



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 6](#)
- [Expanding the APIC Cluster Using the Add Node Option, on page 6](#)
- [Contracting the APIC Cluster Using the GUI, on page 9](#)
- [Contracting the APIC Cluster Using the Delete Node Option, on page 9](#)
- [Commissioning and Decommissioning Cisco APIC Controllers, on page 10](#)
- [Shutting Down the APICs in a Cluster, on page 14](#)
- [Cold Standby, on page 15](#)
- [Migration of APICs, on page 17](#)

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. 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 decommissioned, 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.

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 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.

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 12](#))
 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 10](#))
- 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 APICs to an existing cluster. This procedure is applicable for releases prior to Cisco APIC release 6.0(2). For expanding a cluster in release 6.0(2), you can use the **Add Node** option, as detailed in the subsequent procedure.

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**.
-

Expanding the APIC Cluster Using the Add Node Option

Use this procedure on an existing Cisco Application Policy Infrastructure Controller (APIC) cluster to expand the cluster using the **Add Node** option, which was introduced in Cisco APIC release 6.0(2). To expand a cluster in Cisco APIC releases prior to 6.0(2), see the previous procedure.

The **Add Node** option is a simpler and direct method to add a Cisco APIC to a cluster.

Before you begin

- Ensure the node *to-be-added* is a *clean* node or is in *factory-reset* state.
- Check the current **Cluster Size** in the **General** pane. If it is N , after successful node addition, the size will be $N+1$.

Step 1 On the menu bar, choose **System > Controllers**. In the **Navigation** pane, expand **Controllers > apic_controller_name > Cluster as Seen by Node**.

Step 2 In the **Active Controllers** pane, click the **Actions** button and select the **Add Node** option.
The **Add Node** screen is displayed.

Step 3 Enter the following details in the **Add Node** screen:
Select the **Controller Type**. Based on your selection, proceed to the relevant substep.
Put a check in the **Enabled** box if you need to support IPv6 addresses.

a) When the Controller Type is **Physical**:

- CIMC details pane
 - IP Address: Enter the CIMC IP address.
 - Username: Enter the username to access CIMC.
 - Password: Enter the password to access CIMC.
 - Click **Validate**. *Validation success* is displayed on successful authentication.

This pane appears only if you configured CIMC. If you did not configure CIMC, instead perform the physical APIC login step of the [Bringing up the Cisco APIC Cluster Using the GUI](#) procedure (step 1b) on the new node to configure out-of-band management.

- General pane
 - Name: Enter a name for the controller.
 - Admin Password: Enter the admin password for the controller.
 - Controller ID: This is auto-populated based on the existing cluster size. If the current cluster size is N , the controller ID is displayed as $N+1$.
 - Serial Number: This is auto-populated after CIMC validation.
 - Force Add: Put a check in the **Enabled** box to add a Cisco APIC that has a release earlier than 6.0(2).
- Out of Band Network pane
 - IPv4 Address: The address is auto-populated.
 - IPv4 Gateway: The gateway address is auto-populated.

Note If you put a check in the **Enabled** box for IPv6 earlier, enter the IPv6 address and gateway.

- Infra Network pane
 - IPv4 Address: Enter the infra network IP address.
 - IPv4 Gateway: Enter the infra network IP address of the gateway.
 - VLAN: Enter a VLAN ID.

b) When the Controller Type is **Virtual**:

- Management IP pane
 - IP Address: Enter the management IP address.
 - Note** The management IP addresses are defined during the deployment of the virtual machines using ESXi/AWS.
 - Enter the username for the virtual APIC.
 - Enter the password for the virtual APIC.
 - Click **Validate**. *Validation success* is displayed on successful authentication.
- General pane
 - Name: User-defined name for the controller.
 - Controller ID: This is auto-populated based on the existing cluster size. If the current cluster size is N , the controller ID is displayed as $N+1$.
 - Serial Number: The serial number of the virtual machine is auto-populated.
 - Force Add: Put a check in the **Enabled** box to add a Cisco APIC that has a release earlier than 6.0(2).
- Out of Band Network pane
 - IPv4 Address: The IP address is auto-populated.
 - IPv4 Gateway: The gateway IP address is auto-populated.
 - Note** If you put a check in the **Enabled** box for IPv6 earlier, enter the IPv6 address and gateway.
- Infra Network pane
 - IPv4 Address: Enter the infra network address.
 - IPv4 Gateway: Enter the IP address of the gateway.
 - VLAN: (Applicable only for *remotely attached* virtual APIC- ESXi) Enter the interface VLAN ID to be used.
 - Note** The Infra L3 Network pane is not displayed when the virtual APIC is deployed using AWS.

Step 4 Click **Apply**.

What to do next

The newly added controller appears in the **Unauthorized Controllers** pane. Wait for a few minutes for the latest controller to appear with the other controllers of the cluster, under the **Active Controllers** pane.

Also, check the **Current Size** and the **Target Size** in the **General** pane. The number displayed is updated with the latest node addition.

Contracting the APIC Cluster Using the GUI

This procedure reduces the cluster size. This procedure is applicable for releases prior to Cisco APIC release 6.0(2). For contracting a cluster in release 6.0(2), you can use the **Delete Node** option as detailed in the subsequent 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_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.

Contracting the APIC Cluster Using the Delete Node Option

Use this procedure to contract a cluster using the **Delete Node** option which has been introduced in Cisco APIC release 6.0(2). To contract a cluster in APIC releases prior to 6.0(2), see the previous procedure.

You can use this procedure to delete one or more than one node from an APIC cluster.

The **Delete Node** option includes two operations— reduces the cluster size, and decommissions the node.



Note A two-node cluster is not supported. You cannot delete one node from a three-node cluster. The minimum recommended cluster size is three.

Step 1 On the menu bar, choose **System > Controllers**. In the **Navigation** pane, expand **Controllers > apic_controller_name > Cluster as Seen by Node**.

Step 2 In the **Active Controllers** pane, select the controller which you want to delete by selecting the required check-box..

Step 3 Click the **Actions** button and select the **Delete Node** option.

Step 4 Click **OK** on the pop-up screen to confirm the deletion.

Selecting the force option has no effect. It is a *no operation* option, as it is not supported on Cisco APIC release 6.0(2).

Note You need to delete the nodes in decreasing order, that is, for example, you can not delete a node with ID 5 before deleting the node with ID 6.

Check the **Current Size** and **Target Size** in the General pane. The size indicated will be one lesser than it was before. If the earlier cluster size was N , now it will be $N-1$.

Note If you are deleting more than node from the cluster, the last node of the cluster is deleted first, followed by the other nodes. **Shrink In Progress** in the **General** pane is set to *Yes* until all the selected nodes are deleted.

What to do next

- After deleting an APIC from the cluster, power the controller down and disconnect it from the fabric.
- Wait for a few minutes, and confirm that the **Health State** of the remaining nodes of the cluster is displayed as *Fully fit* before further action.

Commissioning and Decommissioning Cisco APIC Controllers

Commissioning a Cisco APIC in the Cluster Using the GUI

Use this procedure for commissioning an APIC. This procedure is applicable for releases prior to Cisco APIC release 6.0(2). For release 6.0(2), the commissioning workflow has changed, please see the subsequent section for details.

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**.
-

Commissioning a Cisco APIC in the Cluster

Use this procedure on an existing Cisco Application Policy Infrastructure Controller (APIC) cluster for commissioning a Cisco APIC in that cluster. This procedure is applicable for Cisco APIC release 6.0(2). From release 6.0(2), the commissioning workflow has been enhanced. It can be used to provision an existing controller and also for RMA (return material authorization).

- 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** Select a decommissioned Cisco APIC from the **Active Controllers** table.
- Step 4** In the **Active Controllers** table, click the **Actions** icon (three dots), which is displayed at the end of the row for each Cisco APIC. From the displayed options, click **Commission**.
The **Commission** dialog box is displayed.
- Step 5** Enter the following details in the **Commission** screen:
Choose the **Controller Type**. Based on your choice, proceed to the relevant substep.
Put a check in the **Enabled** check box if you need to support IPv6 addresses.
- a) When the Controller Type is **Physical**:
- CIMC details pane
 - IP Address: Enter the CIMC IP address.
 - Username: Enter the username to access CIMC.
 - Password: Enter the password to access CIMC.
 - Click **Validate**. *Validation success* is displayed on successful authentication.
- This pane appears only if you configured CIMC. If you did not configure CIMC, instead perform the physical APIC login step of the [Bringing up the Cisco APIC Cluster Using the GUI](#) procedure (step 1b) on the new node to configure out-of-band management.
- General pane
 - Name: The name of the controller. The name is entered automatically after the CIMC validation.
 - Admin Password: Enter the admin password for the controller.
 - Controller ID: This is auto-populated based on the Cisco APIC that was decommissioned. The ID of the decommissioned node is assigned.
 - Serial Number: This is auto-populated after CIMC validation.

- Pod ID: Enter the ID number of the pod for the Cisco APIC.

- Out of Band Network pane

- IPv4 Address: Enter the IPv4 address of the out-of-band network.
- IPv4 Gateway: Enter the IPv4 gateway address of the out-of-band network.

Note If you have selected the **Enabled** check box for IPv6 earlier, enter the IPv6 address and gateway.

b) When the Controller Type is **Virtual** :

- Virtual Instance: Enter the management IP and click **Validate**.

Note The management IP addresses are defined during the deployment of the VMs using ESXi/ AWS.

- General pane

- Name: A user-defined name for the controller.
- Controller ID: This is auto-populated based on the Cisco APIC that was decommissioned. The ID of the decommissioned node is assigned.
- Serial Number: The serial number of the VM is auto-populated.

- Out of Band Network pane

- IPv4 Address: The IP address is auto-populated.
- IPv4 Gateway: The gateway IP address is auto-populated.

Note If you have selected the **Enabled** check box for IPv6 earlier, enter the IPv6 address and gateway.

- Infra Network pane

- IPv4 Address: Enter the infra network address.
- IPv4 Gateway: Enter the IP address of the gateway.
- VLAN: (applicable only for *remotely attached* virtual APIC- ESXi) enter the interface VLAN ID to be used.

Note The Infra L3 Network pane is not displayed when the virtual APIC is deployed using AWS.

Step 6 Click **Apply**.

Step 7 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. This procedure is applicable for APIC releases prior to Cisco APIC release 6.0(2). To decommission an APIC in release 6.0(2), see the subsequent procedure.



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.

Decommissioning a Cisco APIC in the Cluster

This procedure decommissions a Cisco APIC in the cluster. This procedure is applicable for Cisco APIC release 6.0(2). To decommission a Cisco APIC for releases prior to release 6.0(2), use the previous procedure.



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.

Step 3 In the **Work** pane, verify the **Health State** in the **Active Controllers** summary table indicates the cluster is **Fully Fit** before continuing.

Step 4 In the **Active Controllers** table, click the **Actions** icon (three dots) displayed at the end of the row for each APIC. Select the **Decommission** option.

The **Decommission** dialog box is displayed.

Step 5 Click **OK**.

The **Enabled** check-box for Force is a *no operation* option, as it is not supported on Cisco APIC release 6.0(2).

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.

Guidelines and Limitations for Standby Cisco APICs

The following are guidelines and limitations for standby Cisco Application Policy Infrastructure Controllers (APICs):

- There must be three active Cisco APICs to add a standby Cisco APIC.
- The standby Cisco APIC need to run with the same firmware version of the cluster when the standby Cisco APICs join the cluster during the initial setup.
- During an upgrade process, after all the active Cisco APICs are upgraded, the standby Cisco APICs are also upgraded automatically.
- During the initial setup, IDs are assigned to the standby Cisco APICs. After a standby Cisco APIC is switched over to an active Cisco APIC, the standby Cisco APIC (new active) starts using the ID of the replaced (old active) Cisco APIC.
- 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 needs to be 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, you can setup another standby Cisco APIC, if needed.
- If **Retain OOB IP address for Standby (new active)** is checked, the standby (new active) Cisco APIC will retain its original standby out-of-band management IP address.
- If **Retain OOB IP address for Standby (new active)** is not checked:
 - If only one active Cisco APIC is down: The standby (new active) Cisco APIC will use the old active Cisco APIC's out-of-band management IP address.
 - If more than one active Cisco APICs are down: The standby (new active) Cisco APIC will try to use the active Cisco APIC's out-of-band management IP address, but it may fail if the shard of out-of-band management IP address configuration for the active Cisco APIC is in the minority state.
- For Cisco ACI Multi-Pod, if the old active Cisco APIC and the standby Cisco APIC use different out-of-band management IP subnets, you must check the option to have the standby (new active) Cisco APIC retain its original standby out-of-band management IP address. Otherwise, you will lose out-of-band management IP connectivity to the standby (new active) Cisco APIC. This situation might happen if the old active Cisco APIC and the standby Cisco APIC are in the different pods.

If out-of-band management IP connectivity is lost because of this reason or if more than one active Cisco APICs are down, you must create a new Static Node Management OOB IP Address to change the new active (previously standby) Cisco APIC out-of-band management IP address. You must have the cluster out of the minority state to make the configuration change.
- 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.
- A standby Cisco APIC does not retain the in-band management IP address when you promote the Cisco APIC to be active. You must manually reconfigure the Cisco APIC to have the correct in-band management IP address.

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.
-

Migration of APICs

Beginning with Cisco APIC release 6.1(1), migration from a physical APIC cluster to a virtual APIC cluster deployed on an ESXi host (using VMware vCenter) is supported. Migration from a virtual APIC cluster (on an ESXi host) to a physical APIC cluster is also supported.

Guidelines and Limitations

Following are the guidelines and limitations for migrating physical APICs to virtual APICs (and vice-versa):

- Physical APICs in layer 2 (directly attached to the fabric) can be migrated to layer 2 virtual APICs and layer 2 virtual APICs can be migrated to layer 2 physical APICs. Physical APICs in layer 3 (remotely attached to the fabric) can be migrated to layer 3 virtual APICs. Virtual APICs in layer 3 can be migrated to layer 3 physical APIC. Migration from a layer 2 APIC to a layer 3 APIC (or vice-versa) is not supported.
- No support for standby node migration. Before migration, remove standby node(s) from the cluster and then migrate.
- No support for migration for mini ACI fabric.
- If upgrade is in progress, do not initiate the migration process.
- If migration is in progress, do not initiate an upgrade.
- If NDO is configured, you must update the connection details on NDO as the migration changes the OOB IP and subnet addresses.
- If SMU is installed on the physical APIC, for Cisco APIC release 6.1(1) the migration (from physical to virtual APIC) is not recommended. You need to upgrade the cluster to an image which has the fix for the SMU before proceeding with migration.
- For app-infra, stop any running job for ELAM/FTRIAGE before migration and re-start after migration is complete.
- Any configuration which uses APIC OOB needs to be updated after the migration process is completed.

How the migration process works

In this section, a high-level explanation of the migration process is provided. For the detailed steps, see the [Migrate a Physical APIC Cluster to a Virtual APIC Cluster \(or Virtual APIC Cluster to a Physical APIC Cluster\)](#) procedure in the subsequent section.

Consider a three-node cluster; three source nodes and correspondingly three target nodes (after migration nodes). The APIC with controller ID 1 is considered as APIC 1. Login to APIC 1 (IP address 172.16.1.1) and initiate the migration process.

APIC	Source Node	Target Node
APIC 1	172.16.1.1	172.16.1.11
APIC 2	172.16.1.2	172.16.1.12
APIC 3	17.16.1.3	172.16.1.13

As soon as the migration is initiated, APIC 1 starts with the migration of APIC 3, followed by APIC 2. After source APIC 2 (IP address 172.16.1.2) is migrated to the target APIC 2 node (172.16.1.12), the target APIC 2 node takes control to enable the migration of APIC 1. This is called the handover process where in the control is passed on from source APIC 1 (172.16.1.1) to target APIC 2 (172.16.1.12). At this stage, a new window is displayed (URL redirect to target APIC 2). This is because, after successful migration, source APIC 1 is no longer part of the cluster (which now has the migrated target APICs).

The migration proceeds in the reverse order, that is, APIC N (APIC 3 in the example) is migrated first, followed by APIC $N-1$ (APIC 2 in the example), so on, and then finally APIC 1.

Migrate a Physical APIC Cluster to a Virtual APIC Cluster (or Virtual APIC Cluster to a Physical APIC Cluster)

Use this procedure to migrate the nodes of a physical APIC cluster to a virtual APIC cluster (or vice-versa).

Before you begin

Following are the required prerequisites before you start with the migration process:

Cluster health

Confirm that the current APIC cluster is *Fully fit*.

Generic

- Ensure that the source and destination APICs' date and time are synchronized.
- Ensure that all the controllers are on Cisco APIC release 6.1(1), and all the switches are running the same version as the controller.

Source and Target Nodes

- For directly connected APIC migration, ensure both source and target nodes are on the same layer 2 network.
- For remotely connected APIC migration, ensure both source and target nodes have infra network connectivity between them. This means the new target APIC should have the correct IPN configuration such that it can interact with the infra network of the fabric.
- Target nodes have the same admin password as the source cluster.
- Target nodes' OOB IP address should be different while all other fields can be same or different from the source node. Infra addresses will remain same for layer 2 (directly attached); for layer 3 (remotely attached) cluster, they can be same or different based on deployments.
- Source cluster and target cluster OOB networking stacks should match. For example, if source cluster is using dual stack (IPv4 and IPv6) for OOB, dual stack (IPv4 and IPv6) address details should be provided for target nodes too.
- Ensure OOB connectivity between the source and destination APICs.
- Ensure the OOB contracts and reachability for the new APIC are configured correctly; the migration process uses the OOB IP address to communicate between the APICs.

For virtual APIC to physical APIC migration

- Ensure the physical APIC nodes are factory reset; use the **acdiag touch setup** and **acdiag reboot** commands.
- For migration with/ without CIMC (applicable for physical APICs):

If...	Then....
Using CIMC	ensure that the physical APIC CIMC addresses are reachable from the OOB network of the virtual APIC.
Not using CIMC	ensure that the OOB IP address is configured manually on the physical APIC after factory-reset and use the OOB option for connectivity.

For Physical to virtual migration

- Ensure that you have deployed the virtual APIC nodes as per the procedure in the [Deploying Cisco Virtual APIC Using VMware vCenter](#) guide.
- If virtual APICs are deployed on a vCenter that is part of a VMM domain, ensure that Infrastructure VLAN is enabled on the AEP configured on the interfaces connected to the ESXi host(s) where the virtual APIC is deployed.

Step 1 On the **Cluster as Seen by Node** screen, click **Migrate** (displayed in the **Cluster Overview** area).

All the available controllers in the cluster are displayed.

Note The **Migrate** button is displayed only on APIC 1 (of the cluster).

Step 2 Click the pencil icon next to the Validate column, to start the migration process of the selected controller.

The **Migrate Node** screen is displayed.

Step 3 Enter the following details in the **Migrate Node** screen:

- a) For the **Controller Type**, select Virtual or Physical, as the case may be (migration from physical APIC to virtual APIC and vice-versa is supported).
- b) For the **Connectivity Type**, select OOB, if you are migrating a physical APIC to a virtual APIC. If you are migrating a virtual APIC to a physical APIC, you can either select the OOB option or the CIMC option.

It is recommended to select the CIMC option for the virtual to physical migration. To use the OOB option, connect to the CIMC address of physical APICs and configure the OOB IP addresses manually before starting the migration process.

Controller Type and **Connectivity Type** are auto selected based on the source controller type. If required, you can modify them.

- c) In the Management IP pane, enter the following (target APIC details) — Management IP address, Username, Password.
or

(applicable only for virtual to physical APIC migration) In the CIMC Details pane, enter the following details of the physical APIC— CIMC IP address, username and password of the node.

- d) Click **Validate**.

After you click Validate, the details displayed in the General and Out of Band management panes change to match the details of the controller. The only editable fields are the Name and Pod ID (applicable only for layer 2), the other fields cannot be modified. For virtual to physical APIC migration, confirm the Admin Password too.

Note If dual stack is supported, fill in the IPv4 and IPv6 addresses.

- e) In the Infra Network pane (applicable only for Layer 3, APIC is remotely attached to the fabric), enter the following:
 - IPv4 Address: the infra network address.
 - IPv4 Gateway: the IP address of the gateway.
 - VLAN: the interface VLAN ID to be used.

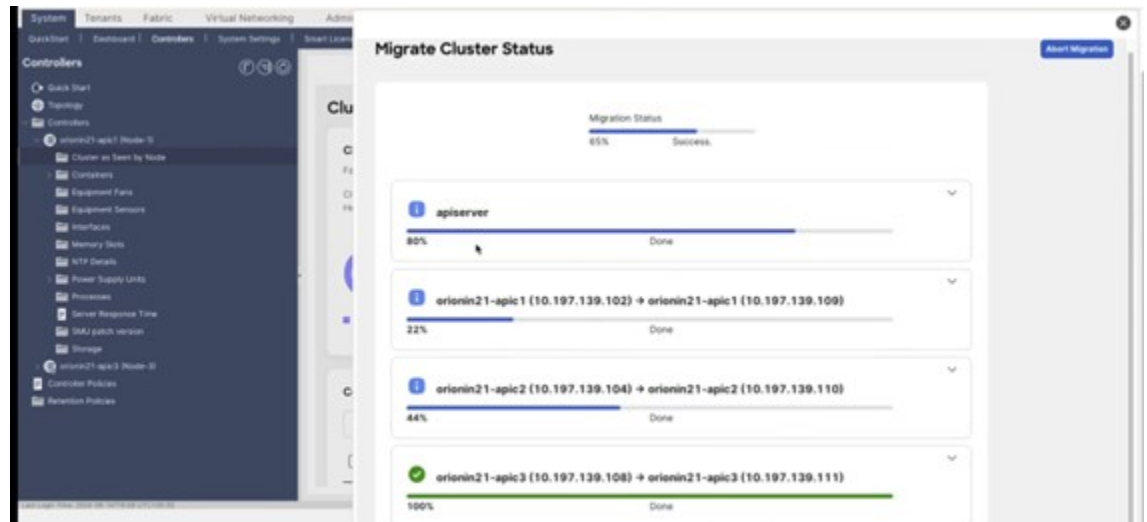
The OOB gateway and IP addresses are auto-populated in the table (based on the validation); click **Apply**. The validation status is displayed as Complete (on the Migrate Nodes screen).

Repeat the same process for the other APICs in the cluster by clicking the pencil icon (next to the Validation column). After providing all the controller details, click the **Migrate** button at the bottom of the Migrate screen.

Checking the Migration Status

Click **Check Migration Status** to check the status of the migration. The **Check Type** section has details of the completed steps. The migration process involves a series of activities, and this is displayed in stages and each stage is indicated with a color-coded bar.

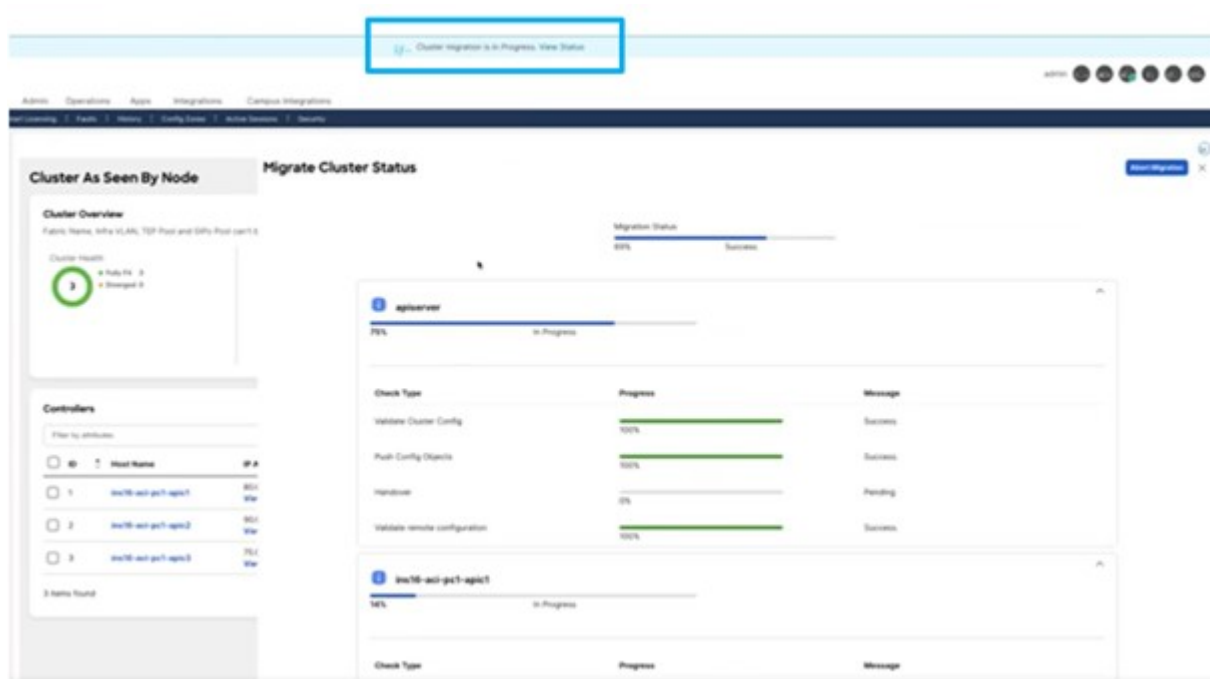
Figure 1: Migration Status



As displayed in the image, the overall migration status is displayed, followed by the status of the `apiserver`. The `apiserver` is the process that orchestrates the whole migration process. Below the `apiserver` status, the controller-wise migration status is displayed. The source IP address and the target IP address of the nodes are also indicated.

The `apiserver` status is indicated as 100% done (green bar) after the handover to APIC 2 is completed. At this stage, a new window is displayed (URL redirected to target APIC 2). Login to target APIC 2. A banner indicating that *Migration is in progress* is displayed at the top of the GUI until the migration is complete. After the handover process, the banner which was displayed on source APIC 1 is displayed on the target APIC 2. Click the **View Status** link on the banner to check the migration status.

Figure 2: Migration Status Banner



You can also abort the migration process from source APIC 1 by clicking the **Abort** button available on the Migrate Cluster Status screen. The **Abort** button is displayed only after a certain period of time after initiating the migration.

After successful migration:

- the migration status is no longer displayed. If the migration has failed, then a failure message is explicitly displayed.
- to confirm if the target cluster is healthy and fully fit, navigate to **System > Controllers > Expand Controllers**. Expand **Controller 1 > Cluster as seen by Node** page.
- verify if all the fabric nodes are in active state; navigate to **Fabric > Fabric Membership**.
- if the Pod ID of the target APIC has changed, inband address for the node needs to be reconfigured on the Tenant Management screen. Navigate to **Tenants > Mgmt > Node Management Addresses** page.

Operations in case of Migration Failure

A migration process may be interrupted due to a failure, or you may choose to abort the migration. In case the migration is not successful, it is recommended to revert or resume migration to either the source or target controller type. It is not recommended to have a APIC cluster in a migration failed state with a mix of physical and virtual controllers. Before attempting revert or resume, follow the steps in the next section, [Basic Troubleshooting](#), to get the cluster to a healthy state.

If you choose to resume— migration process is continued. On the **Migrate Node** screen (source APIC 1):

1. Enter the details of all the target nodes based on the controller type you want to migrate to.
2. Click **Migrate**.

If you choose to revert— migration process is restarted. Migration process needs to be restarted after getting the controllers of the cluster to the initial (source) IP addresses.

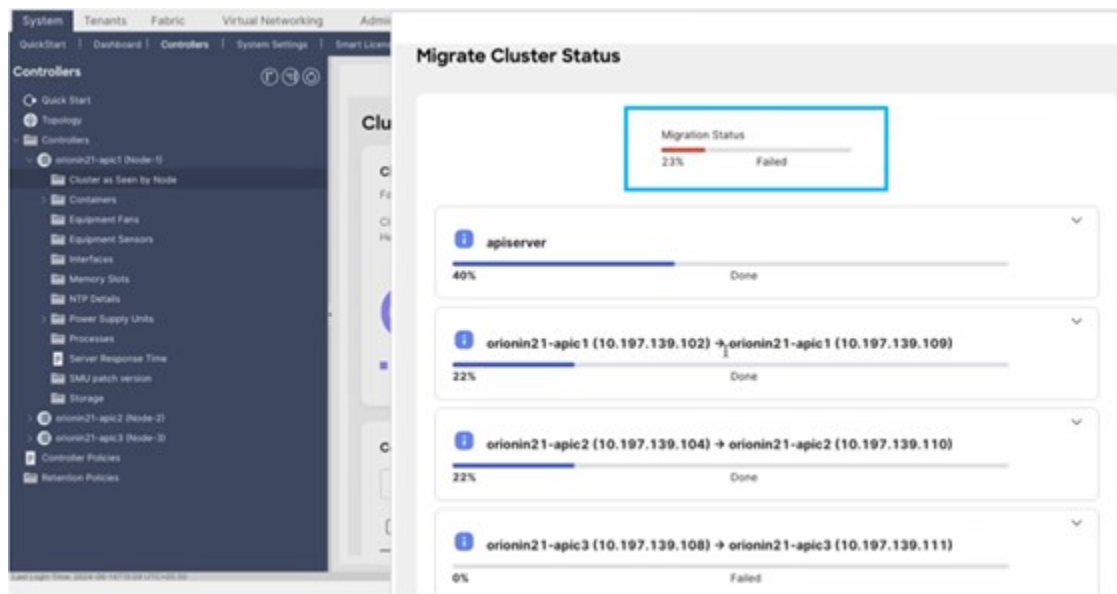
1. Factory reset each of the source APIC nodes which are being migrated using **acidiag touch setup** and **acidiag reboot** commands.
2. On the **Migrate Node** screen, enter the source APIC details for all the nodes, as the migration process reverts the previously migrated APICs to the source controller type.
3. Click **Migrate**.



Note If the migration process fails after the handover process (control is passed on to target APIC 2 from source APIC 1), the migration cannot be resumed or reverted.

As mentioned earlier, the various sub-stages of migration with their completion progress are indicated as bars, for each controller. In case of failure at any stage, collect the relevant tech-support details and contact Cisco TAC for further assistance. To collect the logs for tech support, navigate to **Admin > Import/Export > Export Policies > On-demand Tech Support > migration_techsupport**.

Figure 3: Migration Status - Failed



Basic Troubleshooting

Consider a three-node cluster where in two nodes have migrated successfully, and a failure is detected during the migration of the third node. Check the status of the failed node. If the controller is not in the *Fully fit* state, the migration could fail.

The procedure to get the cluster to a healthy state:

-
- Step 1** (for migration failures with APIC 1) Check the cluster health from target APIC 2 by navigating to, **System > Controllers**. Select **Controller 2 > Cluster as seen by Node** screen.
- or
- (for migration failures with APIC 2 to *N*) Check the cluster health from source APIC 1 by navigating to, **System > Controllers**. Select **Controller 1 > Cluster as seen by Node** screen.
- Step 2** If APIC 1 (or any other node of the cluster) is not *Fully fit*, click the three dots adjacent to the serial number of the controller. Select **Maintenance > Decommission**. Click **Force Decommission** as the node is not in a *Fully fit* state. Connect to the source APIC node *N* using SSH and factory-reset the node using the following commands - **acdiag touch setup, acdiag reboot**.
- Step 3** From source APIC 1 navigate to, **System > Controllers**. Click **Controllers > Controller 1 > Cluster as seen by Node** screen.
- or
- From target APIC 2 navigate to, **System > Controllers**. Click **Controllers > Controller 2 > Cluster as seen by Node** screen.
- Step 4** To commission a controller, click the three dots adjacent to the serial number of the controller. Select **Maintenance > Commission**. Enter the details, as required. Refer to the *Commissioning a Node* procedure (described earlier in this chapter). The only difference here is that the controller ID is pre-populated with the number corresponding to the ID of the controller in the cluster.
- After the controller is commissioned, the cluster is indicated as *Fully fit*.
- Step 5** Check the status of the cluster after commissioning the failed node. If the cluster is in a healthy state, resume the migration by clicking **Migrate** on the **Cluster As Seen By Node** screen. If the migration fails again, contact Cisco TAC for further assistance.
-