



Cisco APIC Cluster Management

This section explains how to expand, contract, commission, and decommission Cisco APIC clusters. For more information about Cisco APIC clusters, see the *Cisco APIC Cluster Management* document:

<http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html>

This section contains the following topics:

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

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:

- Cisco recommends that you have at least 3 active APICs in a cluster, along with additional standby APICs. A cluster size of 3, 5, or 7 APICs is recommended. A cluster size of 4 or 6 APICs is not recommended.
- 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

Procedure

- 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_controller_name** that is within the cluster that you wish to expand.
- The **Cluster as Seen by Node** window appears in the **Work** pane with three tabs: **APIC Cluster**, **APIC-X**, and **Standby APIC**. 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 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.

Procedure

- 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

Procedure

- 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_controller_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 three tabs: **APIC Cluster**, **APIC-X**, and **Standby APIC**. 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 APIC controllers. A cluster of one, three, or more APIC controllers 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 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.

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.

Procedure

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

Procedure

- 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 three tabs: **APIC Cluster**, **APIC-X**, and **Standby APIC**. 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 controller is in the operational state and the health state is **Fully Fit**.
-

Decommissioning a Cisco APIC Controller in the Cluster Using the GUI

Procedure

- 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**. You must choose an **apic_controller_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 three tabs: **APIC Cluster**, **APIC-X**, and **Standby APIC**.
- 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 an APIC controller from the cluster, power the APIC controller down and disconnect it from the fabric. Before returning the APIC controller to service, perform a factory reset on the APIC 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 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.

Before you begin

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

Procedure

-
- Step 1** Identify the APIC that you want to replace.
- Step 2** Note the configuration details of the APIC to be replaced by using the **acdiag avread** command.
- Step 3** 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 4** To commission the new APIC, follow these steps:
- a) Disconnect the old APIC from the fabric.
 - b) Connect the replacement APIC to the fabric.

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

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

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

Verifying the Cisco APIC Cluster Using the CLI

Cisco 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.

Procedure

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	<p>Leaf switches provide infra connectivity between each APIC by detecting APICs via LLDP. This step checks wiring issues between a leaf and an APIC that is detected during LLDP discovery.</p> <p>Any issues in here implies a leaf switch cannot provide infra connectivity for an APIC as it doesn't have a valid information. For example, APIC UUID mismatch means the new APIC2 has a different UUID than the previously known APIC2.</p> <p>UUID – Universally Unique ID, or chassis ID in some outputs</p>
Checking AD Processes	APIC clustering is handled by the Appliance Director process on each APIC. This step checks if the process is running correctly.
Checking All Apics in Commission State	To complete the APIC clustering, all APICs need to be commissioned.
Checking All Apics in Active State	To complete the APIC clustering, all commissioned APICs need to be active. If it is not active, the APIC may not be up yet.
Checking Fabric Nodes: Inactive switches	The 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 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 APICs become Fully-Fit. Otherwise, the status will be Data Layer Partially Diverged, and so on.
Checking Shard Convergence	When 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 APIC in the cluster. This step shows if all shards have an optimal leader. If there is an issue in here when all APICs are up, you may reach out to Cisco TAC for further troubleshooting.

Step	Description
Ping OOB IPs	This step is to check if all 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 APIC. APIC clustering is performed through infra connectivity instead of OOB.
Checking APIC Versions	All APICs should be on a same version to complete clustering.
Checking SSL	All APICs need to have a valid SSL that should be built-in when an APIC is shipped as an appliance. Without a valid SSL, the server cannot operate APIC OS correctly.
