Configuring the Cisco APIC Using the CLI

Cluster Management Guidelines

The Cisco Application Policy Infrastructure Controller (APIC) cluster comprises multiple Cisco APICs that provide operators a unified real time monitoring, diagnostic, and configuration management capability for the Cisco Application Centric Infrastructure (ACI) fabric. To assure optimal system performance, use the following guidelines when making changes to the Cisco APIC cluster:

- Prior to initiating a change to the cluster, always verify its health. When performing planned changes to the cluster, all controllers in the cluster should be healthy. If one or more of the Cisco APICs' health status in the cluster is not "fully fit," remedy that situation before proceeding. Also, assure that cluster controllers added to the Cisco APIC are running the same version of firmware as the other controllers in the Cisco APIC cluster.
- We recommend that you have at least 3 active Cisco APICs in a cluster, along with additional standby Cisco APICs. Cisco APIC clusters can have from 3 to 7 active Cisco APICs. Refer to the Verified Scalability Guide to determine how many active Cisco APICs are required for your deployment.
- Disregard cluster information from Cisco APICs that are not currently in the cluster; they do not provide accurate cluster information.
- Cluster slots contain a Cisco APIC `ChassisID`. Once you configure a slot, it remains unavailable until you decommission the Cisco APIC with the assigned `ChassisID`.

- If a Cisco APIC firmware upgrade is in progress, wait for it to complete and the cluster to be fully fit before proceeding with any other changes to the cluster.
- When moving a Cisco APIC, first ensure that you have a healthy cluster. After verifying the health of the Cisco APIC cluster, choose the Cisco APIC that you intend to shut down. After the Cisco APIC has shut down, move the Cisco APIC, re-connect it, and then turn it back on. From the GUI, verify that the all controllers in the cluster return to a fully fit state.

Note Only move one Cisco APIC at a time.

- When moving a Cisco APIC that is connected to a set of leaf switches to another set of leaf switches or when moving a Cisco APIC to different port within the same leaf switch, first ensure that you have a healthy cluster. After verifying the health of the Cisco APIC cluster, choose the Cisco APIC that you intend to move and decommission it from the cluster. After the Cisco APIC is decomissioned, move the Cisco APIC and then commission it.
- Before configuring the Cisco APIC cluster, ensure that all of the Cisco APICs are running the same firmware version. Initial clustering of Cisco APICs running differing versions is an unsupported operation and may cause problems within the cluster.
- Unlike other objects, log record objects are stored only in one shard of a database on one of the Cisco APICs. These objects get lost forever if you decommission or replace that Cisco APIC.
- When you decommission a Cisco APIC, the Cisco APIC loses all fault, event, and audit log history that was stored in it. If you replace all Cisco APICs, you lose all log history. Before you migrate a Cisco APIC, we recommend that you manually backup the log history.

Replacing a Cisco APIC in a Cluster Using the CLI

Note
- For more information about managing clusters, see Cluster Management Guidelines.
- When you replace an Cisco APIC, the password will always be synced from the cluster. When replacing APIC 1, you will be asked for a password but it will be ignored in favor of the existing password in the cluster. When replacing Cisco APIC 2 or 3, you will not be asked for a password.

Before you begin

Before replacing a Cisco Application Policy Infrastructure Controller (APIC), ensure that the replacement Cisco APIC is running the same firmware version as the Cisco APIC to be replaced. If the versions are not the same, you must update the firmware of the replacement Cisco APIC before you begin. Initial clustering of Cisco APICs running differing versions is an unsupported operation and may cause problems within the cluster.

Step 1 Identify the Cisco APIC that you want to replace.

Step 2 Note the configuration details of the Cisco APIC to be replaced by using the **acidiag avread** command.

Step 3 Decommission the Cisco APIC using the **decommission controller** *controller-id* command in config mode.

Decommissioning the Cisco APIC removes the mapping between the APIC ID and Chassis ID. The new Cisco APIC typically has a different APIC ID, so you must remove this mapping in order to add a new Cisco APIC to the cluster.

Beginning with Cisco APIC release 6.0(2), an optional argument (`force`) is added to the **decommission** command to allow forcing the decommission operation. The revised command is **decommission controller** *controller-id* [**force**], with the following behaviors:

- When `force` is not declared, the decommission proceeds only if the cluster is not in an unhealthy or upgrade state, where a decommission may not be proper.
- When `force` is declared, the decommission proceeds regardless of the cluster state.

For example, `decommission controller 3 force` decommissions APIC3 regardless of the cluster state.

Step 4 To commission the new Cisco APIC, follow these steps:

a) Disconnect the old Cisco APIC from the fabric.
b) Connect the replacement Cisco APIC to the fabric.

The new Cisco APIC appears in the Cisco APIC GUI menu **System** > **Controllers** > ***apic_controller_name*** > **Cluster as Seen by Node** in the **Unauthorized Controllers** list.

c) Commission the new Cisco APIC using the **controller** *controller-id* **commission** command.
d) Boot the new Cisco APIC.
e) Allow several minutes for the new Cisco APIC information to propagate to the rest of the cluster.

The new Cisco APIC appears in the Cisco APIC GUI menu **System** > **Controllers** > ***apic_controller_name*** > **Cluster as Seen by Node** in the **Active Controllers** list.

What to do next

For each decommissioned controller, verify that the operational state of the controller is unregistered and that the controller is no longer in service in the cluster.

Note If a decommissioned Cisco APIC is not promptly removed from the fabric, it might be rediscovered, which could cause problems. In that case, follow the instructions in Reducing the APIC Cluster Size to remove the controller.

Reducing the APIC Cluster Size

Follow these guidelines to reduce the Cisco Application Policy Infrastructure Controller (APIC) cluster size and decommission the Cisco APICs that are removed from the cluster:

Note Failure to follow an orderly process to decommission and power down Cisco APICs from a reduced cluster can lead to unpredictable outcomes. Do not allow unrecognized Cisco APICs to remain connected to the fabric.

- Reducing the cluster size increases the load on the remaining Cisco APICs. Schedule the Cisco APIC size reduction at a time when the demands of the fabric workload will not be impacted by the cluster synchronization.
- If one or more of the Cisco APICs' health status in the cluster is not "fully fit," remedy that situation before proceeding.
- Reduce the cluster target size to the new lower value. For example if the existing cluster size is 6 and you will remove 3 controllers, reduce the cluster target size to 3.
- Starting with the highest numbered controller ID in the existing cluster, decommission, power down, and disconnect the APIC one by one until the cluster reaches the new lower target size.

  Upon the decommissioning and removal of each controller, the Cisco APIC synchronizes the cluster.

Note After decommissioning a Cisco APIC from the cluster, promptly power it down and disconnect it from the fabric to prevent its rediscovery. Before returning it to service, do a wiped clean back to factory reset.

If the disconnection is delayed and a decommissioned controller is rediscovered, follow these steps to remove it:

1. Power down the Cisco APIC and disconnect it from the fabric.
2. In the list of Unauthorized Controllers, reject the controller.
3. Erase the controller from the GUI.

- Cluster synchronization stops if an existing Cisco APIC becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.
- Depending on the amount of data the Cisco APIC must synchronize upon the removal of a controller, the time required to decommission and complete cluster synchronization for each controller could be more than 10 minutes per controller.

Note Complete the entire necessary decommissioning steps, allowing the Cisco APIC to complete the cluster synchronization accordingly before making additional changes to the cluster.

Contracting the Cisco APIC Cluster

Contracting the Cisco APIC cluster is the operation to decrease any size mismatches, from a cluster size of N to size N -1, within legal boundaries. As the contraction results in increased computational and memory

load for the remaining APICs in the cluster, the decommissioned APIC cluster slot becomes unavailable by operator input only.

During cluster contraction, you must begin decommissioning the last APIC in the cluster first and work your way sequentially in reverse order. For example, APIC4 must be decommissioned before APIC3, and APIC3 must be decommissioned before APIC2.

Switching Over Active APIC with Standby APIC Using CLI

Use this procedure to switch over an active APIC with a standby APIC.

Step 1 **replace-controller replace** *ID number Backup serial number*

Replaces an active APIC with an standby APIC.

Example:

```
apic1#replace-controller replace 2 FCH1804V27L
Do you want to replace APIC 2 with a backup? (Y/n): Y
```

Step 2 **replace-controller reset** *ID number*

Resets fail over status of the active controller.

Example:

```
apic1# replace-controller reset 2
Do you want to reset failover status of APIC 2? (Y/n): Y
```

Verifying Cold Standby Status Using the CLI

To verify the Cold Standby status of APIC, log in to the APIC as admin and enter the command **show controller**.

```
apic1# show controller
Fabric Name          : vegas
Operational Size     : 3
Cluster Size         : 3
Time Difference      : 496
Fabric Security Mode : strict

 ID    Pod   Address          In-Band IPv4     In-Band IPv6               OOB IPv4         OOB IPv6
                   Version             Flags  Serial Number     Health
 ----  ----  ---------------  ---------------  -------------------------  ---------------
-------------------------  ------------------  -----  ----------------  ------------------
 1*    1     10.0.0.1         0.0.0.0          fc00::1                    172.23.142.4
fe80::26e9:b3ff:fe91:c4e0  2.2(0.172)          crva-  FCH1748V0DF       fully-fit
 2     1     10.0.0.2         0.0.0.0          fc00::1                    172.23.142.6
fe80::26e9:bf8f:fe91:f37c  2.2(0.172)          crva-  FCH1747V0YF       fully-fit
 3     1     10.0.0.3         0.0.0.0          fc00::1                    172.23.142.8
fe80::4e00:82ff:fead:bc66  2.2(0.172)          crva-  FCH1725V2DK       fully-fit
 21~         10.0.0.21
                                      -----  FCH1734V2DG
```

```
Flags - c:Commissioned | r:Registered | v:Valid Certificate | a:Approved | f/s:Failover fail/success
(*)Current (~)Standby
```

Registering an Unregistered Switch Using the CLI

Use this procedure to register a switch from the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the CLI.

Note This procedure is identical to "Adding a Switch Before Discovery Using the CLI". When you execute the command, the system determines if the node exists and, if not, adds it. If the node exists, the system registers it.

Procedure

	Command or Action	Purpose
Step 1	[**no**] **system switch-id** *serial-number switch-id name* **pod** *id* **role leaf node-type tier-2-leaf**	Adds the switch to the pending registration list.

Adding a Switch Before Discovery Using the CLI

Use this procedure to add a switch to the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the CLI.

Note This procedure is identical to "Registering an Unregistered Switch Using the CLI". When you execute the command, the system determines if the node exists and, if not, adds it. If the node does exist, the system registers it.

[**no**] **system switch-id** *serial-number switch-id name* **pod** *id* **role leaf node-type tier-2-leaf**

Adds the switch to the pending registration list.

Removing a Switch to Maintenance Mode Using the CLI

Use this procedure to remove a switch to maintenance mode using the CLI.

Note While the switch is in maintenance mode, CLI 'show' commands on the switch show the front panel ports as being in the up state and the BGP protocol as up and running. The interfaces are actually shut and all other adjacencies for BGP are brought down, but the displayed active states allow for debugging.

[**no**]**debug-switch** *node_id or node_name*

Removes the switch to maintenance mode.

Inserting a Switch to Operation Mode Using the CLI

Use this procedure to insert a switch to operational mode using the CLI.

[**no**]**no debug-switch** *node_id or node_name*

Inserts the switch to operational mode.

Configuring a Remote Location Using the NX-OS Style CLI

In the ACI fabric, you can configure one or more remote destinations for exporting techsupport or configuration files.

SUMMARY STEPS

1. **configure**
2. [**no**] **remote path** *remote-path-name*
3. **user** *username*
4. **path** {**ftp** | **scp** | **sftp**} *host*[*:port*] [**remote-directory**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure** **Example:** `apic1# configure`	Enters global configuration mode.
Step 2	[**no**] **remote path** *remote-path-name* **Example:** `apic1(config)# remote path myFiles`	Enters configuration mode for a remote path.

	Command or Action	Purpose
Step 3	**user** *username* **Example:** `apic1(config-remote)# user admin5`	Sets the user name for logging in to the remote server. You are prompted for a password.
Step 4	**path** {**ftp** | **scp** | **sftp**} *host*[*:port*] [**remote-directory**] **Example:** `apic1(config-remote)# path sftp filehost.example.com:21 remote-directory /reports/apic`	Sets the path and protocol to the remote server. You are prompted for a password.

Examples

This example shows how to configure a remote path for exporting files.

```
apic1# configure
apic1(config)# remote path myFiles
apic1(config-remote)# user admin5
You must reset the password when modifying the path:
Password:
Retype password:
apic1(config-remote)# path sftp filehost.example.com:21 remote-directory /reports/apic
You must reset the password when modifying the path:
Password:
Retype password:
```

Finding Your Switch Inventory Using the NX-OS CLI

This section explains how to find your switch model and serial numbers using the NX-OS CLI.

Find your switch inventory as follows:

Example:

```
switch# show hardware
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Documents: http://www.cisco.com/en/US/products/ps9372/tsd_products_support_series_home.html
Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php

Software
  BIOS:      version 07.56
  kickstart: version 12.1(1h) [build 12.1(1h)]
```

```
  system:    version 12.1(1h) [build 12.1(1h)]
  PE:        version 2.1(1h)
  BIOS compile time:       06/08/2016
  kickstart image file is: /bootflash/aci-n9000-dk9.12.1.1h.bin
  kickstart compile time:  10/01/2016 20:10:40 [10/01/2016 20:10:40]
  system image file is:    /bootflash/auto-s
  system compile time:     10/01/2016 20:10:40 [10/01/2016 20:10:40]


Hardware
  cisco N9K-C93180YC-EX ("supervisor")
   Intel(R) Xeon(R) CPU  @ 1.80GHz with 16400384 kB of memory.
  Processor Board ID FDO20101H1W

  Device name: ifav41-leaf204
  bootflash:    62522368 kB

Kernel uptime is 02 day(s), 21 hour(s), 42 minute(s), 31 second(s)

Last reset at 241000 usecs after Sun Oct 02 01:27:25 2016
  Reason: reset-by-installer
  System version: 12.1(1e)
  Service: Upgrade

plugin
  Core Plugin, Ethernet Plugin
--------------------------------
Switch hardware ID information
--------------------------------

Switch is booted up
Switch type is : Nexus C93180YC-EX Chassis
Model number is N9K-C93180YC-EX
H/W version is 0.2010
Part Number is 73-15298-01
Part Revision is 1
Manufacture Date is Year 20 Week 10
Serial number is FDO20101H1W
CLEI code is 73-15298-01

--------------------------------
Chassis has one slot
--------------------------------

Module1 ok
  Module type is : 48x10/25G
  1 submodules are present
  Model number is N9K-C93180YC-EX
  H/W version is 0.2110
  Part Number is 73-17776-02
  Part Revision is 11
  Manufacture Date is Year 20 Week 10
  Serial number is FDO20101H1W
  CLEI code is 73-17776-02

GEM ok
  Module type is : 6x40/100G Switch
  1 submodules are present
  Model number is N9K-C93180YC-EX
  H/W version is 0.2110
  Part Number is 73-17776-02
  Part Revision is 11
  Manufacture Date is Year 20 Week 10
  Serial number is FDO20101H1W
```

```
  CLEI code is 73-17776-02

--------------------------------------
Chassis has  2 PowerSupply Slots
--------------------------------------

PS1 shut
  Power supply type is : 54.000000W 220v AC
  Model number is NXA-PAC-650W-PE
  H/W version is 0.0
  Part Number is 341-0729-01
  Part Revision is A0
  Manufacture Date is Year 19 Week 50
  Serial number is LIT19500ZEK
  CLEI code is 341-0729-01

PS2 ok
  Power supply type is : 54.000000W 220v AC
  Model number is NXA-PAC-650W-PE
  H/W version is 0.0
  Part Number is 341-0729-01
  Part Revision is A0
  Manufacture Date is Year 19 Week 50
  Serial number is LIT19500ZEA
  CLEI code is 341-0729-01

--------------------------------------
Chassis has  4 Fans
--------------------------------------

FT1 ok

 Fan1(sys_fan1)(fan_model:NXA-FAN-30CFM-F)                                  is inserted but info
is not available

FT2 ok

 Fan2(sys_fan2)(fan_model:NXA-FAN-30CFM-F)                                  is inserted but info
is not available

FT3 ok

 Fan3(sys_fan3)(fan_model:NXA-FAN-30CFM-F)                                  is inserted but info
is not available

FT4 ok

 Fan4(sys_fan4)(fan_model:NXA-FAN-30CFM-F)                                  is inserted but info
is not available


================================================================================
```

Verifying the Cisco APIC Cluster Using the CLI

Cisco Application Policy Infrastructure Controller (APIC) release 4.2.(1) introduces the **cluster_health** command, which enables you to verify the Cisco APIC cluster status step-by-step. The following output example demonstrates a scenario where everything is fine except for one node (ID 1002), which is inactive.

Note To use the **cluster_health** command, you must be logged in as admin.

To verify the cluster status:

```
F1-APIC1# cluster_health
Password:

Running...

Checking Wiring and UUID: OK
Checking AD Processes: Running
Checking All Apics in Commission State: OK
Checking All Apics in Active State: OK
Checking Fabric Nodes: Inactive switches: ID=1002(IP=10.1.176.66/32)
Checking Apic Fully-Fit: OK
Checking Shard Convergence: OK
Checking Leadership Degration: Optimal leader for all shards
Ping OOB IPs:
APIC-1: 172.31.184.12 - OK
APIC-2: 172.31.184.13 - OK
APIC-3: 172.31.184.14 - OK
Ping Infra IPs:
APIC-1: 10.1.0.1 - OK
APIC-2: 10.1.0.2 - OK
APIC-3: 10.1.0.3 - OK
Checking APIC Versions: Same (4.2(0.261a))
Checking SSL: OK

Done!
```

Table 1: Cluster_Health Verification Steps

Step	Description
Checking Wiring and UUID	Leaf switches provide infra connectivity between each Cisco APIC by detecting the Cisco APICs using LLDP. This step checks wiring issues between a leaf and a Cisco APIC that is detected during LLDP discovery. Any issues in here implies a leaf switch cannot provide infra connectivity for a Cisco APIC as it doesn't have a valid information. For example, a Cisco APIC UUID mismatch means the new APIC2 has a different UUID than the previously known APIC2. UUID – Universally Unique ID, or chassis ID in some outputs
Checking AD Processes	Cisco APIC clustering is handled by the Appliance Director process on each Cisco APIC. This step checks if the process is running correctly.
Checking All APICs in Commission State	To complete the Cisco APIC clustering, all Cisco APICs need to be commissioned.

Step	Description
Checking All APICs in Active State	To complete the Cisco APIC clustering, all commissioned Cisco APICs need to be active. If it is not active, the Cisco APIC may not be up yet.
Checking Fabric Nodes: Inactive switches	The Cisco APIC's communication are through infra connectivity provided by leaf and spine switches. This step checks inactive switches to ensure switches are providing infra connectivity.
Checking APIC Fully-Fit	When Cisco APICs have established IP reachability to each other through infra network, it will synchronize its database to each other. When the synchronization completes, the status of all Cisco APICs become "Fully-Fit." Otherwise, the status will be "Data Layer Partially Diverged," and so on.
Checking Shard Convergence	When Cisco APICs are not fully-fit, database shards need to be checked to see which service is not fully synchronized. If there is any service that has problems in synchronization, you may reach out to Cisco TAC for further troubleshooting.
Checking Leadership Degration	In ACI, each database shard has one leader shard distributed to each Cisco APIC in the cluster. This step shows if all shards have an optimal leader. If there is an issue in here when all Cisco APICs are up, you may reach out to Cisco TAC for further troubleshooting.
Ping OOB IPs	This step is to check if all Cisco APICs are up and operational by pinging the OOB IP which is configured separately from clustering.
Ping Infra IPs	This step is to check if there is infra connectivity between each Cisco APIC. Cisco APIC clustering is performed through infra connectivity instead of OOB.
Checking APIC Versions	All Cisco APICs should be on a same version to complete clustering.
Checking SSL	All Cisco APICs need to have a valid SSL that should be built-in when a Cisco APIC is shopped as an appliance. Without a valid SSL, the server cannot operate the Cisco APIC OS correctly.