



Cisco APIC Cluster Management

- [APIC Cluster Overview, on page 1](#)
- [Expanding the Cisco APIC Cluster, on page 1](#)
- [Contracting the Cisco APIC Cluster, on page 2](#)
- [Cluster Management Guidelines, on page 2](#)
- [Expanding the APIC Cluster Using the GUI, on page 5](#)
- [Contracting the APIC Cluster Using the GUI, on page 6](#)
- [Commissioning and Decommissioning Cisco APIC Controllers, on page 7](#)
- [Shutting Down the APICs in a Cluster, on page 8](#)
- [Cold Standby, on page 9](#)

APIC Cluster Overview

The Cisco Application Policy Infrastructure Controller (APIC) appliance is deployed in a cluster. A minimum of three controllers are configured in a cluster to provide control of the Cisco ACI fabric. The ultimate size of the controller cluster is directly proportionate to the size of the ACI deployment and is based on transaction-rate requirements. Any controller in the cluster can service any user for any operation, and a controller can be transparently added to or removed from the cluster.

This section provides guidelines and examples related to expanding, contracting, and recovering the 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 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. In most cases, we recommend a cluster size of 3, 5, or 7 Cisco APICs. We recommend 4 Cisco APICs for a two site multi-pod fabric that has between 80 to 200 leaf switches.
- 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 an Cisco APIC cluster is split into two or more groups, the ID of a node is changed and the changes are not synchronized across all Cisco APICs. This can cause inconsistency in the node IDs between Cisco APICs and also the affected leaf nodes may not appear in the inventory in the Cisco APIC GUI. When you split a Cisco APIC cluster, decommission the affected leaf nodes from a Cisco 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 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.

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 in the Cluster Using the GUI, on page 7](#))
 3. Commission the replacement APIC using the same configuration and image of the APIC being replaced (see [Commissioning a Cisco APIC in the Cluster Using the GUI, on page 7](#))
- 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 APIC Cluster Using the GUI

This procedure adds one or more Cisco Application Policy Infrastructure Controllers (APICs) to an existing cluster.

Before you begin

You must first set up any Cisco APIC that you will add to the cluster. For information about setting up a Cisco APIC, see [Setting up the Cisco APIC](#).

-
- Step 1** On the menu bar, choose **System > Controllers**.
- Step 2** In the **Navigation** pane, expand **Controllers > apic_name > Cluster as Seen by Node**.
For *apic_name*, you must choose a Cisco APIC that is within the cluster that you wish to expand.
The **Cluster as Seen by Node** window appears in the **Work** pane with the **APIC Cluster** and **Standby APIC** tabs. In the **APIC Cluster** tab, the controller details appear. This includes the current cluster target and current sizes, the administrative, operational, and health states of each controller in the cluster.
- Step 3** Verify that the health state of the cluster is **Fully Fit** before you proceed with contracting the cluster.
- Step 4** In the **Work** pane, click **Actions > Change Cluster Size**.
- Step 5** 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** You cannot have a cluster size of two Cisco APICs. You can have a cluster of one, three, or more Cisco APICs.
- Step 6** In the **Confirmation** dialog box, click **Yes**.
In the **Work** pane, under **Properties**, the **Target Size** field must display your target cluster size.
- Step 7** Physically connect all the Cisco APICs that are being added to the cluster.
In the **Work** pane, in the **Cluster > Controllers** area, the Cisco APICs 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 8** Verify that the Cisco APICs are in operational state, and the health state of each controller is **Fully Fit**.
-

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 as Seen by Node**.
You must choose an *apic_name* that is within the cluster and not the controller that is being decommissioned.
The **Cluster as Seen by Node** window appears in the **Work** pane with the **APIC Cluster** and **Standby APIC** tabs. In the **APIC Cluster** tab, the controller details appear. 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 APICs. A cluster of one, three, or more APICs is acceptable.
- Step 5** From the **Active Controllers** area of the **Work** pane, 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 Right-click on the controller you want to decommission and choose **Decommission**. When 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.

Commissioning and Decommissioning Cisco APIC Controllers

Commissioning a Cisco APIC 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**.

The **Cluster as Seen by Node** window appears in the **Work** pane with the **APIC Cluster** and **Standby APIC** tabs. In the **APIC Cluster** tab, the controller details appear. This includes the current cluster target and current sizes, the administrative, operational, and health states of each controller in the cluster.

Step 3 From the **APIC Cluster** tab of 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, right-click the decommissioned controller that is displaying **Unregistered** in the **Operational State** column and choose **Commission**.

The controller is highlighted.

Step 5 In the **Confirmation** dialog box, click **Yes**.

Step 6 Verify that the commissioned Cisco APIC is in the operational state and the health state is **Fully Fit**.

Decommissioning a Cisco APIC in the Cluster Using the GUI

This procedure decommissions a Cisco Application Policy Infrastructure Controller (APIC) in the cluster.



Note 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.

-
- Step 1** On the menu bar, choose **System** > **Controllers**.
- Step 2** In the **Navigation** pane, expand **Controllers** > **apic_name** > **Cluster as Seen by Node**.
You must choose an **apic_name** that is within the cluster and not the controller that is being decommissioned.
The **Cluster as Seen by Node** window appears in the **Work** pane with the controller details and the **APIC Cluster** and **Standby APIC** tabs.
- Step 3** In the **Work** pane, verify in the **APIC Cluster** tab that the **Health State** in the **Active Controllers** summary table indicates the cluster is **Fully Fit** before continuing.
- Step 4** In the **Active Controllers** table located in the **APIC Cluster** tab of the **Work** pane, right-click on the controller you want to decommission and choose **Decommission**.
The **Confirmation** dialog box displays.
- Step 5** 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**
- After decommissioning a Cisco APIC from the cluster, power the controller down and disconnect it from the fabric. Before returning the Cisco APIC to service, perform a factory reset on the controller.
 - 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 Cisco APIC, you must reboot the controller for Layer 4 to Layer 7 services. You must perform the reboot before re-commissioning the controller.
-

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 a Cisco APIC Cluster

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

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

Important Notes

- The standby Cisco APICs are automatically updated with firmware updates to keep the backup Cisco APIC at same firmware version as the active cluster.
- During an upgrade process, after all the active Cisco APICs are upgraded, the standby Cisco APICs are also upgraded automatically.
- Temporary IDs are assigned to the standby Cisco APICs. After a standby Cisco APIC is switched over to an active Cisco APIC, a new ID is assigned.
- The admin login is not enabled on the standby Cisco APICs. To troubleshoot a Cold Standby Cisco APIC, you must log in to the standby using SSH as *rescue-user*.
- During the switch over, the replaced active Cisco APIC is powered down to prevent connectivity to the replaced Cisco APIC.
- Switch over fails under the following conditions:

- If there is no connectivity to the standby Cisco APIC.
- If the firmware version of the standby Cisco APIC is not the same as that of the active cluster.
- After switching over a standby Cisco APIC to be 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 the standby Cisco APIC after a fail over:
 - The standby (new active) Cisco APIC may not retain its out of band address if more than 1 active Cisco APICs are down or unavailable.
 - The standby (new active) Cisco APIC may not retain its out of band address if it is in a different subnet than the active Cisco APIC. This limitation is only applicable for Cisco APIC release 2.x.
 - The standby (new active) Cisco APIC may not retain its IPv6 out of band address. This limitation is not applicable starting from Cisco APIC release 3.1x.
 - The standby (new active) Cisco APIC may not retain its out of band address if you have configured a non-static OOB management IP address policy for the replacement (old active) Cisco APIC.
 - The standby (new active) Cisco APIC may not retain its out of band address if it is not in a pod that has an active Cisco APIC.



Note If you want to retain the standby Cisco APIC's out of band address despite the limitations, you must manually change the OOB policy for the replaced Cisco APIC after the replace operation had completed successfully.

- There must be three active Cisco APICs to add a standby Cisco APIC.
- The standby Cisco APIC does not participate in policy configuration or management.
- No information is replicated to the standby Cisco APICs, not even the administrator credentials.

Verifying Cold Standby Status Using the 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 an Active APIC with a Standby APIC Using the GUI

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

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 controllers other than the one being replaced are **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.
-

