ACI Upgrade Best Practices

ACI TAC Troubleshooting Series – Part I

Linda Wang
Technical Leader - ACI
September 2021

Agenda

- ACI Pre-upgrade Best Practice
- ACI Upgrade Tools
- Common Upgrade Problems
- Unsupported Operations During Upgrade
- Troubleshooting Failures During Upgrade
- Post Upgrade Validation
- QA

Why Upgrade ACI?

- Security updates, Bug fixes, Software enhancement
- New software features, Hardware support
- Higher scalability .etc

Examples:

1. Customer should upgrade to ACI 4.2(6), or above which has fixes for SSD related issues (CSCVx19640/CSCvt36458/FN-72145/FN-70538).
2. Auto-EPLD upgrade feature is delivered in ACI 5.2(1) release
3. Auto conversion of Spine/Leaf from NXOS to ACI is committed in 5.2(3) release

ACI Pre-Upgrade Best Practice

Check Basic Information on Your Fabric

- Clear all your faults
- Perform a configuration export **with AES Encryption**
- Verify access to out-of-band IP addresses of all your ACI nodes (all your APIC nodes and switch nodes)
- Verify CIMC access for all your APICs
- Verify console access for all your switches
- Understand **Changes in Behavior** in Release Notes of both APIC and ACI switches for versions between the target and current version
- Understand **Open Issues** and **Known Issues** in Release Notes of both APIC and ACI switches for the target version

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-pre-upgrade-checklists.html

Configuration Export with AES Encryption

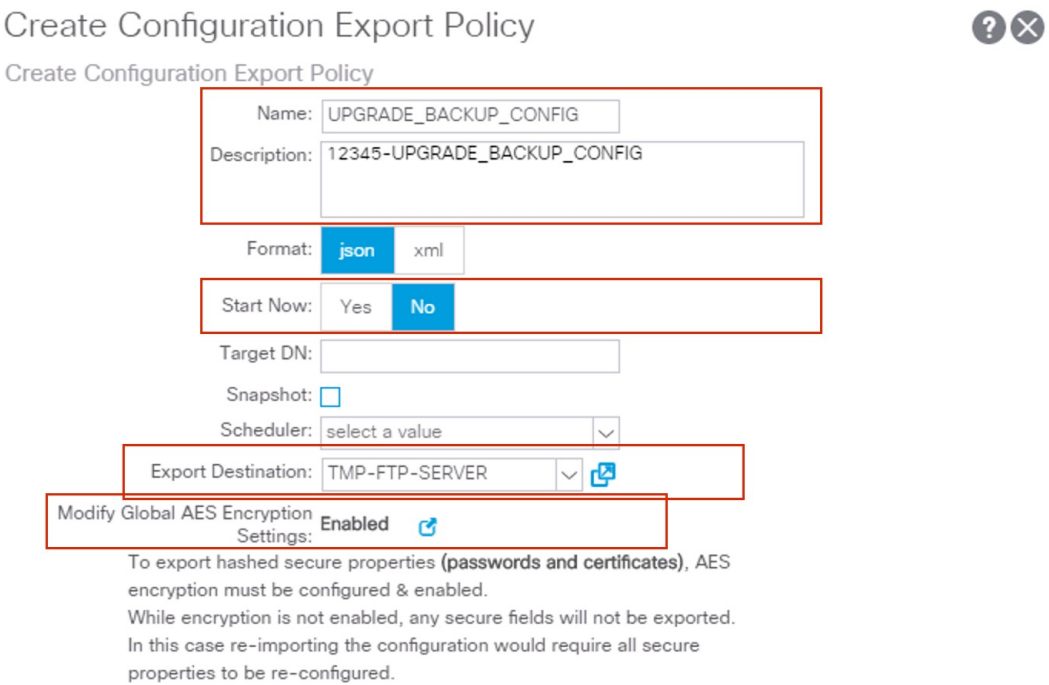

Pre-Upgrade Best Practice

- Verify `/firmware` is less than 75%
- Test browser access to APIC, CIMC KVM console with Java and HTML or CIMC CLI
- VPC switch pairs/BGP RR in different maintenance groups
- NTP is reachable
- Verify boot variable is set up correctly
- Review behavior changes of new version and evaluate potential impact (read release notes)
- Stage ACI upgrade in lab before applying change in production

Configurations That Must Be Disabled Prior To Upgrades

The following features must **be disabled** prior to upgrades:

- App Center apps
- Maintenance Mode through **Fabric** > **Inventory** > **Fabric Membership** > **Maintenance (GIR)**
- Config Zone
- Rogue Endpoint (only when the running version is 14.1(x) or when upgrading to 14.1(x))

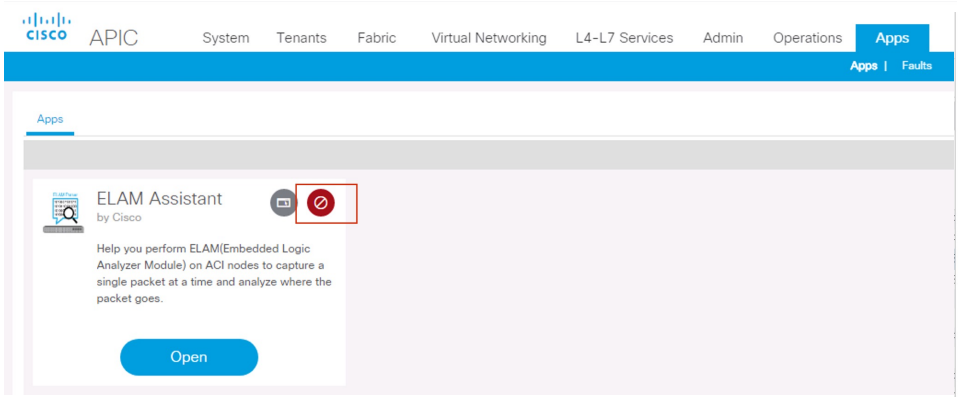

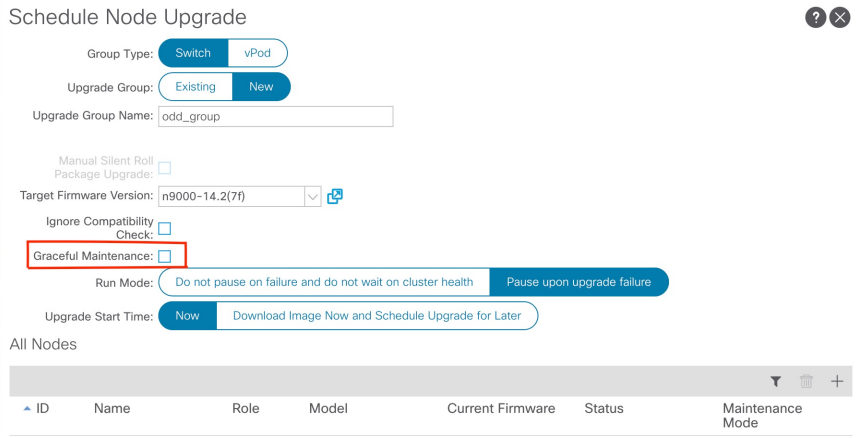

ACI Upgrade Tools

APIC Upgrade/Downgrade Support Matrix

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 4.0(1)

To release 5.2(2)

Current release: 4.0(1)

Target release: 5.2(2) [↗]

Recommended path: 4.0(1) → 4.2(1) → 5.2(2) [Hide Alternate Paths]
4.0(1) → 4.2(1) → 5.2(2)
4.0(1) → 4.1(2) → 5.2(2)
4.0(1) → 4.0(3) → 5.2(2)

Procedure:

- Upgrade the Cisco APICs. Unless otherwise stated, we recommend upgrading to the latest letter release in the target release train.
- After the Cisco APICs are upgraded successfully, upgrade the switches using 2 or more maintenance groups.
- After the APICs and the switches are upgraded successfully, upgrade the Cisco ACI Virtual Edge or Cisco AVS.
- Once an APIC cluster upgrade has been triggered, you must allow the APIC cluster upgrade to complete before attempting any Target Version rollback. If any APIC in the cluster encounters an issue during the upgrade, you must first resolve that issue and get all of the APICs in the cluster to reach the Target Version that you set. Changing the Target Version while the APIC upgrades are in process is not supported and can produce unexpected results.

Caveats:

- This tool provides most important upgrade-related reference information. It does not provide an exhaustive list of all caveats and guidelines.
  Ensure that you have read the target release's *Release Notes* and *Upgrade Guide* for more specific information that may not be listed here.
- You can have at most 2 different releases in the fabric at any given time. The entire fabric must be upgraded to each intermediate release before any one element is upgraded to the next hop. This means upgrading the APICs and then upgrading the switches to each intermediate release.
- If you are using Cisco AVE, you must upgrade or downgrade it to the recommended version for each intermediate APIC release before continuing with your fabric upgrade.
- **NOTE:** If you have Remote Leaf Switches, you must enable "direct traffic forwarding" before you upgrade to APIC release 5.0(1) or later.
  If you are upgrading from a release prior to 4.1(2), the option will not be available. In this case, you must first upgrade to a 4.2(x) release, enable "direct traffic forwarding" option, then upgrade to your target 5.x release.

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

Pre-Upgrade Validation Tools

- Pre-Upgrade Validator (APIC)
- Pre-Upgrade Validator (App Center app)
  - Allows users in older APIC firmware to perform the latest validations.
- Standalone ACI-Pre-Upgrade-Validation Script
  - For any feature not currently implemented in the Pre-Upgrade Validator, a standalone script can be run directly on the APIC to validate any existing issues prior to upgrading. The script supports all versions of software.
  - See https://github.com/datacenter/ACI-Pre-Upgrade-Validation-Script for more details about the script.

See live demo

State Checker: Introduction

- State Checker (also known as StateChangeChecker) is a Cisco ACI application that allows operators to snapshot a collection of managed objects (MO) in the fabric and perform snapshot comparisons.
- This allows operators to answer questions such as:
  - ✓ What changed in my fabric ?
  - ✓ Are my critical objects the same after maintenance ?
  - ✓ Did any route change on any node ?
  - ✓ Are all the local endpoint learns still present ?

- This application can be installed directly on the Cisco APIC or deployed in standalone mode.

State Checker: Get Fabric State Snapshot With Demo

- Use state checker to get current Fabric state snapshot – Prior to upgrade
- Run state checker in standalone mode

```
docker run --name statechecker -p 5000:443 -d agccie/statechecker:latest
```

  - http://localhost:5000
  - Run on another host – not APIC

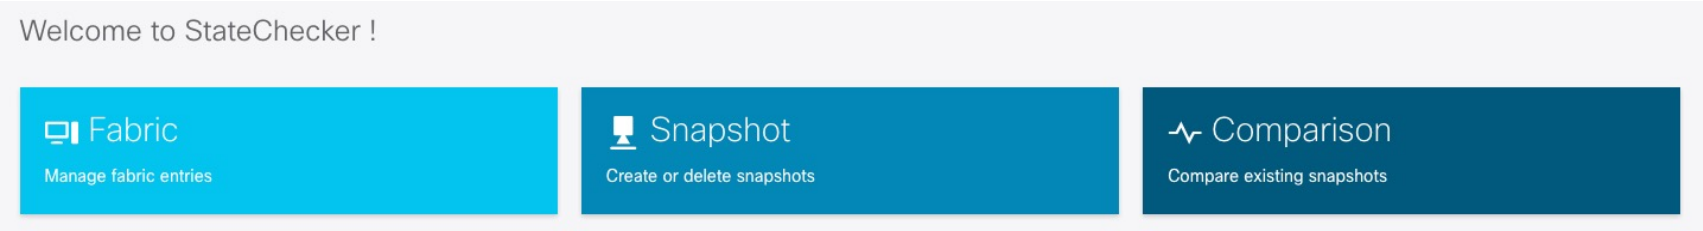

https://statechecker.readthedocs.io/en/latest/install.html#standalone-mode

Common Upgrade Problems

What Could Go Wrong During Upgrade

- Upgrade failed: Did not start, failed compatibility check (hardware and software), or out of disk space
- Upgrade failed: Stuck mid-upgrade
- Upgrade completed: APIC cluster is not fully-fit, Services are not available
- Upgrade completed: Network/software behavior change (something is wrong)
- Upgrade complete: Some services/port not available

Switch Upgrade Stuck in the queue and won't kick in

- If any manual intervention of the fabric was attempted before the upgrade within 30 minutes or after upgrade, the scheduler needs to be cleared. By manual intervention, we mean anything that includes
  - Removing a switch
  - Replacing a switch
  - Editing group membership
  - Decommissioning
- The symptom would be some switch would never upgrade and always wait in the queue. Create a new firmware group and put that switch to the new group may help.

Upgrade failed due to /root partition usage over 75%

- If root partition has been used more than 75% (which by default should be 34%), the installer.py will fail to complete.
- So always check if there is a faults "F1527" for partition /, then go to the root mode figure out which files have used the space.
- Take away: Clear all faults seen prior to upgrade, e.g. APIC Disk Space Usage (F1527, F1528, F1529)

Post-Upgrade Dataplane/EP learning/Spanning-Tree issues due to overlapping VLAN pools

1 Before upgrade

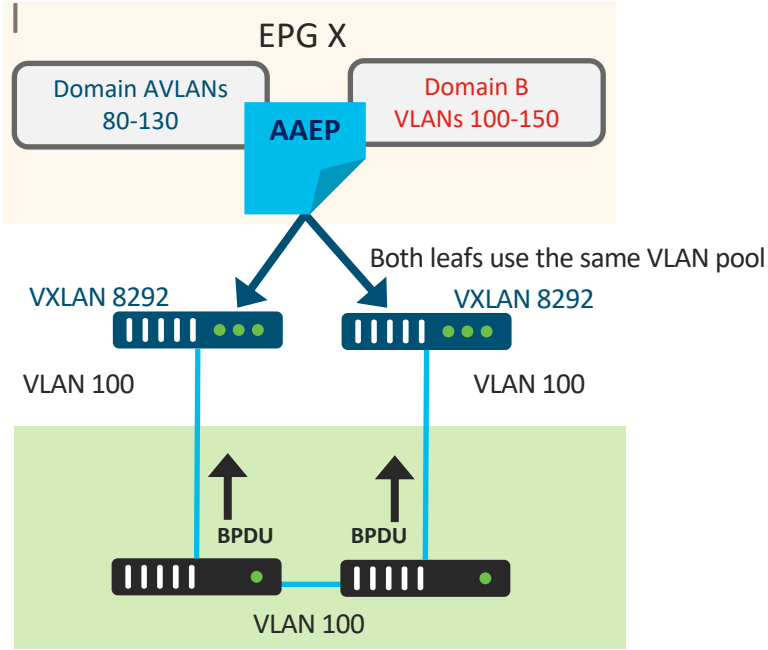

2 After upgrade

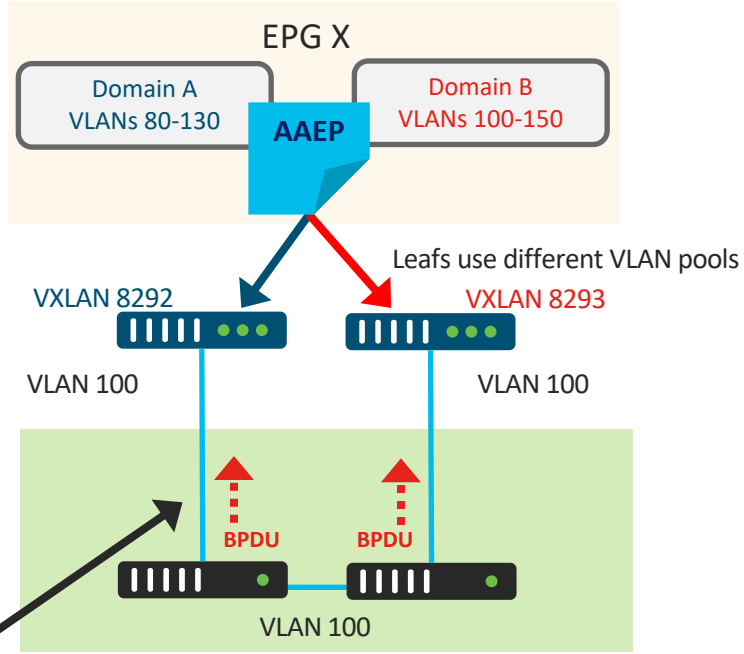

BPDUs are dropped
Layer 2 Loop

Incident Post Upgrade due to Software Behavior Change

Extra validations which are added in newer version may cause some misconfigurations stop working post upgrade, Issue symptoms may vary depending on which version is running with the misconfiguration.

- Use case 1:

  Having a configuration where a given subnet is defined under both an EPG as well as its corresponding BD, however both subnet definitions have a different combination of "scope" definitions: [private, public, shared], this is a misconfiguration. See example below.

```
dn        : uni/tn-TN-PROD/BD-APP/subnet-[192.168.1.126/27]
scope       : public,shared
dn        : uni/tn-TN-PROD/ap-AP-APP-/epg-EPG-APP/subnet-[192.168.1.126/27]
scope       : shared
```

- Use case 2:

  Some overlapping subnets are configured as BD and as external subnet in the same VRF, post upgrade you might see some security zoning rules disabled on the boarder leafs

```
uni/tn-SYD/BD-PROD/subnet-[10.48.32.1/24]
```

```
uni/tn-SYD/out-Internet/instP-Internet/extsubnet-[10.48.32.0/24]
```

Unsupported Operations During Upgrade

No Clean Reloads During the Upgrade

- Do not attempt a clean reload during the upgrade.
- Takeaway: Please be patience during the upgrade
  - Time taken per APIC may vary in a cluster , we have had a case where APIC 1 and 2 in a 3 apic cluster took x time, when APIC 3 crossed x time, customer panicked and reloaded APIC3. On reviewing logs later, APIC 3 was still in the data conversion process.

Rolling back Target Version is Unsupported

- Do not attempt a rollback by changing the ACI cluster target version back to the initial version.
  a. Example: Upgrade from 3.1 to 4.1, APIC x fails upgrade. Do not change target version back to 3.1 as a troubleshooting step, all focus should be on why APIC x failed to upgrade
  b. Changing the target version mid-upgrade is unsupported and can ultimately result in database shards being deleted off of APICs.
- Take away: Don't panic and just rollback, please engage TAC to troubleshoot the upgrade failure.

Operations Allowed During Mixed Versions

Supported Operations with Mixed Versions for Each Upgrade Path

Upgrade Path		Supported Operations
From	To	
2.2(x)	Any versions in the supported upgrade path	•Exporting configuration •Collecting techsupport •Physical network change (i.e. reboot, cable replacement etc.) •Policy changes for features introduced prior to the major release*
2.3(x) or later	Any versions in the supported upgrade path	•Exporting configuration •Collecting techsupport •Physical network change (i.e. reboot, cable replacement etc.) •Policy changes for features introduced prior to the major release* •Policy changes for features in Supported Operations with Mixed Versions for Upgrades from Release 2.3(x) or Later

* This operation is supported only when the upgrade is within the same release train (for example, an upgrade from 3.2(5d) to 3.2(5f), where the releases are still part of the 3.2(5) release train, but the upgrade occurs between the **d** and the **f** versions of that release train).

APIC L3/M3 supports 4.x only

- APIC L3/M3 supports 4.x only (cannot downgrade to older release)

Troubleshooting Failures During Upgrade

Upgrade installer log locations

- How to monitor and not panic if no progress for a long time..
- While the upgrade is running, it is written into **/root/insieme_installer.log** and once the installation is complete, it is appended to **/firmware/insieme_installer.log**.
- The data conversion logs are in /firmware/dataconv.log and **/firmware/dataconv_detail.log**.
- **Switches installer's** log are saved **/mnt/pss/installer.log,** DME logs are saved **/var/sysmgr/tmp_logs/**
- **Later versions include a new location and new filenames for upgrade installer logs:**

/firmware/logs/<timestamp>/**insieme_3x_to_4x_installer.log**
/firmware/logs/<timestamp>/**insieme_4x_installer.log**
/firmware/logs/<timestamp>/**atom_installer.log**

Full example:
/firmware/logs/2021-07-26T20:27:41-90/insieme_4x_installer.log

Useful Troubleshooting Commands for Upgrade

- show firmware upgrade status
- show firmware upgrade status detail – CLI version of GUI
- show firmware upgrade status switch-group Leafs-Odd detail

```
apic1# show firmware upgrade status switch-group odd-leaf-group
 Pod         Node        Current-Firmware      Target-Firmware       Status                    Upgrade-Progress(%)
 ----------  ----------  --------------------  --------------------  ------------------------  --------------------
 1           101         n9000-13.1(2m)        n9000-14.1(2g)        version not compatible    0
```

- moquery -c maintUpgStatus
- moquery -d topology/pod-1/node-101/sys/fwstatuscont/upgjob

DME processes involved During APIC upgrade

- The DMEs mainly responsible for the upgrade are BootMgr, PolicyManager, PolicyElement and ApplianceElement. Since 2.3 and above, the user configurations go into PolicyDistributor, which talks to PolicyManager.
- The logs for the DME processes are saved into **/var/log/dme/log**

Cluster became "data layer partially diverged" after APIC upgrade

- Use case:

  After the upgrade from 2.2(4r) to 4.1(1i) all the three APIC were in "Data Layer Partially Diverged" state

- Troubleshooting steps:
  - Check system faults
  - Check IP connectivity on infra network between APICs
  - Check CLI output: acidiag rvread
    - If some services (DME process) are not healthy
      1. Check process name (man acidiag)
      2. Check that process is running on each APIC (ps -eaf | grep svc)
      3. Check for core (show cores)
      4. Check DME process logs for error, or
      5. Collect techsupport logs and get TAC engaged for further troubleshooting

Post Upgrade Validation

Fabric State Snapshot (State Checker)

Use state checker to get current Fabric state snapshot – post upgrade

- ✓ What changed in my Fabric between windows?
- ✓ Are my critical objects the same after maintenance?
- ✓ Are all my routes/endpoints still present?

Post upgrade verification

- Use state checker compare post upgrade Fabric snapshot to pre upgrade Fabric snapshot

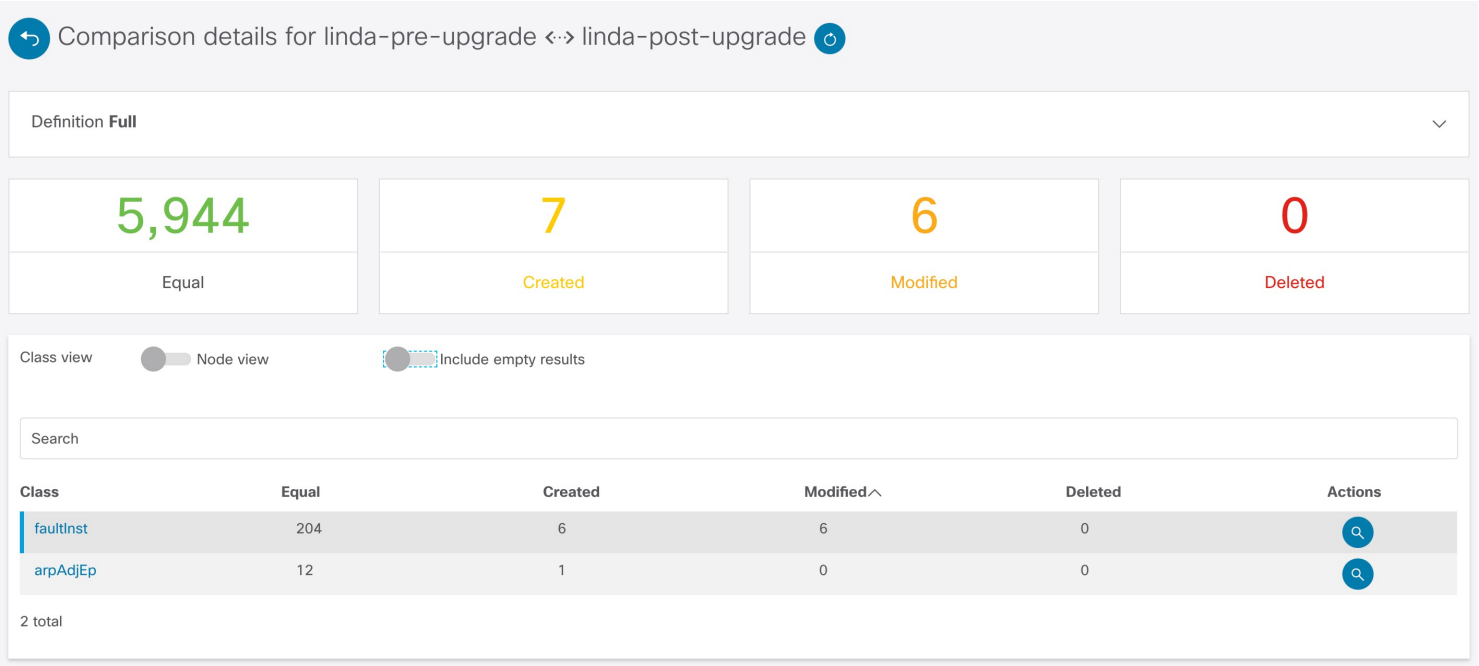

Verify Network Availability

- Execute test plan/verify application/network availability
  - Execute post upgrade test plan
  - Verify network Fabric connectivity
  - Verify network devices connectivity
  - Verify ACI integration such as VMM, AVE, MSO, and others
  - Verify application availability

References

Cisco Documentations

ACI Upgrade Downgrade Guide (CIMC upgrade procedure is also included)

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-pre-upgrade-checklists.html

ACI Upgrade Best Practices and Troubleshooting

https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/213618-aci-upgrade-best-practices-and-troublesh.html

Cisco ACI Upgrade Checklist

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

APIC Upgrade/Downgrade Support Matrix

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

ACI Pre-Upgrade Validation script on Github

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

APIC CIMC Upgrade Procedure on cisco community

https://community.cisco.com/t5/data-center-documents/apic-cimc-upgrade-procedure/ta-p/3216002

Questions?


CISCO