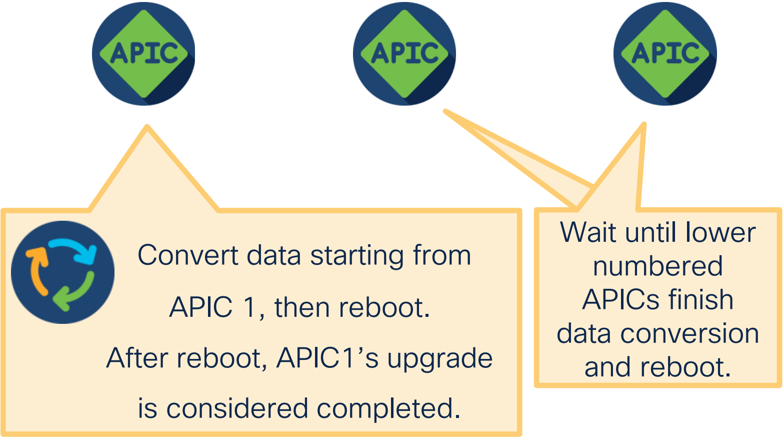

Convert data starting from APIC 1, then reboot.

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

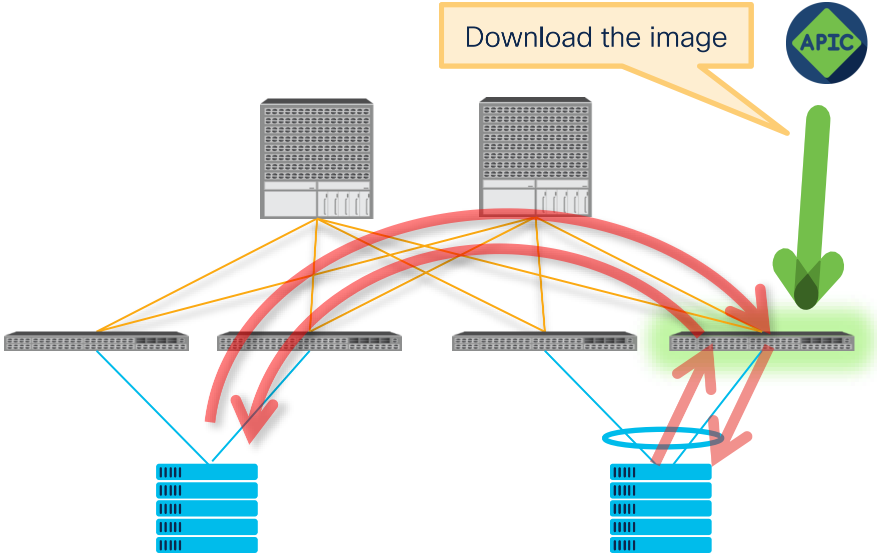

ACI Switch Upgrade Flow

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

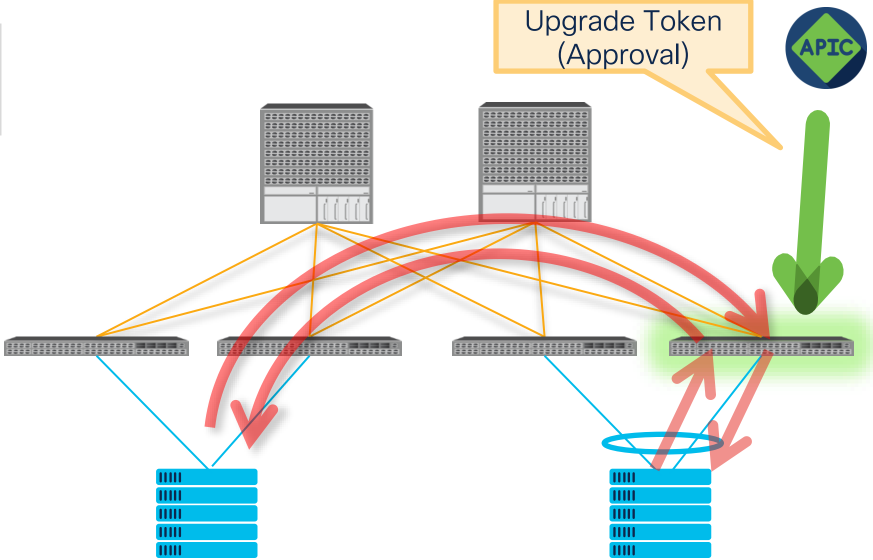

ACI Switch Upgrade Flow

- Image Download
- Queuing
- **Preparation**
  - The switch extracts the image.
  - The switch sets the boot var and so on.
- Reboot
- Boot Up

No Traffic Impact

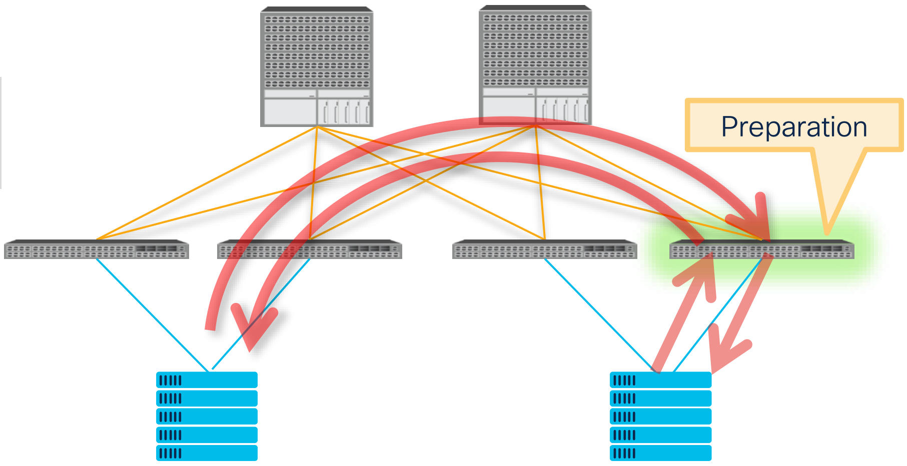

ACI Switch Upgrade Flow

Image Download

Queuing

Preparation

Reboot

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

Boot Up

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

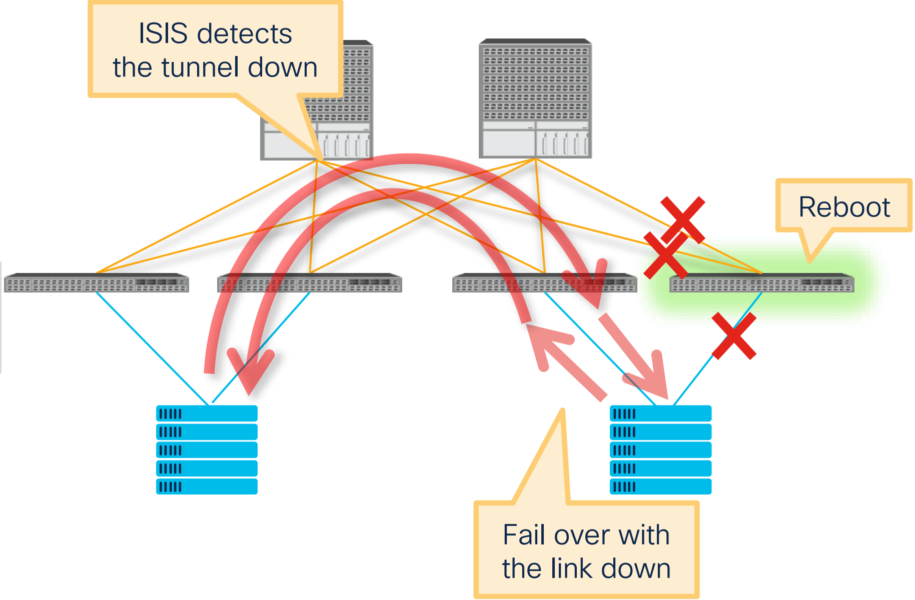

ISIS detects the tunnel down

Reboot

Fail over with the link down

ACI Switch Upgrade Flow

Image Download

Queuing

Preparation

Reboot

Boot Up

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

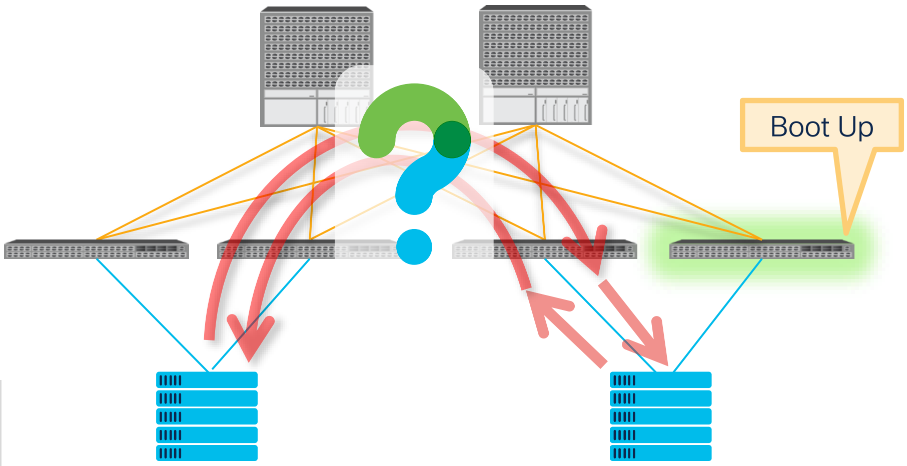

ACI Switch Upgrade Flow (Boot Up Sequence)

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

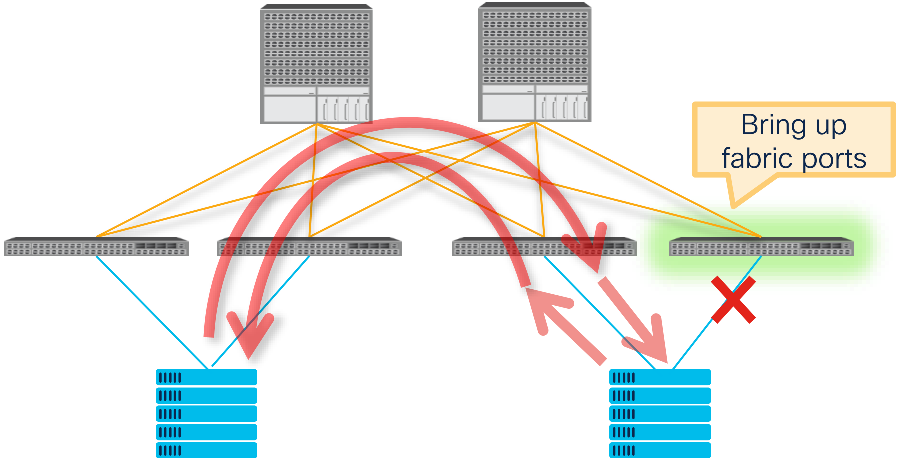

ACI Switch Upgrade Flow (Boot Up Sequence)

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

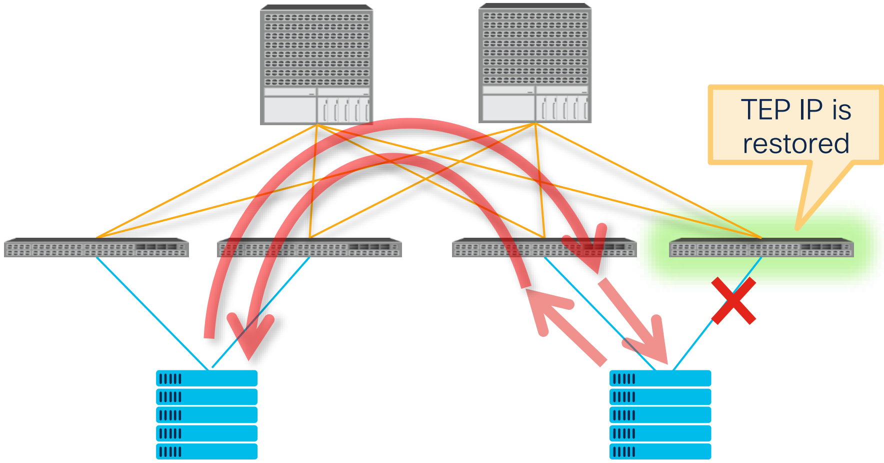

ACI Switch Upgrade Flow (Boot Up Sequence)

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

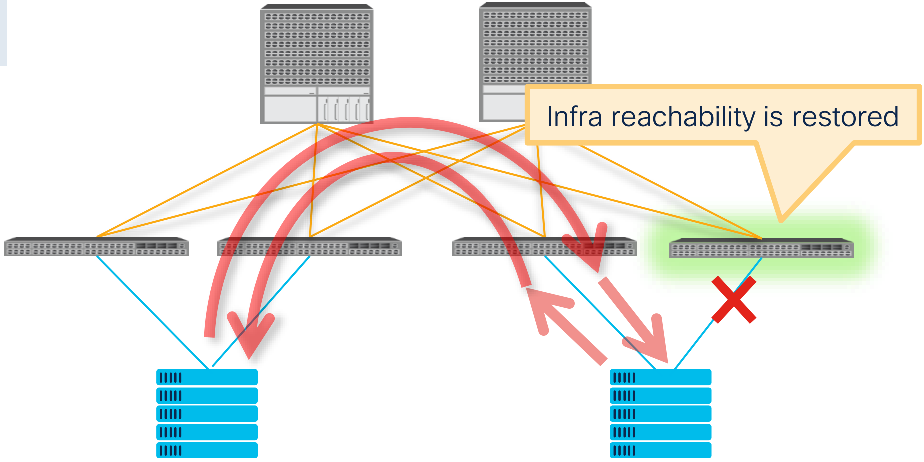

Infra reachability is restored

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

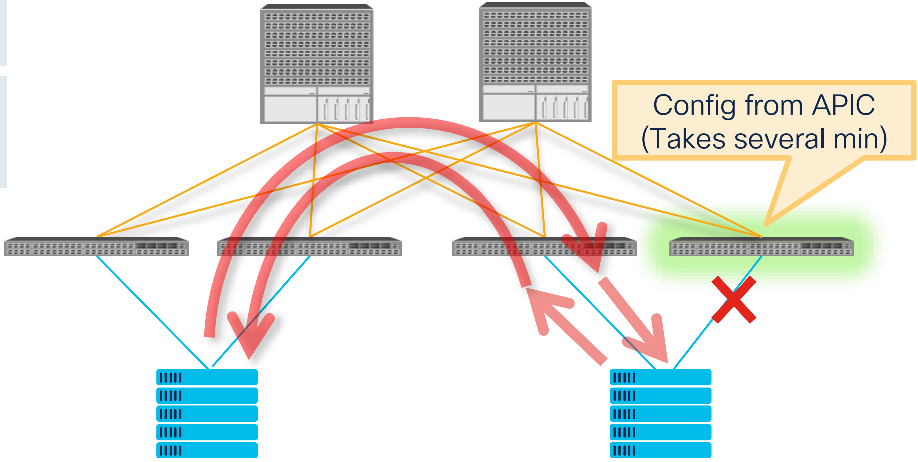

ACI Switch Upgrade Flow (Boot Up Sequence)

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

07

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

No Traffic Flow Change

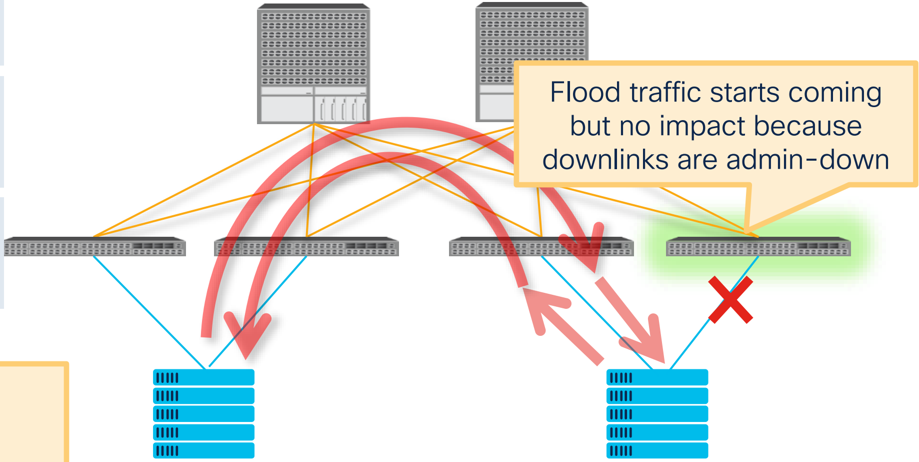

Flood traffic starts coming but no impact because downlinks are admin-down

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

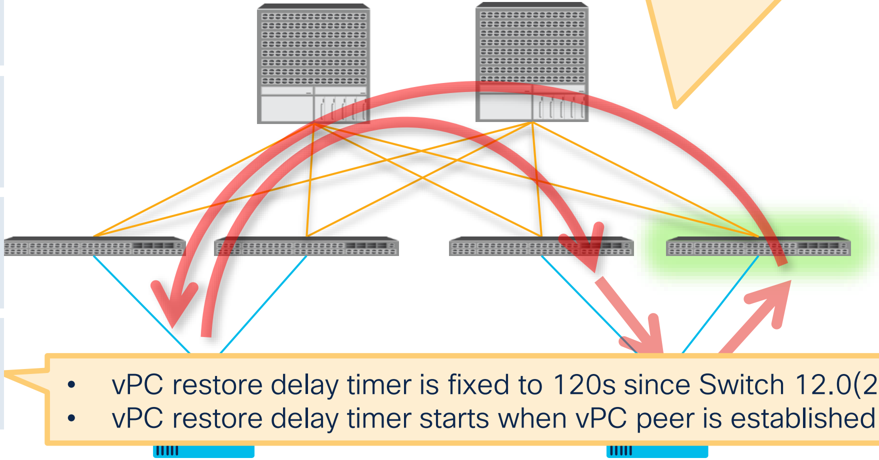

Ready to receive traffic
- VLANs are deployed
  - For VMM, depends on Resolution Immediacy
- Contracts are deployed
  - Depends on Deployment Immediacy
- Spine-Proxy is ready
- Flood handling (FTAG) is ready

- 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

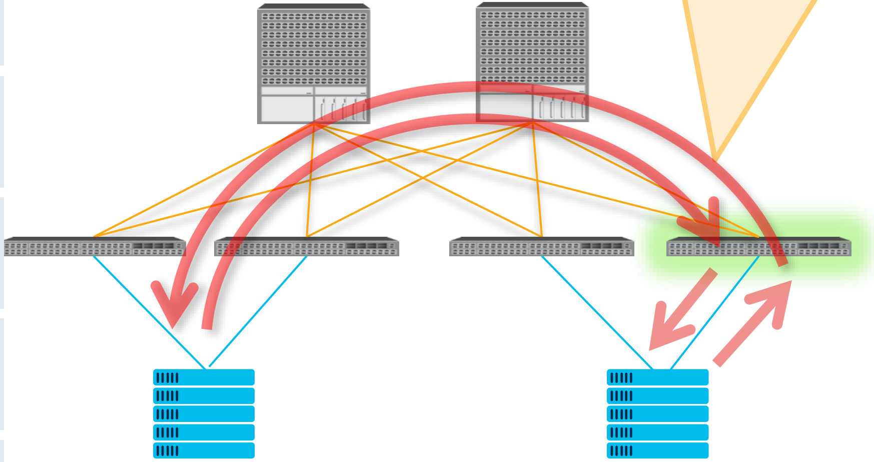

Traffic flow is back to the previous status

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

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

The rest is the same as without graceful option.

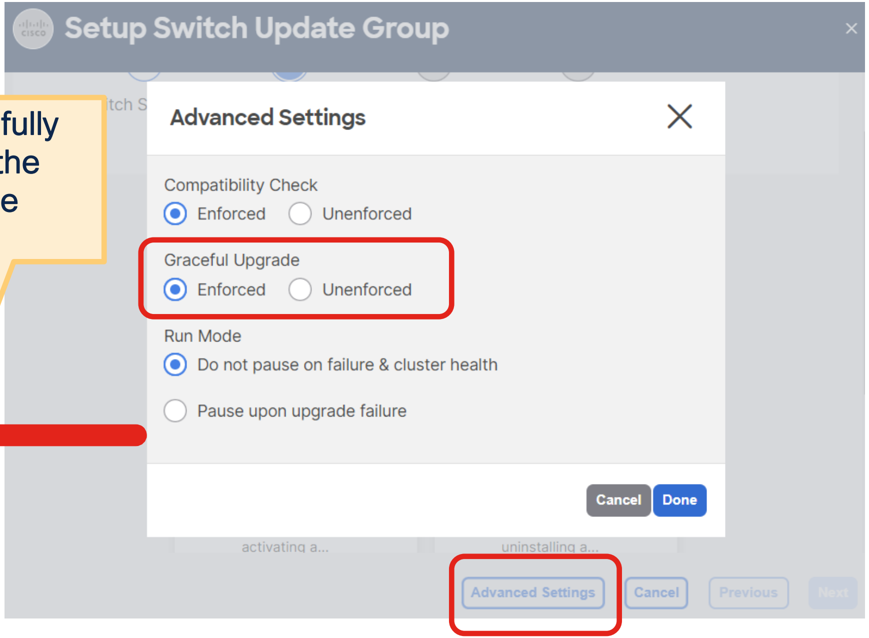

Setup Switch Update Group

Advanced Settings

Compatibility Check
● Enforced ○ Unenforced

Graceful Upgrade
● Enforced ○ Unenforced

Run Mode
● Do not pause on failure & cluster health
○ Pause upon upgrade failure

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

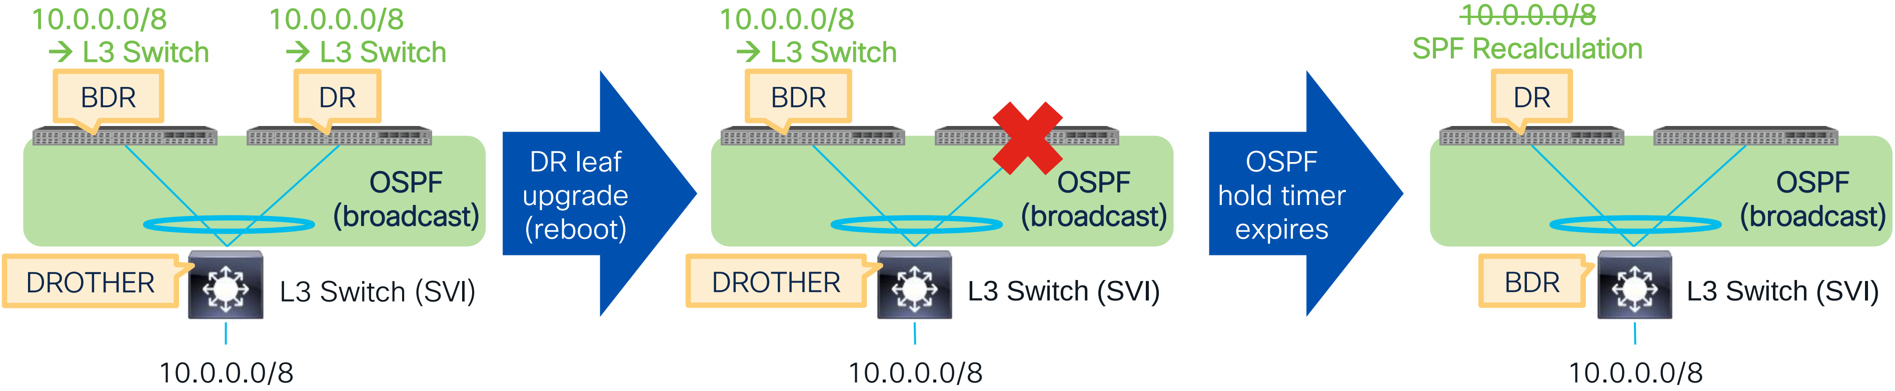

10.0.0.0/8
→ L3 Switch

10.0.0.0/8
→ L3 Switch

BDR

DR

OSPF
(broadcast)

DROTHER

L3 Switch (SVI)

10.0.0.0/8

DR leaf
upgrade
(reboot)

10.0.0.0/8
→ L3 Switch

BDR

OSPF
(broadcast)

DROTHER

L3 Switch (SVI)

10.0.0.0/8

OSPF
hold timer
expires

~~10.0.0.0/8~~
SPF Recalculation

DR

OSPF
(broadcast)

BDR

L3 Switch (SVI)

10.0.0.0/8

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

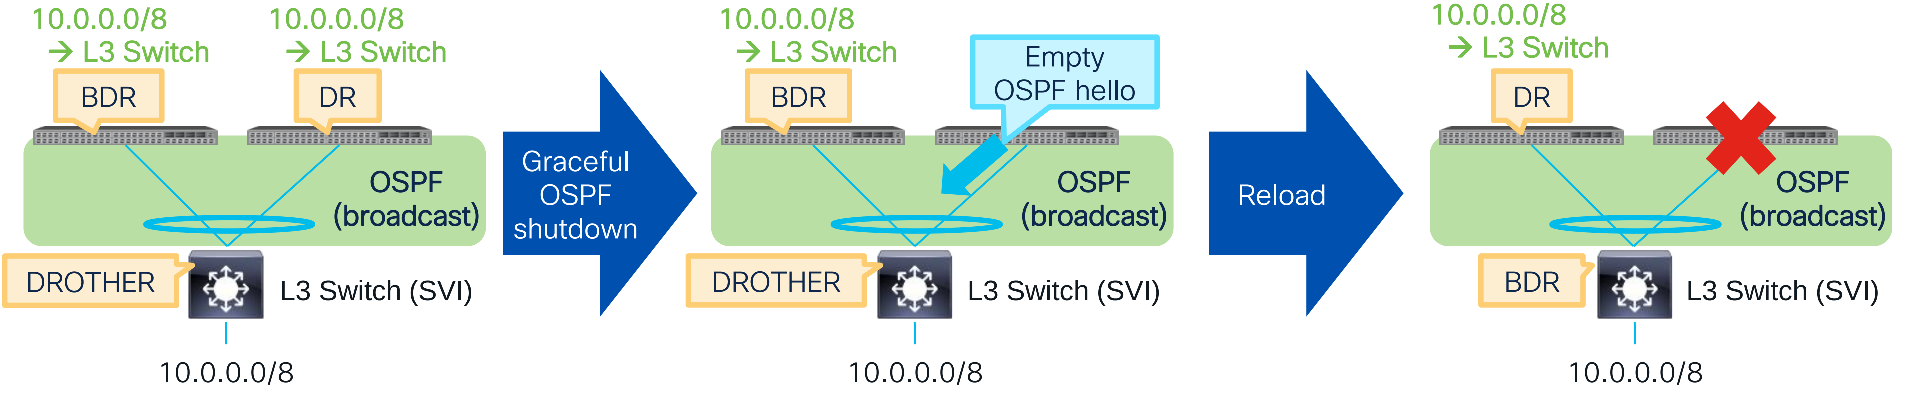

No Traffic Impact

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

TIP

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

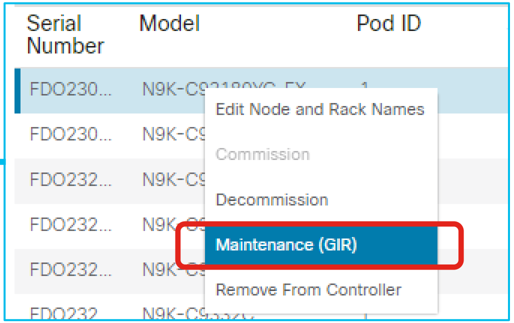

Serial Number	Model	Pod ID
FDO230...	N9K-C93180YC-FX	1
FDO230...	N9K-C9	
FDO232...	N9K-C9	
FDO232...	N9K-C9	
FDO232...	N9K-C9	
FDO232	N9K-C9	

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.

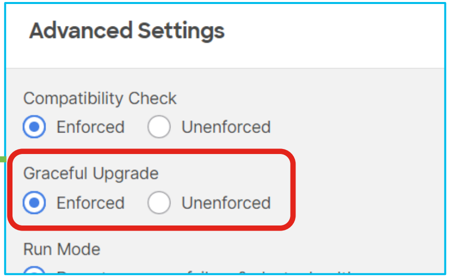

Advanced Settings

Compatibility Check
● Enforced ○ Unenforced

Graceful Upgrade
● Enforced ○ Unenforced

Run Mode

Auto Firmware Update

Auto Firmware Update for APIC

6.0(2) NEW

Use Case 1: APIC Replacement

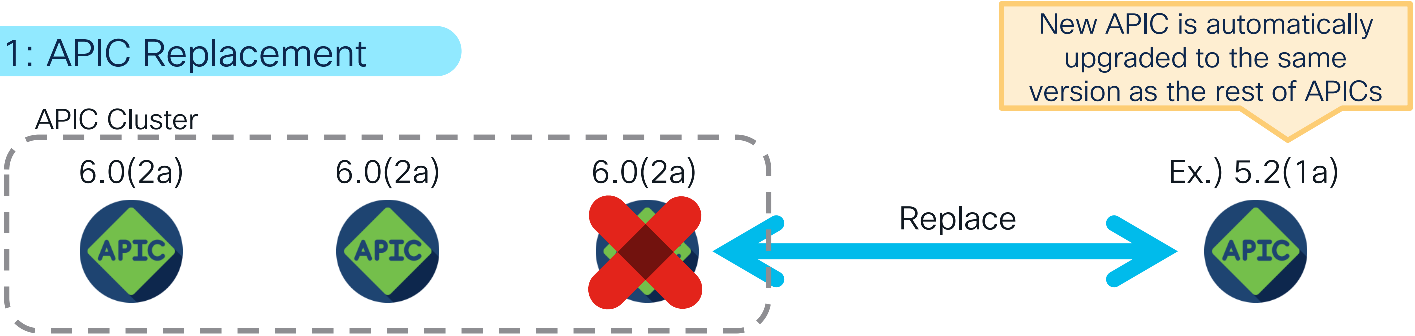

New APIC is automatically upgraded to the same version as the rest of APICs

Use Case 2: Cluster Expansion

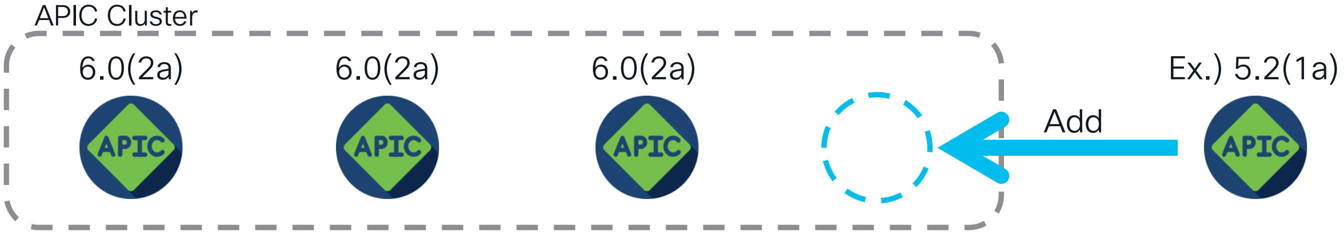

Auto Firmware Update for Switches

Enforcing Version Consistency

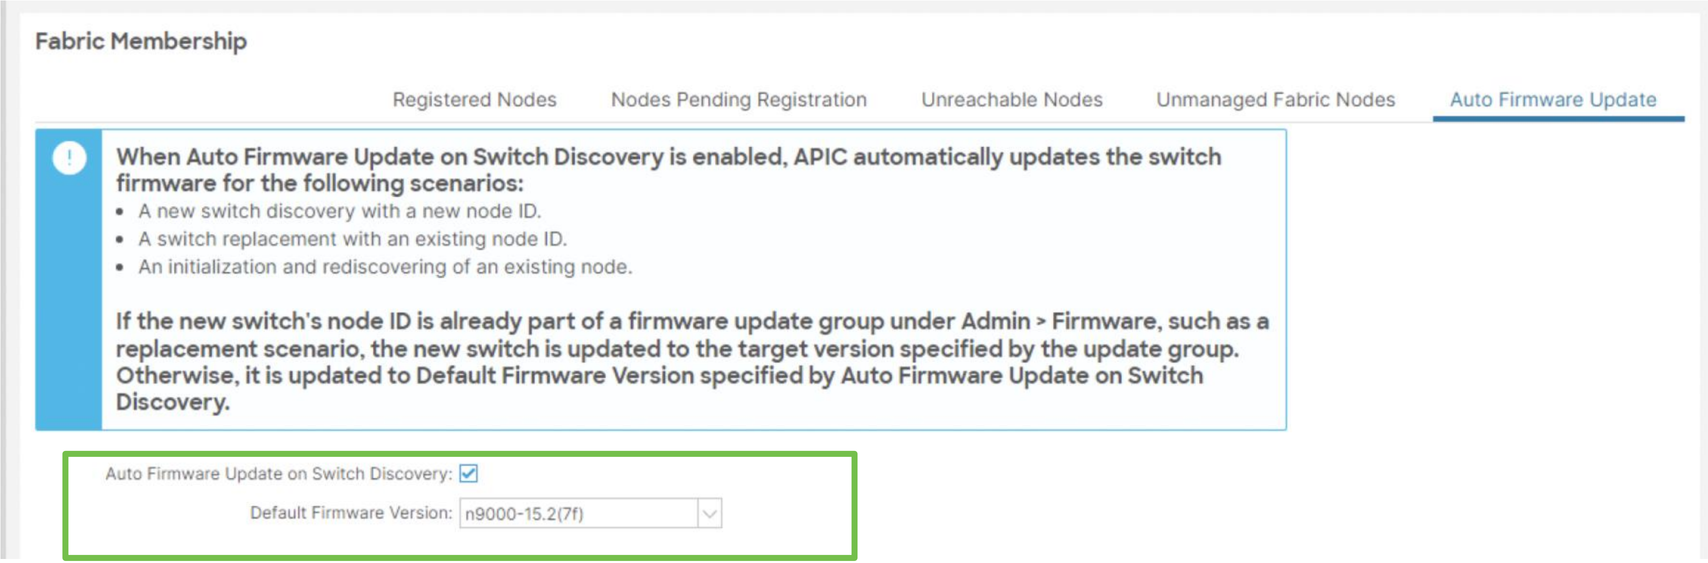

When Auto Firmware Update on Switch Discovery is enabled, APIC automatically updates the switch firmware for the following scenarios:

- A new switch discovery with a new node ID.
- A switch replacement with an existing node ID.
- An initialization and rediscovering of an existing node.

If the new switch's node ID is already part of a firmware update group under Admin > Firmware, such as a replacement scenario, the new switch is updated to the target version specified by the update group. Otherwise, it is updated to Default Firmware Version specified by Auto Firmware Update on Switch Discovery.

Auto Firmware Update on Switch Discovery: ☑

Default Firmware Version: n9000-15.2(7f)

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

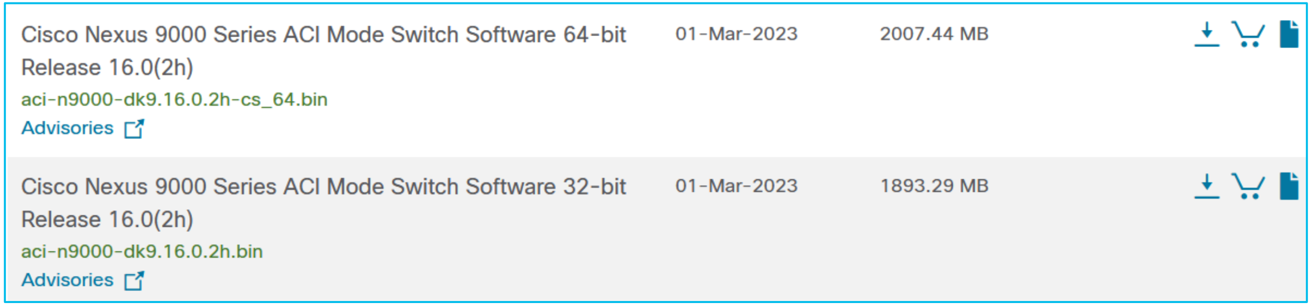

Cisco Nexus 9000 Series ACI Mode Switch Software 64-bit Release 16.0(2h) aci-n9000-dk9.16.0.2h-cs_64.bin Advisories	01-Mar-2023	2007.44 MB	
Cisco Nexus 9000 Series ACI Mode Switch Software 32-bit Release 16.0(2h) aci-n9000-dk9.16.0.2h.bin Advisories	01-Mar-2023	1893.29 MB	

Why?

Scalability

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

Which Image?

Depends on the memory size

24G or less -> 32 bit
Otherwise -> 64 bit

Note: Fixed per switch model in ACI Switch 16.0(2)

Feature Differences?

None

Features (scale) are handled based on switch model.
No feature differences specific to image type (32/64 bits).

cisco Live!

Upgrade Procedure

Regular One (ex. 4.2 -> 5.2)

0 Upload Images

1. Upload target APIC image to APICs
2. Upload target switch images to APICs

1 APIC Upgrade

1. Upgrade APIC cluster

2 Switch Upgrade (through APIC)

1. Download the switch images from APICs to switches
2. Upgrade switches

Upgrade Procedure

Specific to pre-6.0(2) -> 6.0(2) or later (ex. 5.2.7 -> 6.0.2)

except 5.2(8), 5.3

0 Upload Images

1. Upload target APIC image to APICs

~~2. Upload target switch images to APICs~~

Do not upload switch images (16.0(2) or later) until APICs are upgraded.

If your APICs are on 5.2(8), 5.3, no need to worry about this

1 APIC Upgrade

1. Upgrade APIC cluster

2 Switch Upgrade (through APIC)

1. Upload target switch images to APICs
2. Download the switch images from APICs to switches
3. Upgrade switches

ACI Upgrade 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-upgrade-downgrade-architecture.html#Cisco_Reference.dita_22480abb-4138-416b-8dd5-ecde23f707b4

Upgrade Procedure

Anything -> 6.0(2) or later (ex. 5.2.7->6.0.4, 6.0.2->6.0.4 etc.)

0 Upload Images

1. Upload target APIC image to APICs
2. Upload target switch images to APICs

1 APIC Upgrade

1. Upgrade APIC cluster

2 Switch Upgrade (through APIC)

1. Upload target switch images to APICs
2. Download the switch images from APICs to switches
3. Upgrade switches

In either case, upload BOTH 32- and 64-bit images to APICs.
APICs will pick the appropriate image for each switch.

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

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”

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

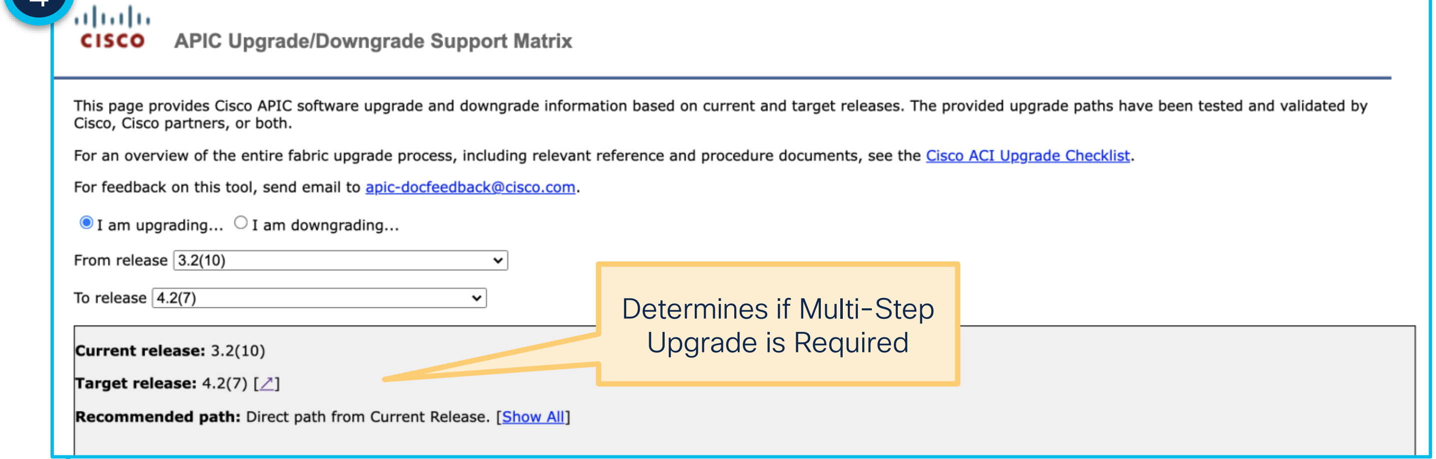

APIC Upgrade/Downgrade Support Matrix

This page provides Cisco APIC software upgrade and downgrade information based on current and target releases. The provided upgrade paths have been tested and validated by Cisco, Cisco partners, or both.

For an overview of the entire fabric upgrade process, including relevant reference and procedure documents, see the Cisco ACI Upgrade Checklist.

For feedback on this tool, send email to apic-docfeedback@cisco.com.

● I am upgrading... ○ I am downgrading...

From release 3.2(10)

To release 4.2(7)

Current release: 3.2(10)

Target release: 4.2(7) [↗]

Recommended path: Direct path from Current Release. [Show All]

Determines if Multi-Step Upgrade is Required

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

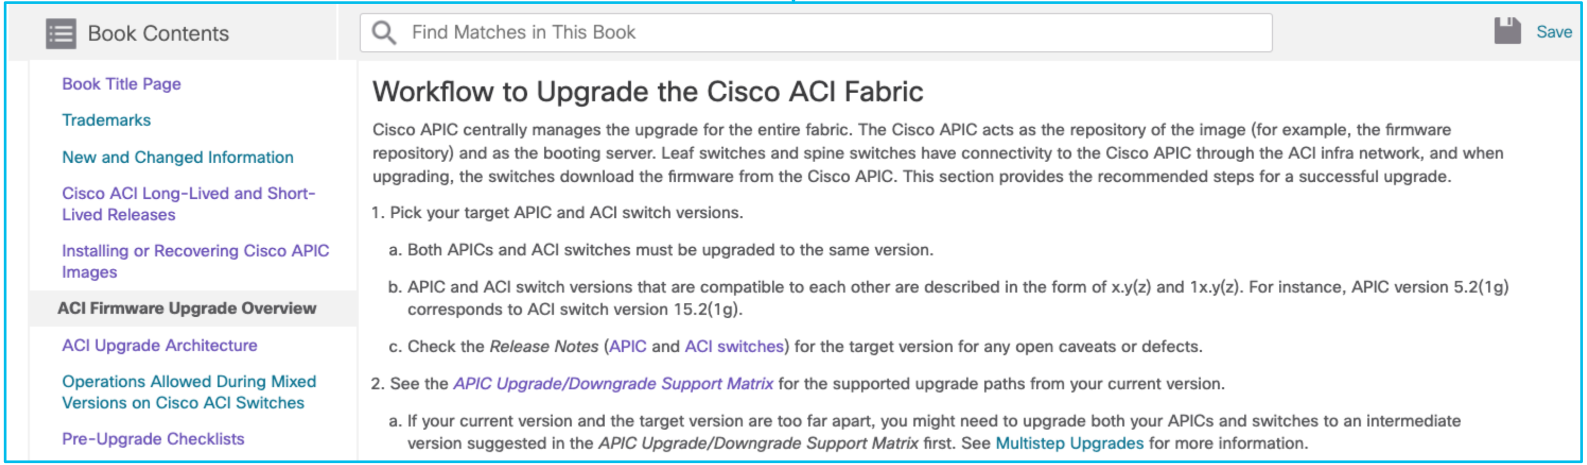
Book Contents
Find Matches in This Book
Save
Book Title Page
Trademarks
New and Changed Information
Cisco ACI Long-Lived and Short-Lived Releases
Installing or Recovering Cisco APIC Images
ACI Firmware Upgrade Overview
ACI Upgrade Architecture
Operations Allowed During Mixed Versions on Cisco ACI Switches
Pre-Upgrade Checklists

Workflow to Upgrade the Cisco ACI Fabric

Cisco APIC centrally manages the upgrade for the entire fabric. The Cisco APIC acts as the repository of the image (for example, the firmware repository) and as the booting server. Leaf switches and spine switches have connectivity to the Cisco APIC through the ACI infra network, and when upgrading, the switches download the firmware from the Cisco APIC. This section provides the recommended steps for a successful upgrade.

1. Pick your target APIC and ACI switch versions.
   a. Both APICs and ACI switches must be upgraded to the same version.
   b. 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).
   c. Check the *Release Notes* (APIC and ACI switches) for the target version for any open caveats or defects.
2. See the *APIC Upgrade/Downgrade Support Matrix* for the supported upgrade paths from your current version.
   a. 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 for more information.

CIMC Version Compatibility

6

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

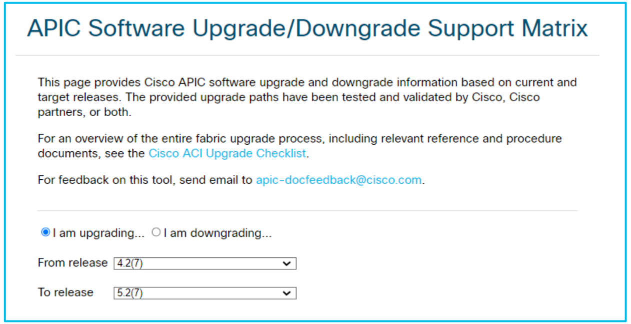

APIC Software Upgrade/Downgrade Support Matrix

This page provides Cisco APIC software upgrade and downgrade information based on current and target releases. The provided upgrade paths have been tested and validated by Cisco, Cisco partners, or both.

For an overview of the entire fabric upgrade process, including relevant reference and procedure documents, see the Cisco ACI Upgrade Checklist.

For feedback on this tool, send email to apic-docfeedback@cisco.com.

● I am upgrading... ○ I am downgrading...

From release 4.2(7)

To release 5.2(7)

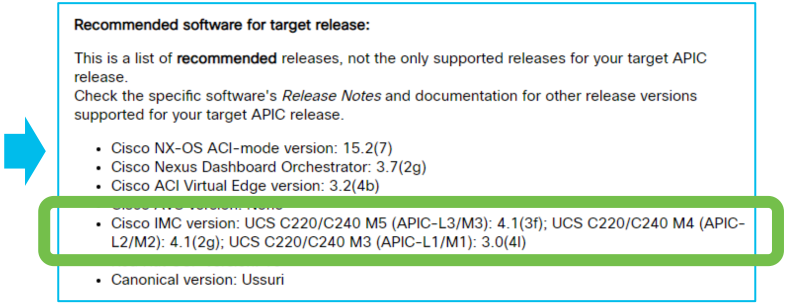

Recommended software for target release:

This is a list of **recommended** releases, not the only supported releases for your target APIC release.
Check the specific software's *Release Notes* and documentation for other release versions supported for your target APIC release.

- Cisco NX-OS ACI-mode version: 15.2(7)
- Cisco Nexus Dashboard Orchestrator: 3.7(2g)
- Cisco ACI Virtual Edge version: 3.2(4b)
- Cisco IMC version: UCS C220/C240 M5 (APIC-L3/M3): 4.1(3f); UCS C220/C240 M4 (APIC-L2/M2): 4.1(2g); UCS C220/C240 M3 (APIC-L1/M1): 3.0(4l)
- Canonical version: Ussuri

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

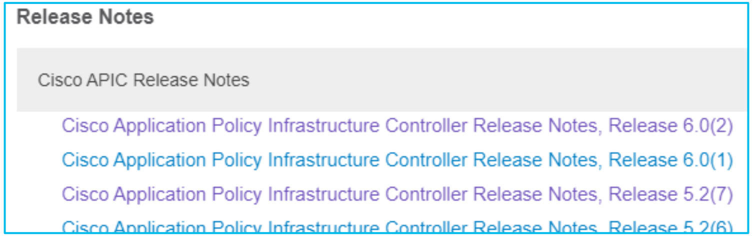

Release Notes

Cisco APIC Release Notes

Cisco Application Policy Infrastructure Controller Release Notes, Release 6.0(2)
Cisco Application Policy Infrastructure Controller Release Notes, Release 6.0(1)
Cisco Application Policy Infrastructure Controller Release Notes, Release 5.2(7)
Cisco Application Policy Infrastructure Controller Release Notes, Release 5.2(6)

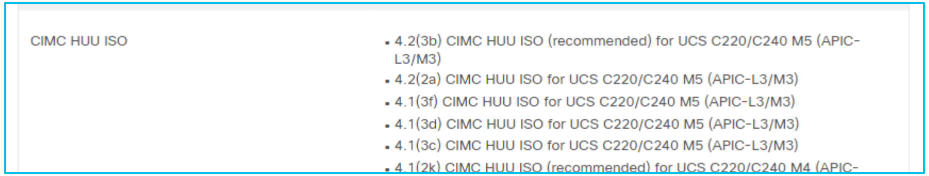

CIMC HUU ISO	- 4.2(3b) CIMC HUU ISO (recommended) for UCS C220/C240 M5 (APIC-L3/M3) - 4.2(2a) CIMC HUU ISO for UCS C220/C240 M5 (APIC-L3/M3) - 4.1(3f) CIMC HUU ISO for UCS C220/C240 M5 (APIC-L3/M3) - 4.1(3d) CIMC HUU ISO for UCS C220/C240 M5 (APIC-L3/M3) - 4.1(3c) CIMC HUU ISO for UCS C220/C240 M5 (APIC-L3/M3) - 4.1(2k) CIMC HUU ISO (recommended) for UCS C220/C240 M4 (APIC-

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”

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

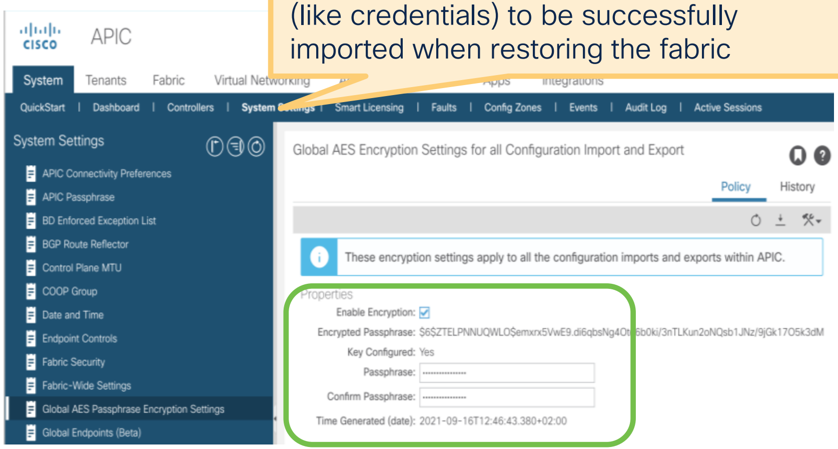

Global AES Encryption Settings for all Configuration Import and Export

These encryption settings apply to all the configuration imports and exports within APIC.

Properties

Enable Encryption: ☑

Encrypted Passphrase: $6$ZTELPNNUQWLO$emxrx5VwE9.di6qbsNg4Ot6b0ki/3nTLKun2oNQsb1JNz/9jGk17O5k3dM

Key Configured: Yes

Passphrase: ················

Confirm Passphrase: ················

Time Generated (date): 2021-09-16T12:46:43.380+02:00

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

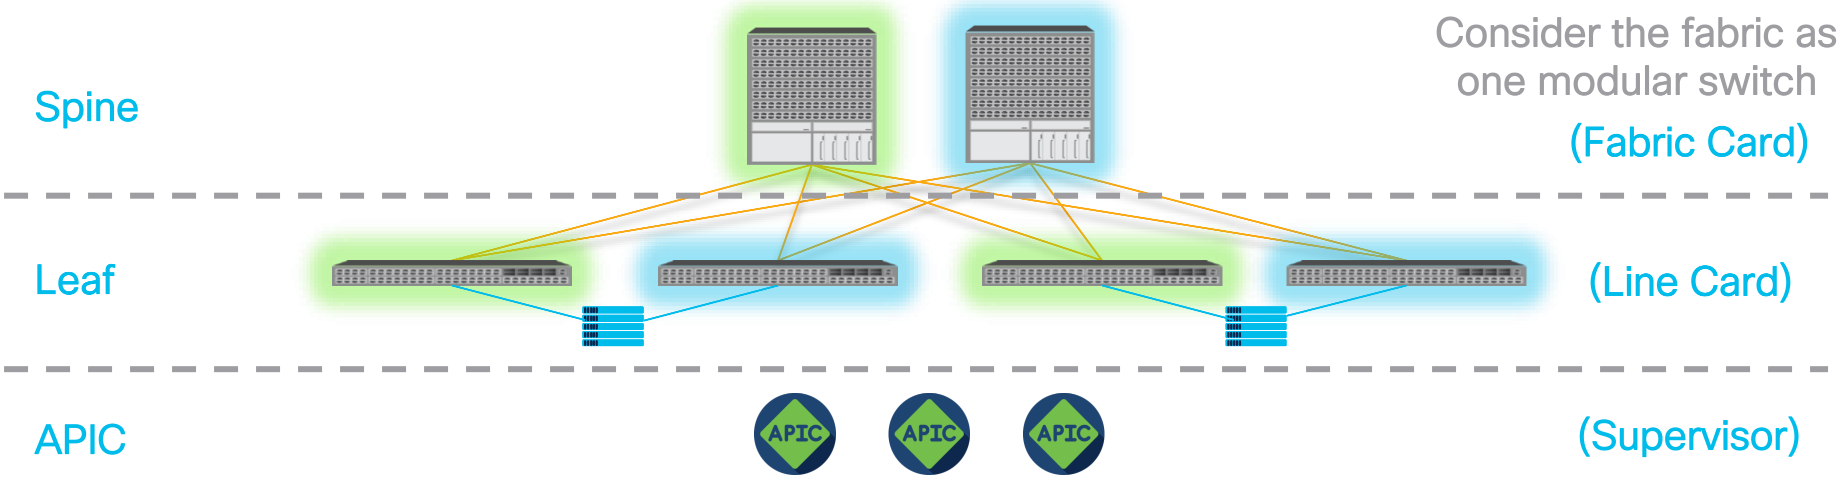

ACI is a solution to manage multiple switches as if it's one huge switch

- 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

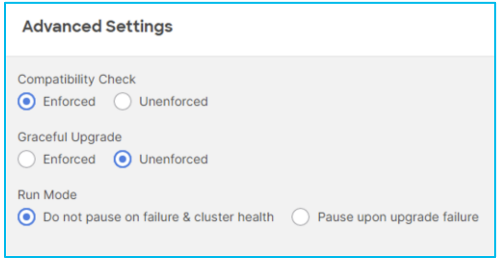

Advanced Settings

Compatibility Check
● Enforced ○ Unenforced

Graceful Upgrade
○ Enforced ● Unenforced

Run Mode
● Do not pause on failure & cluster health ○ Pause upon upgrade failure

Rule of Thumb

Change defaults only when you must.

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

- **<u>Graceful Upgrade (default: Unenforced)</u>**
  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

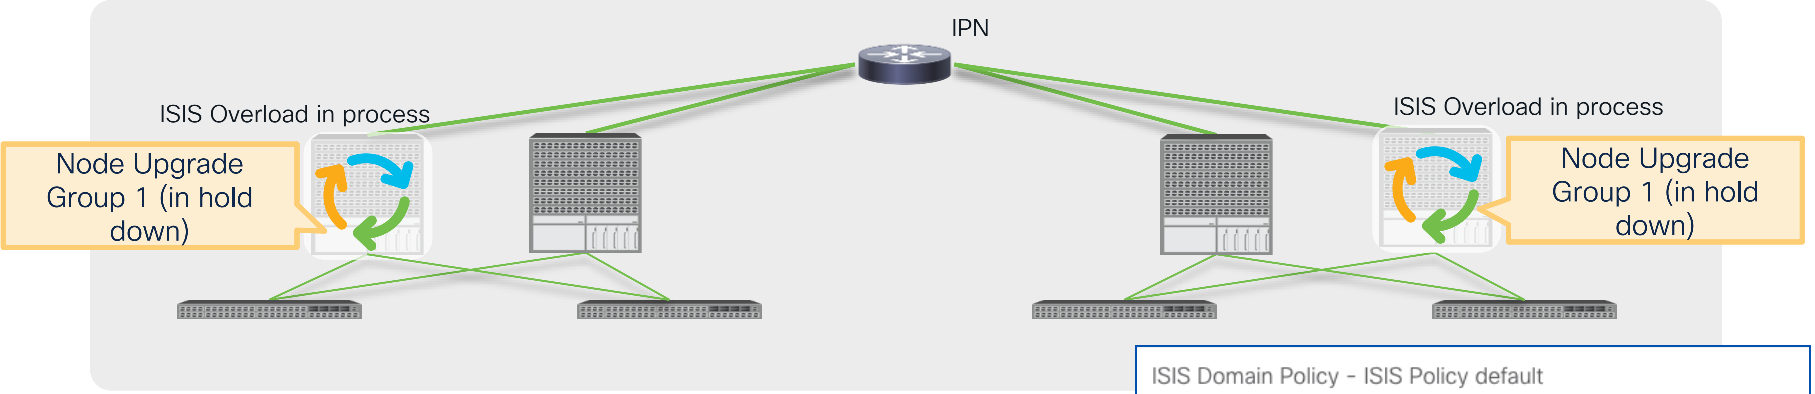
IPN
ISIS Overload in process
Node Upgrade Group 1 (in hold down)
ISIS Overload in process
Node Upgrade Group 1 (in hold down)

- 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.

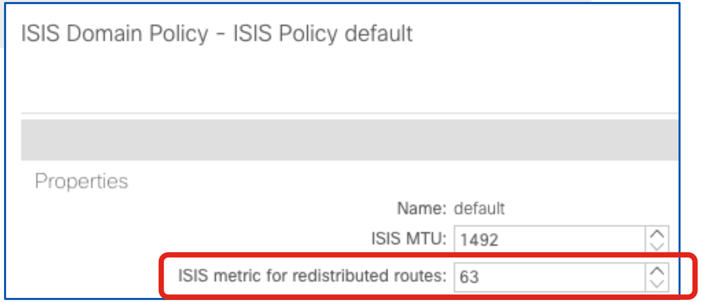

ISIS Domain Policy - ISIS Policy default

Properties

Name: default

ISIS MTU: 1492

ISIS metric for redistributed routes: 63

Settings > ISIS Policy (Default Config)

Helpful Tips for Multi-Pod / Multi-Site

ISIS Metric Policy Configuration

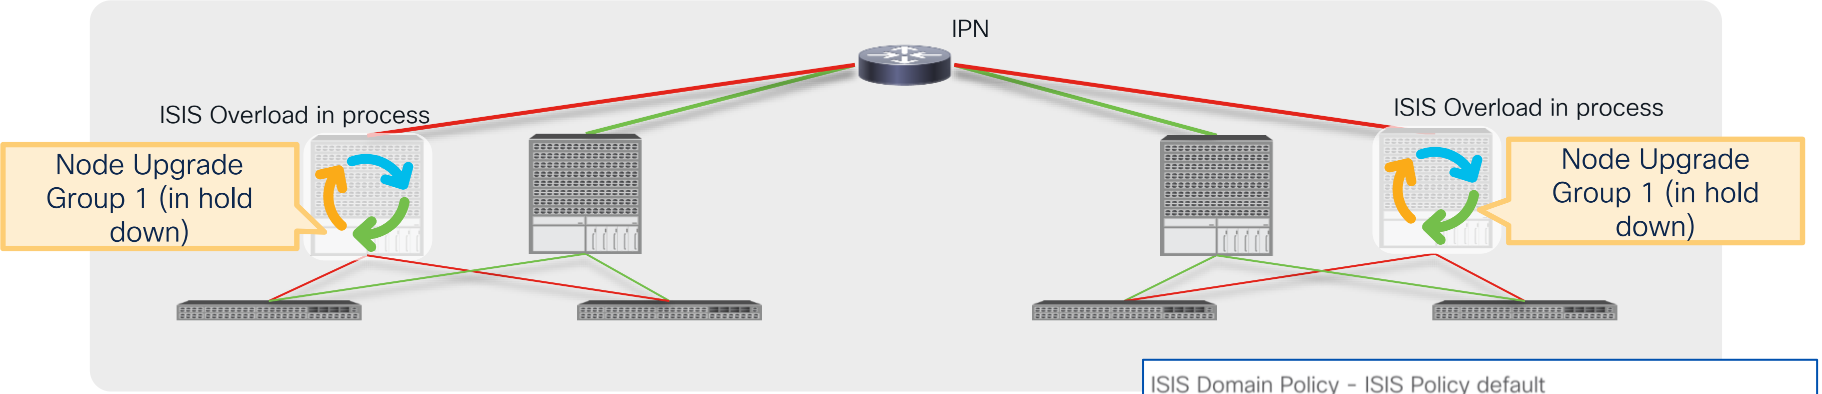
IPN
ISIS Overload in process
Node Upgrade Group 1 (in hold down)
ISIS Overload in process
Node Upgrade Group 1 (in hold down)

- 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.

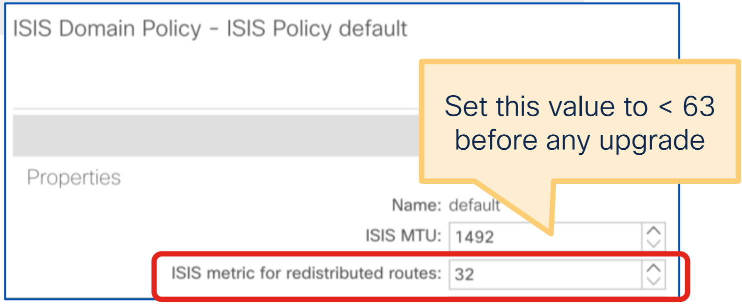

ISIS Domain Policy - ISIS Policy default

Properties

Name: default

ISIS MTU: 1492

ISIS metric for redistributed routes: 32

Set this value to < 63 before any upgrade

Settings > ISIS Policy

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?

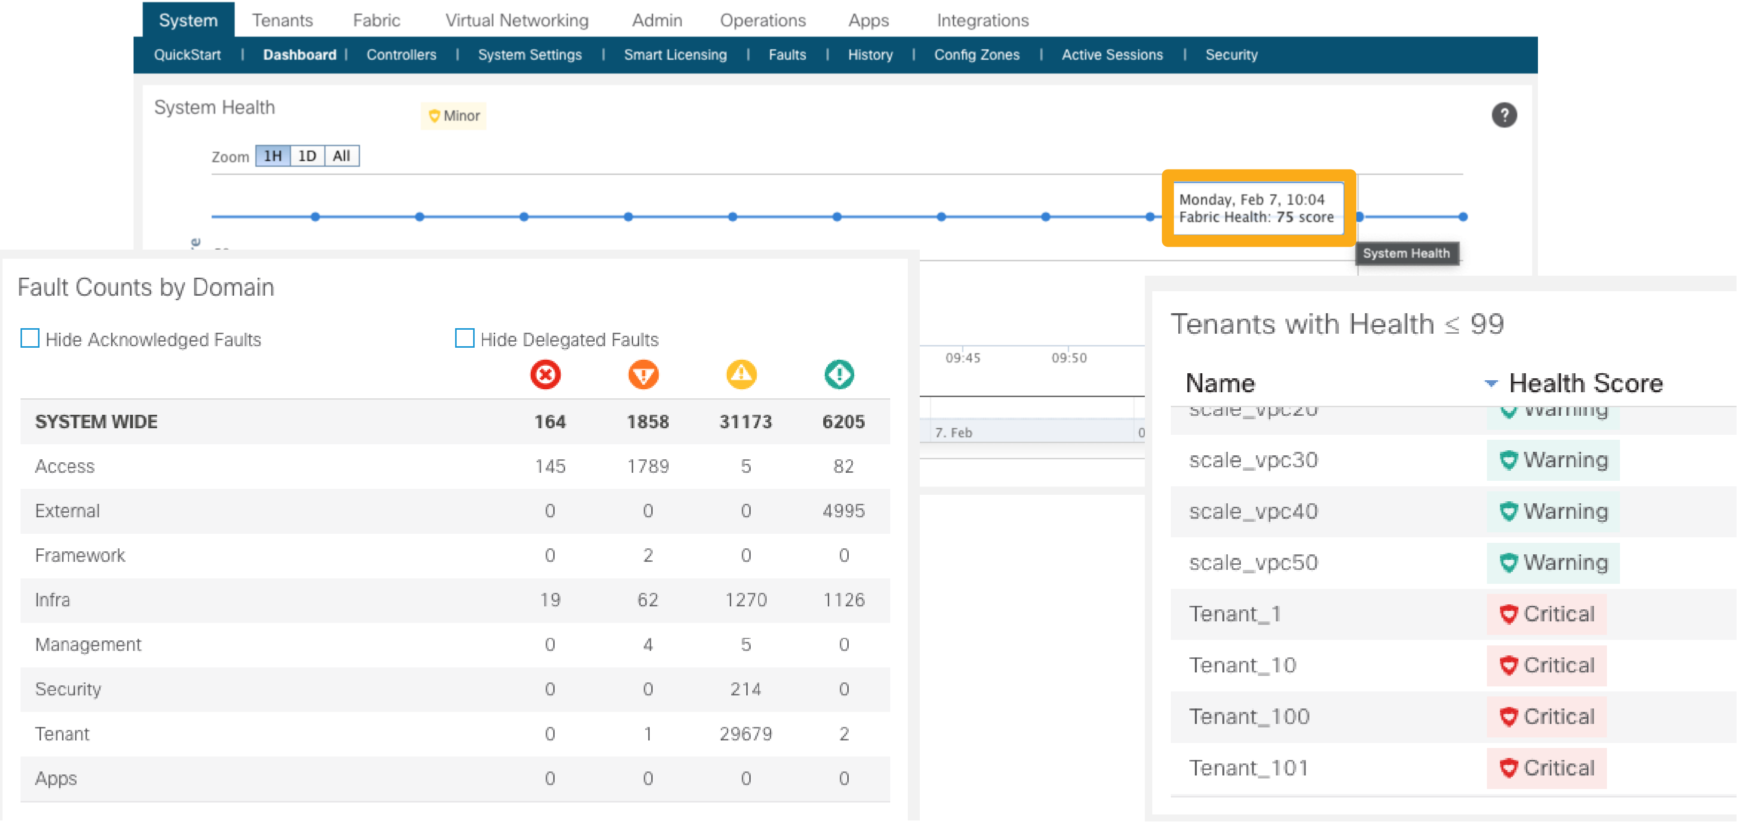

Fault Counts by Domain

□ 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.

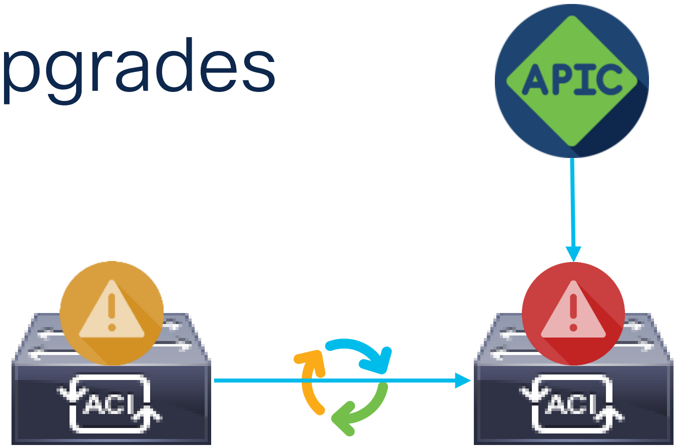

Faults raised but functioning normally.

After upgrade, previous working config can be changed to “faulted” 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

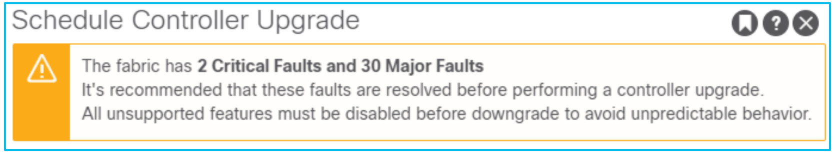

Schedule Controller Upgrade

The fabric has **2 Critical Faults and 30 Major Faults**
It's recommended that these faults are resolved before performing a controller upgrade.
All unsupported features must be disabled before downgrade to avoid unpredictable behavior.

APIC 4.2(1) – 4.2(3)

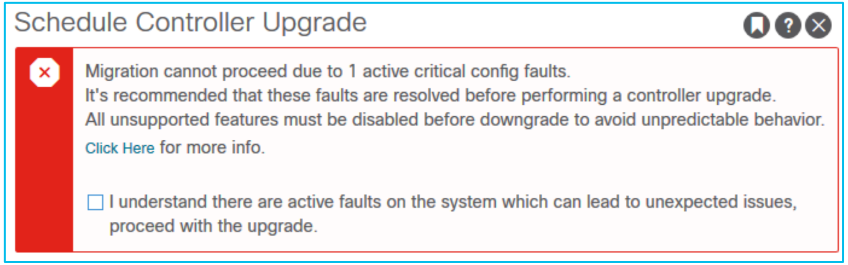

Schedule Controller Upgrade

Migration cannot proceed due to 1 active critical config faults.
It's recommended that these faults are resolved before performing a controller upgrade.
All unsupported features must be disabled before downgrade to avoid unpredictable behavior.
Click Here for more info.

□ I understand there are active faults on the system which can lead to unexpected issues, proceed with the upgrade.

APIC 4.2(4)

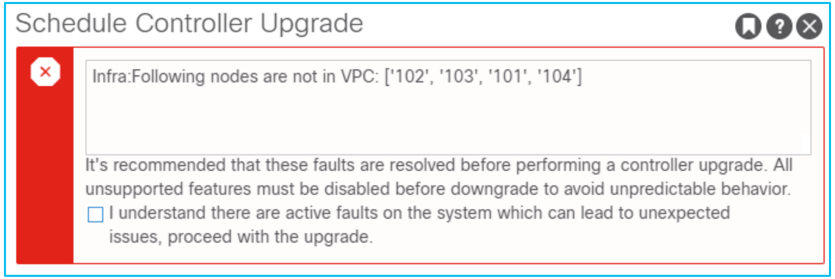

Schedule Controller Upgrade

Infra:Following nodes are not in VPC: ['102', '103', '101', '104']

It's recommended that these faults are resolved before performing a controller upgrade. All unsupported features must be disabled before downgrade to avoid unpredictable behavior.
□ I understand there are active faults on the system which can lead to unexpected issues, proceed with the upgrade.

- 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 – Script

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

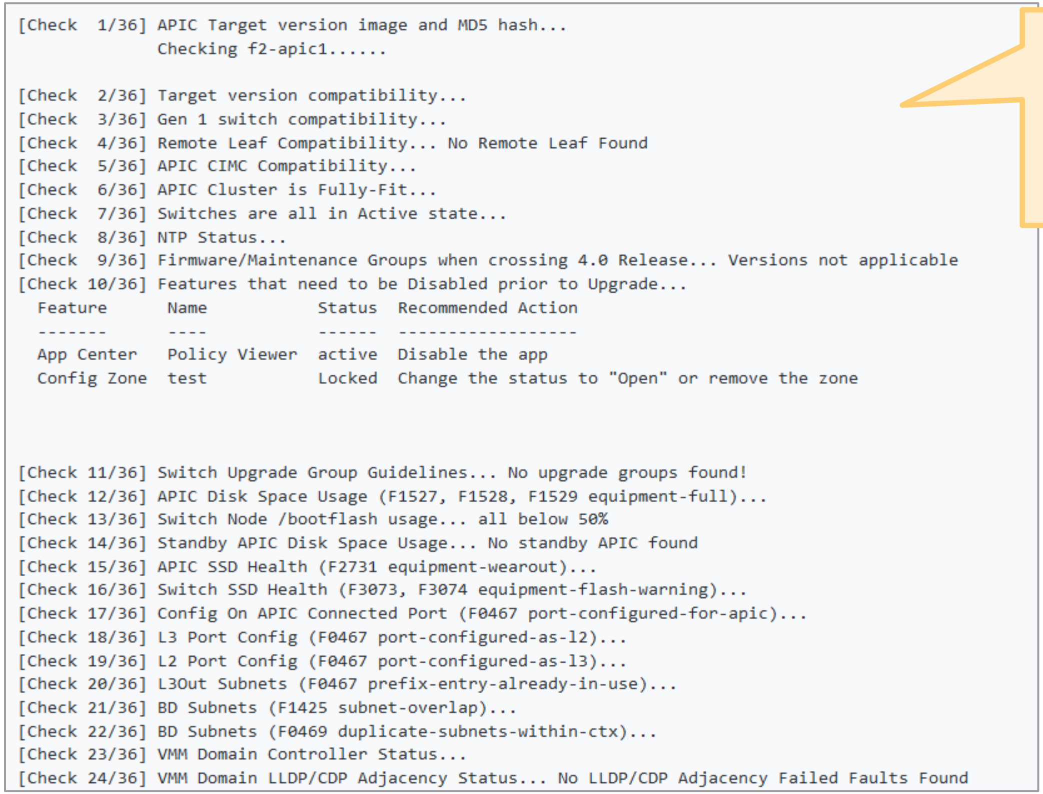

```
[Check  1/36] APIC Target version image and MD5 hash...
              Checking f2-apic1......

[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 (Preferred)

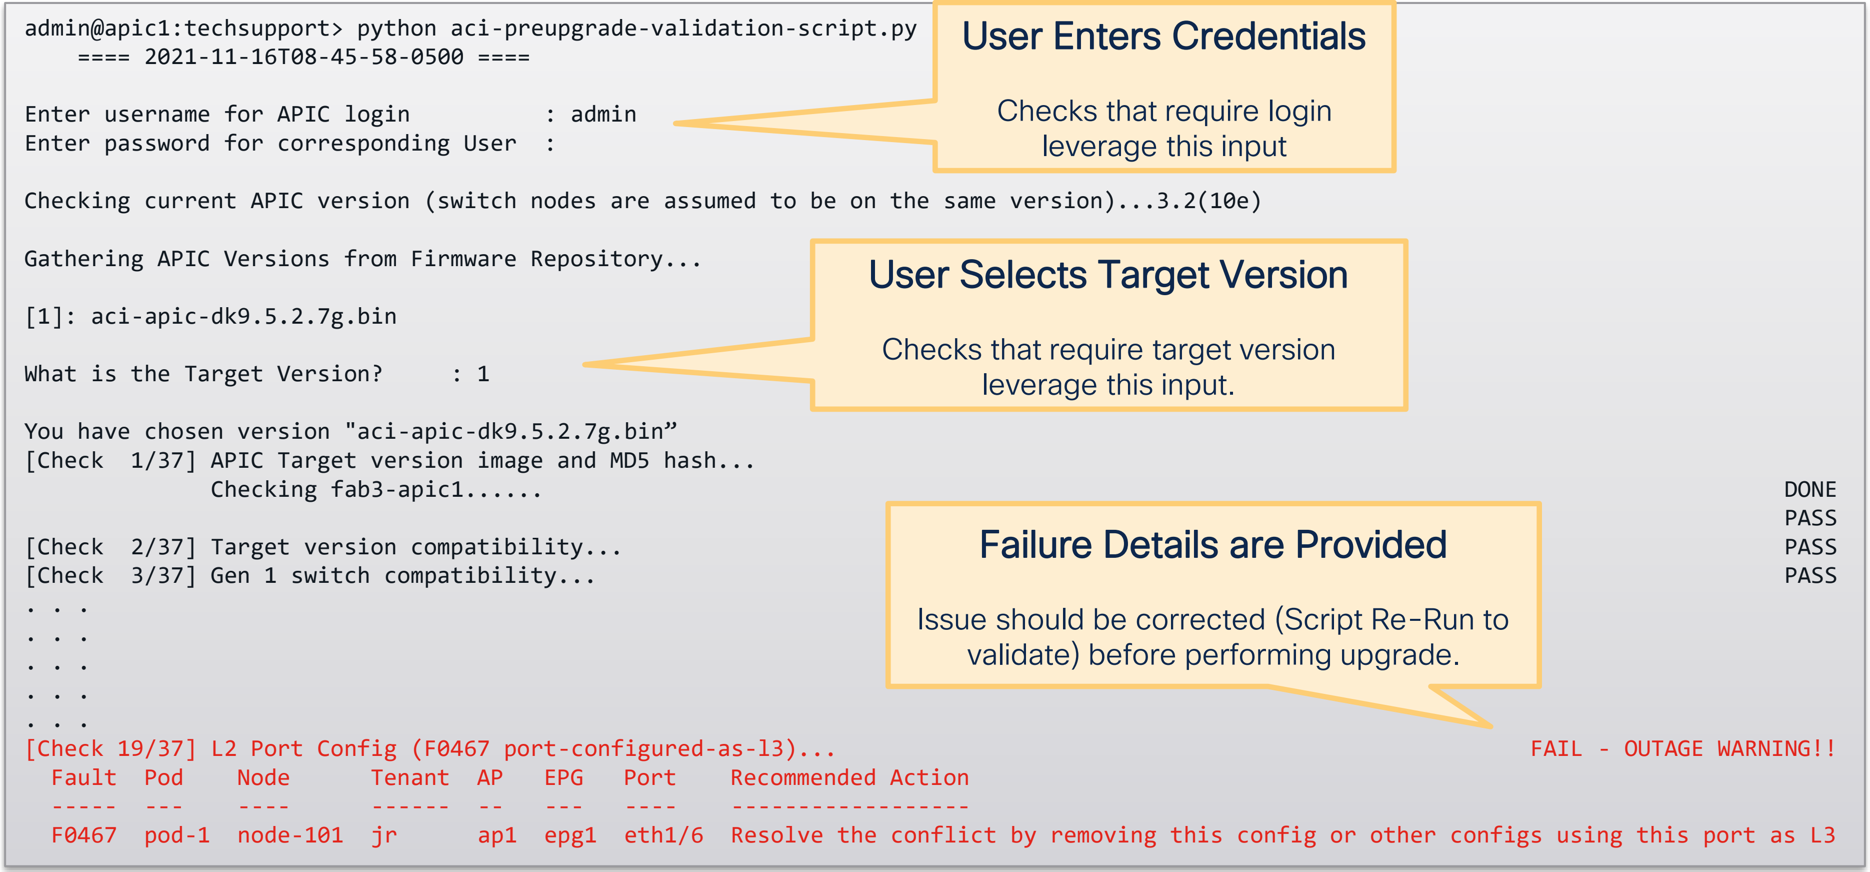

```
admin@apic1:techsupport> python aci-preupgrade-validation-script.py
    ==== 2021-11-16T08-45-58-0500 ====

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  1/37] APIC Target version image and MD5 hash...
              Checking fab3-apic1......                                                                          DONE
                                                                                                                 PASS
[Check  2/37] Target version compatibility...                                                                    PASS
[Check  3/37] Gen 1 switch compatibility...                                                                      PASS
. . .
. . .
. . .
. . .
. . .
[Check 19/37] 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 (Preferred)

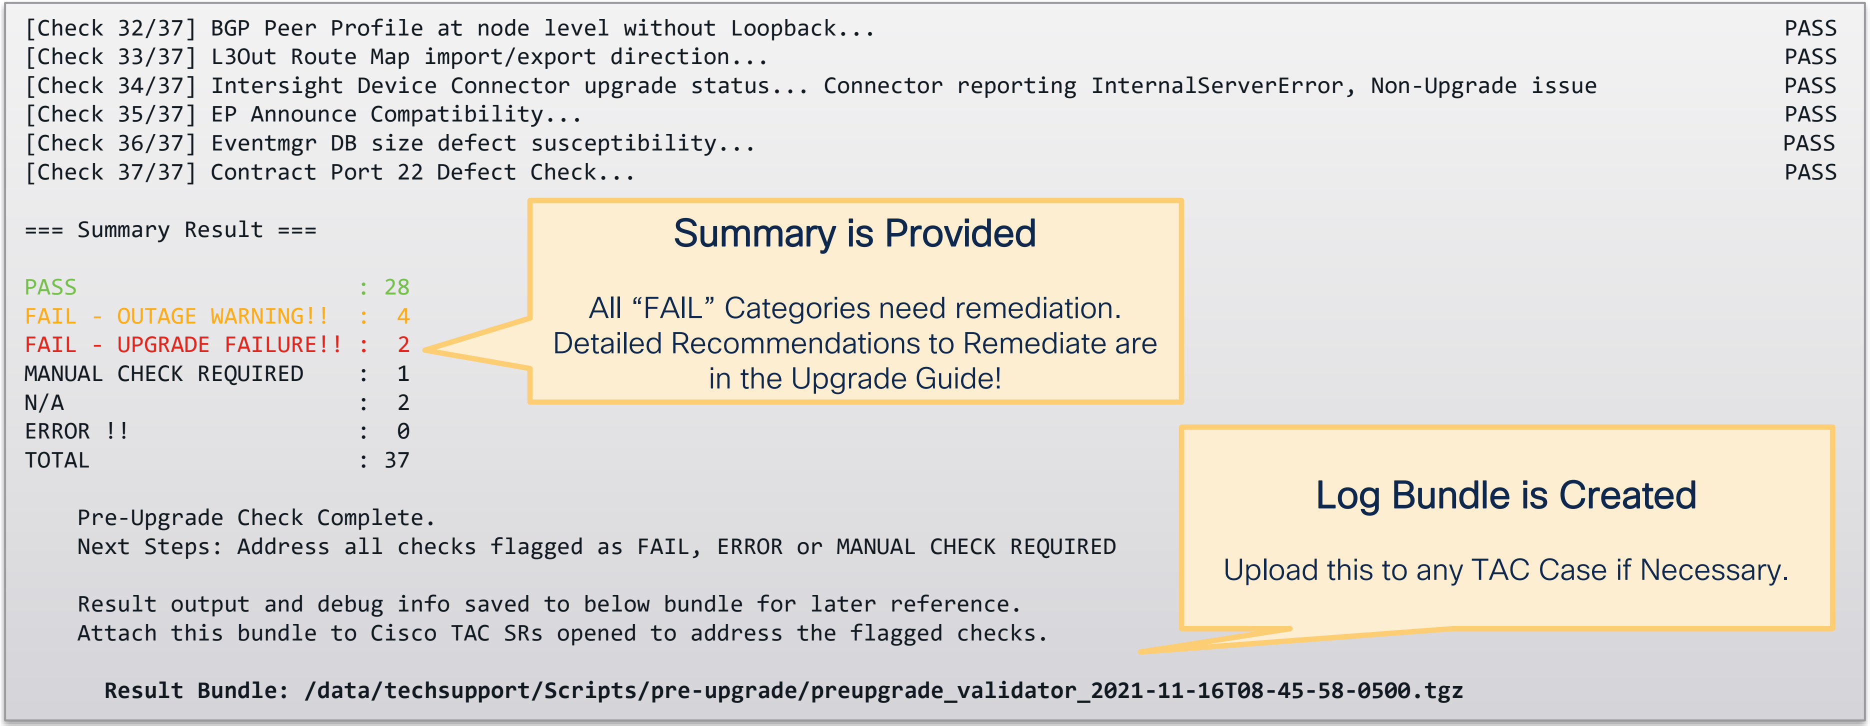

```
[Check 32/37] BGP Peer Profile at node level without Loopback...                                                       PASS
[Check 33/37] L3Out Route Map import/export direction...                                                               PASS
[Check 34/37] Intersight Device Connector upgrade status... Connector reporting InternalServerError, Non-Upgrade issue  PASS
[Check 35/37] EP Announce Compatibility...                                                                             PASS
[Check 36/37] Eventmgr DB size defect susceptibility...                                                                PASS
[Check 37/37] Contract Port 22 Defect Check...                                                                         PASS

=== Summary Result ===

PASS                       : 28
FAIL - OUTAGE WARNING!!    :  4
FAIL - UPGRADE FAILURE!!   :  2
MANUAL CHECK REQUIRED      :  1
N/A                        :  2
ERROR !!                   :  0
TOTAL                      : 37

    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_2021-11-16T08-45-58-0500.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.

Nexus Dashboard Insights (Optional)

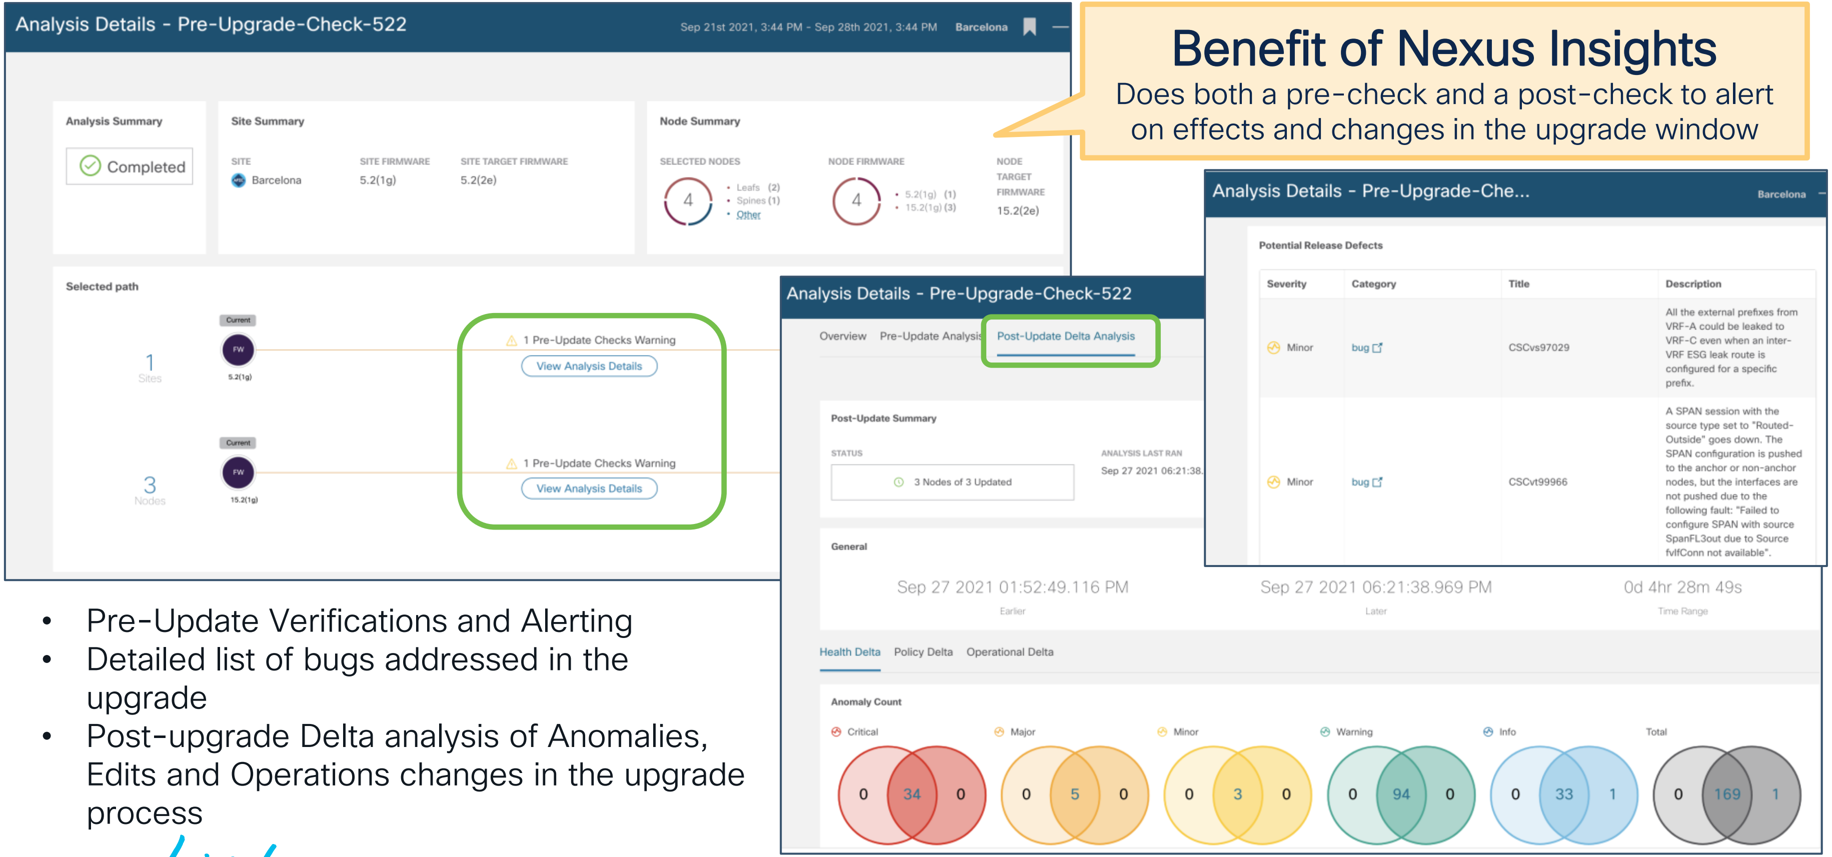

Analysis Details - Pre-Upgrade-Check-522

Sep 21st 2021, 3:44 PM - Sep 28th 2021, 3:44 PM Barcelona

Analysis Summary

✓ Completed

Site Summary

SITE	SITE FIRMWARE	SITE TARGET FIRMWARE
Barcelona	5.2(1g)	5.2(2e)

Node Summary

SELECTED NODES: 4
- Leafs (2)
- Spines (1)
- Other

NODE FIRMWARE: 4
- 5.2(1g) (1)
- 15.2(1g) (3)

NODE TARGET FIRMWARE: 15.2(2e)

Selected path

1 Sites — Current FW 5.2(1g) — ⚠ 1 Pre-Update Checks Warning — View Analysis Details

3 Nodes — Current FW 15.2(1g) — ⚠ 1 Pre-Update Checks Warning — View Analysis Details

Analysis Details - Pre-Upgrade-Check-522

Overview | Pre-Update Analysis | Post-Update Delta Analysis

Post-Update Summary

STATUS: 3 Nodes of 3 Updated

ANALYSIS LAST RAN: Sep 27 2021 06:21:38.

General

Earlier	Later	Time Range
Sep 27 2021 01:52:49.116 PM	Sep 27 2021 06:21:38.969 PM	0d 4hr 28m 49s

Health Delta | Policy Delta | Operational Delta

Anomaly Count

Critical	Major	Minor	Warning	Info	Total
0 / 34 / 0	0 / 5 / 0	0 / 3 / 0	0 / 94 / 0	0 / 33 / 1	0 / 169 / 1

Analysis Details - Pre-Upgrade-Che... Barcelona

Potential Release Defects

Severity	Category	Title	Description
Minor	bug	CSCvs97029	All the external prefixes from VRF-A could be leaked to VRF-C even when an inter-VRF ESG leak route is configured for a specific prefix.
Minor	bug	CSCvt99966	A SPAN session with the source type set to "Routed-Outside" goes down. The SPAN configuration is pushed to the anchor or non-anchor nodes, but the interfaces are not pushed due to the following fault: "Failed to configure SPAN with source SpanFL3out due to Source fvIfConn not available".

Benefit of Nexus Insights

Does both a pre-check and a post-check to alert on effects and changes in the upgrade window

- 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.

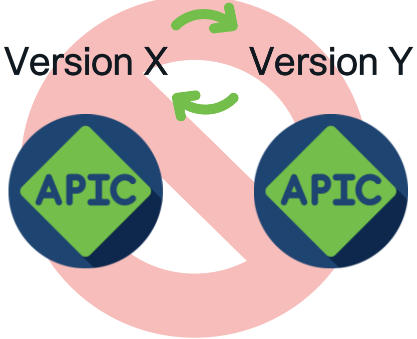

Final Tip

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

The bridge to possible

Thank you

cisco Live!

Let’s go