



#CiscoLive



Why You Shouldn't Fear Upgrading Your ACI Fabric Part I - Under the Hood

Takuya Kishida, Joseph Ristaino Cloud Networking Business Unit BRKDCN-2810a



#CiscoLive

Cisco Webex App

Questions?

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

How

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

Webex spaces will be moderated by the speaker until June 17, 2022.



https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2810a

cisco / ille



Agenda

- ACI Firmware Upgrade Types
- Upgrade Architecture APIC
- Upgrade Architecture Switches
- (Bonus) Upgrade Enhancements



ACI Firmware Upgrade Types

cisco live!

ACI Firmware Upgrade Types



Regular Upgrade







ACI Firmware Upgrade Types (Regular)



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-upgradedowngrade/Cisco-APIC-Installation-ACI-Upgrade-Downgrade-Guide/m-operationsallowed-during-mixed-versions-on-cisco-aci-switches.html





ACI Firmware Upgrade Types (SMU)





ACI Firmware Upgrade Types (EPLD/FPGA)



cisco / ille

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 ive!

APIC Upgrade Architecture











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



APIC Upgrade Architecture



cisco ile

APIC Upgrade Architecture



is considered completed.

numbered **APICs** finish data conversion and reboot.



ACI Switch Upgrade Architecture

cisco live!







cisco live!







Boot Up

Various traffic flow optimizations

01

02

03

04

05

06

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





Boot Up

· Various traffic flow optimizations

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



03

04

05

06

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





Boot Up

· Various traffic flow optimizations

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



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



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



05

06









cisco live!

Boot Up

· Various traffic flow optimizations

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



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



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



(05)

06

• Starts downloading configurations from an APIC





Boot Up

· Various traffic flow optimizations



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



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



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



• Starts downloading configurations from an APIC



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





ISIS multicast overload timer

- Leaf nodes Fixed 1min
- Spine nodes When FTAG tree is created
 (Fixed 1 min prior to Switch 14)
 - (Fixed 1 min prior to Switch 14.2(1))



Spine ISIS multicast overload timer (CSCvp79708)



Why not a fixed 1 min?



Boot l	Jp
--------	----

- · Various traffic flow optimizations
- Bring up fabric links
 - Bring up APIC connected down links
 - Admin down other down links



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



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



• Starts downloading configurations from an APIC



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



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

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.



Boot Up

· Various traffic flow optimizations

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



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



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



• Starts downloading configurations from an APIC



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



- Full configuration has been downloaded
- ✓ Bring up access links (downlinks)
- $\checkmark~$ 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



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

	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
Reboot	✓ 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
Reboot	 Leaf – all down links including APIC connected links
	 ✓ Spine – all IPN links
	2. Wipe the config and reboot (i.e. clean reboot)

Traffic Disruption without Graceful Upgrade OSPF DR reboot example





With Graceful Upgrade

OSPF DR reboot example



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.



Upgrade Enhancements

cisco live!

ACI Upgrade Enhancement Quick Summary

		Supported APIC versions	3.2	4.1(1)	4.2(*)	4.2(5)	5.0/5.1	5.2(1)	5.2(3)	Switch version requirements
Upgrade Time Dptimization Visibility Visibility Operation Optimization	Switch Image Pre-download		\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	14.1(1) or later	
Upgrade Time		Multi-Pod Parallel Switch Upgrade				\checkmark	\checkmark	\checkmark	\checkmark	No requirements
Optimization		Unlimited Parallel Switch Upgrade By Default				V	✓	V	✓	No requirements
Vicibility		APIC Detailed Install Stage				\checkmark	\checkmark	\checkmark	\checkmark	N/A
		Switch Image Download Progress				\checkmark	\checkmark	\checkmark	\checkmark	14.5(1) or later
		Pre-Upgrade Validation*			\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	No requirements
Operation		SMU Support						\checkmark	\checkmark	15.2(1) or later
Optimization		Auto EPLD/FPGA upgrade						\checkmark	\checkmark	15.2(1) or later
Operation Optimization		NXOS to ACI auto conversion via POAP							\checkmark	15.2(3) or later

cisco ive!

*APIC Built-in validation. The standalone script supports all APIC versions. More to come in BRKDCN-2810-b.

#CiscoLive BRKDCN-2810a

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public 38





Switch Image Download from APIC to switches



Upgrade multiple pods/switches in parallel

cisco

Switch Image Pre-Download with a scheduler





- 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



Firmware													?
								Summary	Infrastructure	Images	Faults	His	tory
											Controllers	No	des
Enforce Bo Default Firm	ootscript Versio vare Version:	on Validat select an	ion: 🔽 option		(switches	w in ACI 4.2(5), downlo s need to be 14.2(5) fo	oad or th	progress his functionality)				Ŏ	***
▲ ID	Name	Role	Model	Current Firmware	Upgrade Group	Download Progress		Status	I	Jpgrade Prog	ress		
Pod1/1001	f2-spine1	spine	N9K-C9332C	n9000-15.0(0.128)	ALL 🕑 Target FW: n9000-15.0(0.139b)		30%	Firmware upgrade queued with	group ALL to				0%
Pod1/1002	f2-spine2	spine	N9K-C9332C	n9000-15.0(0.128)	ALL 🔁 Target FW: n9000-15.0(0.139b)		30%	Firmware upgrade queued with	group ALL to				0%
Pod1/101	f2-leaf1	leaf	N9K-C93180YC-FX	n9000-15.0(0.128)	ALL 🕑 Target FW: n9000-15.0(0.139b)		30%	Firmware upgrade queued with	group ALL to				0%
Pod1/102	f2-leaf2	leaf	N9K-C93180YC-FX	n9000-15.0(0.128)	ALL 🕑 Target FW: n9000-15.0(0.139b)		30%	Firmware upgrade queued with	group ALL to				0%
Pod1/103	f2-leaf3	leaf	N9K-C93240YC-FX2	n9000-15.0(0.128)	ALL 🕑 Target FW: n9000-15.0(0.139b)		30%	Firmware upgrade queued with	group ALL to				0%
Pod1/104	f2-leaf4	leaf	N9K-C93240YC-FX2	n9000-15.0(0.128)	ALL 🕑 Target FW: n9000-15.0(0.139b)		30%	Firmware upgrade queued with	group ALL to				0%

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

Firmware												?
							Summary	Infrastructure	Images	Faults	Hist	tory
										Controllers	Nor	des
						Remain empty					Q	***
Enforce Bo	ootscript Version Validatio	on: 📃										
▲ ID	Name	Role	Model	Current Firmware	Upgrade Group	Download Progress	Status		Upgra	ade Progress		
Pod1/1001	F3-P1-Spine-1001	spine	N9K-C9364C	n9000-14.2(4i)	ALL 🔁 Target FW: n9000-15.0(1i)		Firmware upgrade q	ueued 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 q	ueued 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 q	ueued 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 q	ueued with group ALL	t			0%
Pod2/2001	F3-P2-Spine-2001	spine	N9K-C9504	n9000-14.2(4i)	ALL C Target FW: n9000-15.0(1i)		Firmware upgrade q	ueued with group ALL	t			0%
Pod2/2002	F3-P2-Spine-2002	spine	N9K-C9504	n9000-14.2(4i)	ALL C Target FW: n9000-15.0(1i)		Firmware upgrade q	ueued with group ALL	t			0%
Pod2/201	F3-P2-Leaf-201	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL [2] Target FW: n9000-15.0(1i)		Firmware upgrade q	ueued with group ALL	t 🧲			0%
Pod2/202	F3-P2-Leaf-202	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL P Target FW: n9000-15 0(1i)		Firmware upgrade q	ueued with group ALL	t 🧲			0%
Pod2/211	F3-P2-RL-211	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL P Target FW: n9000-15 0(1i)		Firmware upgrade q	ueued 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 of APIC is 4.2(5) or later

Upgrade multiple pods/switches in parallel



						Summary	Infrastructure	Images	Faults	History
									Controllers	Nodes
Enforce Bo	ootscript Version Valida	tion: 🗌					One pod at	a time		0 %-
ID	Name	Role	Model	Current Firmware	 Upgrade Group 	Status		Upg	rade Progress	
Pod1/101	F3-P1-Leaf-101	leaf	N9K-C93180YC-EX	n9000-14.1(2g)	ODD 🗗 Target FW: n9000-14.1(2x)	Firmware upgrade in progress with group ODD to desire	ed version n9000-14.1	(2x)		60%
Pod1/111	F3-P1-RL-111	leaf	N9K-C9336C-FX2	n9000-14.1(2g)	ODD 🗗 Target FW: n9000-14.1(2x)	Firmware upgrade in progress with group ODD to desir	ed version n9000-14.1	(2x)		60%
Pod2/201	F3-P2-Leaf-201	leaf	N9K-C93180YC-EX	n9000-14.1(2g)	ODD 🗗 Target FW: n9000-14.1(2x)	Firmware upgrade queued with group ODD to desired v	version n9000-14.1(2x)			0%
Pod2/211	F3-P2-RL-211	leaf	N9K-C93180YC-EX	n9000-14.1(2g)	ODD 🗗 Target FW: n9000-14.1(2x)	Firmware upgrade queued with group ODD to desired v	version n9000-14.1(2x)			0%
Pod1/1001	F3-P1-Spine-1001	spine	N9K-C9364C	n9000-14.1(2g)	ODD 🗗 Target FW: n9000-14.1(2x)	Firmware upgrade in progress with group ODD to desir	ed version n9000-14.1	(2x)		60%
Pod2/2001	F3-P2-Spine-2001	spine	N9K-C9504	n9000-14.1(2g)	ODD 🗗 Target FW: n9000-14.1(2x)	Firmware upgrade queued with group ODD to desired v	version n9000-14.1(2x)			0%
Pod2/202	F3-P2-Leaf-202	leaf	N9K-C93180YC-EX	n9000-14.1(2g)		Not Scheduled				

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

#CiscoLive BRKDCN-2810a

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public 43

Unlimited Parallel Upgrade



rmware												
							Summary	Infrastructure	e Images	Faults	Histo	or
s at o										Controllers	Nod	le
15 at 01	lice										Q	4 -
Enforce Bo	otscript Version Validat	tion: 🗌										
▲ ID	Name	Role	Model	Current Firmware	Upgrade Group	Download Progress	Status		Ipgrade Progres	iS		
Pod1/1001	F3-P1-Spine-1001	spine	N9K-C9364C	n9000-14.2(4i)	ALL 🔁 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A	_			4
Pod1/1002	F3-P1-Spine-1002	spine	N9K-C9364C	n9000-14.2(4i)	ALL 🔁 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A)			4
Pod1/101	F3-P1-Leaf-101	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL 🗗 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A				4
Pod1/111	F3-P1-RL-111	leaf	N9K-C9336C-FX2	n9000-14.2(4i)	ALL 🕑 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A				4
Pod2/2001	F3-P2-Spine-2001	spine	N9K-C9504	n9000-14.2(4i)	ALL 🗗 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A				4
Pod2/2002	F3-P2-Spine-2002	spine	N9K-C9504	n9000-14.2(4i)	ALL 🛃 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A				4
Pod2/201	F3-P2-Leaf-201	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL 🕑 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A				4
Pod2/202	F3-P2-Leaf-202	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL 🛃 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A				4
Pod2/211	F3-P2-RL-211	leaf	N9K-C93180YC-EX	n9000-14.2(4i)	ALL 🔁 Target FW: n9000-15.0(1i)		Firmware upgrade in progress v	with group A				4
Pod2/212	F3-P2-RL-212	leaf					unknown					

- 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

cisco ile

Reference

 Cisco APIC Installation and ACI Upgrade and Downgrade Guide
 bttps://www.cisco.com/c/on/us/td//docs/don/aci/anic/all/

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/all/api c-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-systemsmanagement/application-policy-infrastructure-controllerapic/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/get ting-started/cisco-apic-getting-started-guide-52x/fabricinitialization-52x.html#d5018e3247a1635



Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.



Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.

E Learn

Cisco U.

IT learning hub that guides teams and learners toward their goals

Cisco Digital Learning

Subscription-based product, technology, and certification training

Cisco Modeling Labs

Network simulation platform for design, testing, and troubleshooting

Cisco Learning Network

Resource community portal for certifications and learning

En Train

Cisco Training Bootcamps Intensive team & individual automation and technology training programs

Cisco Learning Partner Program

Authorized training partners supporting Cisco technology and career certifications

Cisco Instructor-led and Virtual Instructor-led training

Accelerated curriculum of product, technology, and certification courses

E Certify

Cisco Certifications and Specialist Certifications

Award-winning certification program empowers students and IT Professionals to advance their technical careers

Cisco Guided Study Groups

180-day certification prep program with learning and support

Cisco Continuing Education Program

Recertification training options for Cisco certified individuals

Here at the event? Visit us at The Learning and Certifications lounge at the World of Solutions



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 <u>www.CiscoLive.com/on-demand</u>



CISCO The bridge to possible

Thank you



#CiscoLive





#CiscoLive