



Cisco ACI Cluster Management

- [Overview, page 1](#)
- [Expanding the Cisco APIC Cluster, page 1](#)
- [Contracting the Cisco APIC Cluster, page 2](#)
- [Cluster Management Guidelines, page 2](#)
- [Expanding the Cluster Examples, page 5](#)
- [Contracting the Cluster Examples, page 6](#)
- [Commissioning and Decommissioning Cisco APIC Controllers, page 8](#)
- [Replacing a Cisco APIC in a Cluster Using the CLI, page 9](#)
- [Shutting Down the APICs in a Cluster, page 9](#)
- [Cold Standby, page 10](#)

Overview

This article provides guidelines and examples related to expanding and contracting the Cisco APIC cluster.

Expanding the Cisco APIC Cluster

Expanding the Cisco APIC cluster is the operation to increase any size mismatches, from a cluster size of N to size N+1, within legal boundaries. The operator sets the administrative cluster size and connects the APICs with the appropriate cluster IDs, and the cluster performs the expansion.

During cluster expansion, regardless of in which order you physically connect the APIC controllers, the discovery and expansion takes place sequentially based on the APIC ID numbers. For example, APIC2 is discovered after APIC1, and APIC3 is discovered after APIC2 and so on until you add all the desired APICs to the cluster. As each sequential APIC is discovered, a single data path or multiple data paths are established, and all the switches along the path join the fabric. The expansion process continues until the operational cluster size reaches the equivalent of the administrative cluster size.

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.

Cluster Management Guidelines

The APIC cluster is comprised of multiple APIC controllers that provide operators a unified real time monitoring, diagnostic, and configuration management capability for the ACI fabric. To assure optimal system performance, follow the guidelines below for making changes to the APIC cluster.



Note

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 APIC controllers' health status in the cluster is not "fully fit", remedy that situation before proceeding. Also, assure that cluster controllers added to the APIC are running the same version of firmware as the other controllers in the APIC cluster.

Follow these general guidelines when managing clusters:

- It is recommended that you have at least 3 active APICs in a cluster, and one or more standby APICs.
- Disregard cluster information from APICs that are not currently in the cluster; they do not provide accurate cluster information.
- Cluster slots contain an APIC `ChassisID`. Once you configure a slot, it remains unavailable until you decommission the APIC with the assigned `ChassisID`.
- If an 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 an APIC, first ensure that you have a healthy cluster. After verifying the health of the APIC Cluster, choose the APIC you intend to shut down. After the APIC has shutdown, move the 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 APIC at a time.

- When an APIC cluster is split into two or more groups, the ID of a node is changed and the changes are not synchronized across all APICs. This can cause inconsistency in the node IDs between APICs and also the affected leaf nodes may not appear in the inventory in the APIC GUI. When you split an APIC cluster, decommission the affected leaf nodes from APIC and register them again, so that the inconsistency in the node IDs is resolved and the health status of the APICs in a cluster are in a fully fit state.

- Before configuring the APIC cluster, ensure that all the APICs are running the same firmware version. Initial clustering of APICs running differing versions is an unsupported operation and may cause problems within the cluster.

This section contains the following topics:

Expanding the APIC Cluster Size

Follow these guidelines to expand the APIC cluster size:

- Schedule the cluster expansion at a time when the demands of the fabric workload will not be impacted by the cluster expansion.
- If one or more of the APIC controllers' health status in the cluster is not "fully fit", remedy that situation before proceeding.
- Stage the new APIC controller(s) according to the instructions in their hardware installation guide. Verify in-band connectivity with a PING test.
- Increase the cluster target size to be equal to the existing cluster size controller count plus the new controller count. For example, if the existing cluster size controller count is 3 and you are adding 3 controllers, set the new cluster target size to 6. The cluster proceeds to sequentially increase its size one controller at a time until all new the controllers are included in the cluster.



Note Cluster expansion stops if an existing APIC controller becomes unavailable. Resolve this issue before attempting to proceed with the cluster expansion.

- Depending on the amount of data the APIC must synchronize upon the addition of each appliance, the time required to complete the expansion could be more than 10 minutes per appliance. Upon successful expansion of the cluster, the APIC operational size and the target size will be equal.



Note Allow the APIC to complete the cluster expansion before making additional changes to the cluster.

Reducing the APIC Cluster Size

Follow these guidelines to reduce the APIC cluster size and decommission the APIC controllers that are removed from the cluster:



Note

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

- Reducing the cluster size increases the load on the remaining APIC controllers. Schedule the APIC controller 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 APIC controllers' 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 controller one by one until the cluster reaches the new lower target size.

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



Note After decommissioning an APIC controller from the cluster, power it down and disconnect it from fabric. Before returning it to service, do a wiped clean back to factory reset.

- Cluster synchronization stops if an existing APIC controller becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.
- Depending on the amount of data the 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 APIC to complete the cluster synchronization accordingly before making additional changes to the cluster.

Replacing Cisco APIC Controllers in the Cluster

Follow these guidelines to replace Cisco APIC controllers:

- If the health status of any Cisco APIC controller in the cluster is not **Fully Fit**, remedy the situation before proceeding.
- Schedule the Cisco APIC controller replacement at a time when the demands of the fabric workload will not be impacted by the cluster synchronization.
- Make note of the initial provisioning parameters and image used on the Cisco APIC controller that will be replaced. The same parameters and image must be used with the replacement controller. The Cisco APIC proceeds to synchronize the replacement controller with the cluster.



Note Cluster synchronization stops if an existing Cisco APIC controller becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.

- You must choose a Cisco APIC controller that is within the cluster and not the controller that is being decommissioned. For example: Log in to Cisco APIC1 or APIC2 to invoke the shutdown of APIC3 and decommission APIC3.
- Perform the replacement procedure in the following order:
 - 1 Make note of the configuration parameters and image of the APIC being replaced.

- 2 Decommission the APIC you want to replace (see [Decommissioning a Cisco APIC Controller in the Cluster Using the GUI](#), on page 8)
 - 3 Commission the replacement APIC using the same configuration and image of the APIC being replaced (see [Commissioning a Cisco APIC Controller in the Cluster Using the GUI](#), on page 8)
- Stage the replacement Cisco APIC controller according to the instructions in its hardware installation guide. Verify in-band connectivity with a PING test.



Note Failure to decommission Cisco APIC controllers before attempting their replacement will preclude the cluster from absorbing the replacement controllers. Also, before returning a decommissioned Cisco APIC controller to service, do a wiped clean back to factory reset.

- Depending on the amount of data the Cisco APIC must synchronize upon the replacement of a controller, the time required to complete the replacement could be more than 10 minutes per replacement controller. Upon successful synchronization of the replacement controller with the cluster, the Cisco APIC operational size and the target size will remain unchanged.



Note Allow the Cisco APIC to complete the cluster synchronization before making additional changes to the cluster.

- The UUID and fabric domain name persist in a Cisco APIC controller across reboots. However, a clean back-to-factory reboot removes this information. If a Cisco APIC controller is to be moved from one fabric to another, a clean back-to-factory reboot must be done before attempting to add such a controller to a different Cisco ACI fabric.

Expanding the Cluster Examples

Expanding the APIC Cluster Using the GUI

Step 1 On the menu bar, choose **SYSTEM > Controllers**. In the **Navigation** pane, expand **Controllers > apic_controller_name > Cluster**.

You must choose an **apic_controller_name** that is within the cluster that you wish to expand.

In the **Work** pane, the cluster details are displayed. This includes the current cluster target and current sizes, the administrative, operational, and health states of each controller in the cluster.

Step 2 Verify that the health state of the cluster is **Fully Fit** before you proceed with contracting the cluster.

Step 3 In the **Work** pane, click **Actions > Change Cluster Size**.

Step 4 In the **Change Cluster Size** dialog box, in the **Target Cluster Administrative Size** field, choose the target number to which you want to expand the cluster. Click **Submit**.

Note It is not acceptable to have a cluster size of two APIC controllers. A cluster of one, three, or more APIC controllers is acceptable.

- Step 5** In the **Confirmation** dialog box, click **Yes**.
In the **Work** pane, under **Properties**, the **Target Size** field must display your target cluster size.
- Step 6** Physically connect all the APIC controllers that are being added to the cluster.
In the **Work** pane, in the **Cluster > Controllers** area, the APIC controllers are added one by one and displayed in the sequential order starting with N + 1 and continuing until the target cluster size is achieved.
- Step 7** Verify that the APIC controllers are in operational state, and the health state of each controller is **Fully Fit**.
-

Expanding the APIC Cluster Using the REST API

The cluster drives its actual size to the target size. If the target size is higher than the actual size, the cluster size expands.

- Step 1** Set the target cluster size to expand the APIC cluster size.

Example:

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=3/>
```

- Step 2** Physically connect the APIC controllers that you want to add to the cluster.
-

Contracting the Cluster Examples

Contracting the APIC Cluster Using the GUI

- Step 1** On the menu bar, choose **SYSTEM > Controllers**. In the **Navigation** pane, expand **Controllers > apic_controller_name > Cluster**.
You must choose an **apic_controller_name** that is within the cluster and not the controller that is being decommissioned.
In the **Work** pane, the cluster details are displayed. This includes the current cluster target and current sizes, the administrative, operational, and health states of each controller in the cluster.
- Step 2** Verify that the health state of the cluster is **Fully Fit** before you proceed with contracting the cluster.
- Step 3** In the **Work** pane, click **Actions > Change Cluster Size**.
- Step 4** In the **Change Cluster Size** dialog box, in the **Target Cluster Administrative Size** field, choose the target number to which you want to contract the cluster. Click **Submit**.
Note It is not acceptable to have a cluster size of two APIC controllers. A cluster of one, three, or more APIC controllers is acceptable.
- Step 5** In the **Work** pane, in the **Controllers** area, choose the APIC that is last in the cluster.

Example:

In a cluster of three, the last in the cluster is three as identified by the controller ID.

Step 6

Click **Actions** > **Decommission**. The **Confirmation** dialog box displays. Click **Yes**.

The decommissioned controller displays **Unregistered** in the **Operational State** column. The controller is then taken out of service and not visible in the **Work** pane any longer.

Step 7

Repeat the earlier step to decommission the controllers one by one for all the APICs in the cluster in the appropriate order of highest controller ID number to the lowest.

Note The operation cluster size shrinks only after the last appliance is decommissioned, and not after the administrative size is changed. Verify after each controller is decommissioned that the operational state of the controller is unregistered, and the controller is no longer in service in the cluster.

You should be left with the remaining controllers in the APIC cluster that you desire.

Contracting the APIC Cluster Using the REST API

The cluster drives its actual size to the target size. If the target size is lower than the actual size, the cluster size contracts.

Step 1

Set the target cluster size so as to contract the APIC cluster size.

Example:

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=1/>
```

Step 2

Decommission APIC3 on APIC1 for cluster contraction.

Example:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=3 adminSt='out-of-service'/>
```

Step 3

Decommission APIC2 on APIC1 for cluster contraction.

Example:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=2 adminSt='out-of-service'/>
```

Commissioning and Decommissioning Cisco APIC Controllers

Commissioning a Cisco APIC Controller in the Cluster Using the GUI

-
- Step 1** From the menu bar, choose **SYSTEM > Controllers**.
- Step 2** In the **Navigation** pane, expand **Controllers > apic_controller_name > Cluster as Seen by Node**.
- Step 3** From the **Work** pane, verify in the **Active Controllers** summary table that the cluster **Health State** is **Fully Fit** before continuing.
- Step 4** From the **Work** pane, click the decommissioned controller that displaying **Unregistered** in the **Operational State** column.
The controller is highlighted.
- Step 5** From the **Work** pane, click **Actions > Commission**.
- Step 6** In the **Confirmation** dialog box, click **Yes**.
- Step 7** Verify that the commissioned Cisco APIC controller is in the operational state and the health state is **Fully Fit**.
-

Decommissioning a Cisco APIC Controller in the Cluster Using the GUI

-
- Step 1** On the menu bar, choose **System > Controllers**.
- Step 2** In the **Navigation** pane, expand **Controllers > apic_controller_name > Cluster as Seen by Node**.
- Step 3** In the **Work** pane, verify that the **Health State** in the **Active Controllers** summary table indicates the cluster is **Fully Fit** before continuing.
- Step 4** In the **Navigation** pane, click an **apic_controller_name** that is within the cluster and not the controller that is being decommissioned.
The controller details appear in the **Work** pane.
- Step 5** In the **Work** pane, click **Actions > Decommission**.
The **Confirmation** dialog box displays.
- Step 6** Click **Yes**.
The decommissioned controller displays **Unregistered** in the **Operational State** column. The controller is then taken out of service and no longer visible in the **Work** pane.
- Note**
- The operation cluster size shrinks only after the last appliance is decommissioned, and not after the administrative size is changed. Verify after each controller is decommissioned that the operational state of the controller is unregistered, and the controller is no longer in service in the cluster.
 - After decommissioning the APIC controller, you must reboot the APIC for Layer 4 to Layer 7 services. Reboot must be done before commissioning back the controller.
-

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 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 APIC 2 or 3, you will not be asked for a password.

Step 1 Identify the APIC that you want to replace.

Step 2 Decommission the APIC using the **controller *controller-id* decommission** command.

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

Step 3 If you want to recommission the same APIC, follow these steps:

- Restart the APIC using the **acdiag reboot** command.
- Verify that the APIC boots without error.
- Commission the APIC using the **controller *controller-id* commission** command.
- Allow several minutes for the new APIC information to propagate to the rest of the cluster.

Step 4 If you want to commission a new APIC, follow these steps:

- Disconnect the APIC from the fabric.
 - Connect the replacement APIC to the fabric.
 - Commission the APIC using the **controller *controller-id* commission** command.
 - Boot the new APIC.
 - Allow several minutes for the new APIC information to propagate to the rest of the cluster.
-

Shutting Down the APICs in a Cluster

Shutting Down all the APICs in a Cluster

Before you shutdown all the APICs in a cluster, ensure that the APIC cluster is in a healthy state and all the APICs are showing fully fit. Once you start this process, we recommend that no configuration changes are done during this process. Use this procedure to gracefully shut down all the APICs in a cluster.

Step 1 Log in to Cisco APIC with appliance ID 1.

Step 2 On the menu bar, choose **System > Controllers**.

Step 3 In the Navigation pane, expand **Controllers > apic_controller_name**.
You must select the third APIC in the cluster.

- Step 4** Right-click the controller and click **Shutdown**.
- Step 5** Repeat the steps to shutdown the second APIC in the cluster.
- Step 6** Log in to Cisco IMC of the first APIC in the cluster to shutdown the APIC.
- Step 7** Choose **Server > Server Summary > Shutdown Server**.
You have now shutdown all the three APICs in a cluster.
-

Bringing Back the APICs in a Cluster

Use this procedure to bring back the APICs in a cluster.

-
- Step 1** Log in to Cisco IMC of the first APIC in the cluster.
- Step 2** Choose **Server > Server Summary > Power On** to power on the first APIC.
- Step 3** Repeat the steps to power on the second APIC and then the third APIC in the cluster.
After all the APICs are powered on, ensure that all the APICs are in a fully fit state. Only after verifying that the APICs are in a fully fit state, you must make any configuration changes on the APIC.
-

Cold Standby

About Cold Standby for APIC Cluster

The Cold Standby functionality for an APIC cluster enables you to operate the APICs in a cluster in an Active/Standby mode. In an APIC cluster, the designated active APICs share the load and the designated standby APICs can act as a replacement for any of the APICs in an active cluster.

As an admin user, you can set up the Cold Standby functionality when the APIC is launched for the first time. We recommend that you have at least three active APICs in a cluster, and one or more standby APICs. As an admin user, you can initiate the switch over to replace an active APIC with a standby APIC.

Important Notes

- The standby APIC is automatically updated with firmware updates to keep the backup APIC at same firmware version as the active cluster.
- During an upgrade process, once all the active APICs are upgraded, the standby APIC is also be upgraded automatically.
- Temporary IDs are assigned to standby APICs. After a standby APIC is switched over to an active APIC, a new ID is assigned.
- Admin login is not enabled on standby APIC. To troubleshoot Cold Standby, you must log in to the standby using SSH as *rescue-user*.

- During switch over the replaced active APIC is powered down, to prevent connectivity to the replaced APIC.
- Switch over fails under the following conditions:
 - If there is no connectivity to the standby APIC.
 - If the firmware version of the standby APIC is not the same as that of the active cluster.
- After switching over a standby APIC to active, if it was the only standby, you must configure a new standby.
- The following limitations are observed for retaining out of band address for standby APIC after a fail over.
 - Standby (new active) APIC may not retain its out of band address if more than 1 active APICs are down or unavailable.
 - Standby (new active) APIC may not retain its out of band address if it is in a different subnet than active APIC. This limitation is only applicable for APIC release 2.x.
 - Standby (new active) APIC may not retain its IPv6 out of band address. This limitation is not applicable starting from APIC release 3.1x.
 - Standby (new active) APIC may not retain its out of band address if you have configured non Static OOB Management IP address policy for replacement (old active) APIC.

**Note**

In case you observe any of the limitations, in order to retain standby APICs out of band address, you must manually change the OOB policy for replaced APIC after the replace operation is completed successfully.

- We recommend keeping standby APICs in same POD as the active APICs it may replace.
- There must be three active APICs in order to add a standby APIC.
- The standby APIC does not participate in policy configuration or management.
- No information is replicated to standby controllers, including admin credentials.

Verifying Cold Standby Status

Verifying Cold Standby Status Using 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

Verifying Cold Standby Status Using GUI

- 1 On the menu bar, choose **System > Controllers**.
- 2 In the **Navigation** pane, expand **Controllers > apic_controller_name > Cluster as Seen by Node**.
- 3 In the **Work** pane, the standby controllers are displayed under **Standby Controllers**.

Switching Over Active APIC with Standby APIC Using GUI

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

Before You Begin

-
- Step 1** On the menu bar, choose **System > Controllers**.
 - Step 2** In the **Navigation** pane, expand **Controllers > apic_controller_name > Cluster as Seen by Node**.
The **apic_controller_name** should be other than the name of the controller being replaced.
 - Step 3** In the **Work** pane, verify that the **Health State** in the **Active Controllers** summary table indicates the active controller is **Fully Fit** before continuing.
 - Step 4** Click an **apic_controller_name** that you want to switch over.
 - Step 5** In the **Work** pane, click **Actions > Replace**.
The **Replace** dialog box displays.
 - Step 6** Choose the **Backup Controller** from the drop-down list and click **Submit**.
It may take several minutes to switch over an active APIC with a standby APIC and for the system to be registered as active.
 - Step 7** Verify the progress of the switch over in the **Failover Status** field in the **Active Controllers** summary table.
-

Switching Over Active APIC with Standby APIC Using CLI

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

SUMMARY STEPS

1. `replace-controller replace ID number Backup serial number`
2. `replace-controller reset ID number`

DETAILED STEPS

	Command or Action	Purpose
Step 1	replace-controller replace <i>ID number Backup serial number</i> Example: <pre>apic1#replace-controller replace 2 FCH1804V27L Do you want to replace APIC 2 with a backup? (Y/n): Y</pre>	Replaces an active APIC with an standby APIC.
Step 2	replace-controller reset <i>ID number</i> Example: <pre>apic1# replace-controller reset 2 Do you want to reset failover status of APIC 2? (Y/n): Y</pre>	Resets fail over status of the active controller.

Switching Over Active APIC with Standby APIC Using REST API

Use this procedure to switch over an active APIC with standby APIC using REST API.

Switch over active APIC with standby APIC.

URL for POST: `https://ip address/api/node/mo/topology/pod-initiator_pod_id/node-initiator_id/av.xml`

Body: `<infraWiNode id=outgoing_apic_id targetMbSn=backup-serial-number/>`

where `initiator_id` = id of an active APIC other than the APIC being replaced.

`pod-initiator_pod_id` = pod ID of the active APIC

`backup-serial-number` = serial number of standby APIC

Example:

```
https://ip address/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=2 targetMbSn=FCH1750V00Q/>
```
