



Troubleshooting the Cisco APIC Cluster

This chapter contains information about cluster faults and possible solutions for common scenarios. For information about managing Cisco APIC clusters, see the appendix in this document.

This chapter contains the following sections:

- [Cluster Troubleshooting Scenarios, on page 1](#)
- [Cluster Faults, on page 4](#)

Cluster Troubleshooting Scenarios

The following table summarizes common cluster troubleshooting scenarios for the Cisco APIC.

Problem	Solution
An APIC node fails within the cluster. For example, node 2 of a cluster of 5 APICs fails.	<p>There are two available solutions:</p> <ul style="list-style-type: none">• Leave the target size and replace the APIC.• Reduce the cluster size to 4, decommission controller 5, and recommission it as APIC 2. The target size remains 4, and the operational size is 4 when the reconfigured APIC becomes active. <p>Note You can add a replacement APIC to the cluster and expand the target and operational size. For instructions on how to add a new APIC, refer to the <i>Cisco APIC Management, Installation, Upgrade, and Downgrade Guide</i>.</p>

Problem	Solution
<p>A new APIC connects to the fabric and loses connection to a leaf switch.</p>	<p>Use the following commands to check for an infra (infrastructure) VLAN mismatch:</p> <ul style="list-style-type: none"> • <code>cat /mit/sys/lldp/inst/if-[eth1--1]/ctrlradj/summary</code>—Displays the VLAN configured on the leaf switch. • <code>cat /mit/sys/lldp/inst/if-[eth1--1]/ctrlradj/summary</code>—Displays the infra (infrastructure) VLANs advertised by connected APICs. <p>If the output of these commands shows different VLANs, the new APIC is not configured with the correct infra (infrastructure) VLAN. To correct this issue, follow these steps:</p> <ul style="list-style-type: none"> • Log in to the APIC using <code>rescue-user</code>. <p>Note Admin credentials do not work because the APIC is not part of the fabric.</p> <ul style="list-style-type: none"> • Erase the configuration and reboot the APIC using the acidiag touch setup command. • Reconfigure the APIC. Verify that the fabric name, TEP addresses, and infra (infrastructure) VLAN match the APICs in the cluster. • Reload the leaf node.
<p>Two APICs cannot communicate after a reboot.</p>	<p>The issue can occur after the following sequence of events:</p> <ul style="list-style-type: none"> • APIC1 and APIC2 discover each other. • APIC1 reboots and becomes active with a new ChassisID (APIC1a) • The two APICs no longer communicate. <p>In this scenario, APIC1a discovers APIC2, but APIC2 is unavailable because it is in a cluster with APIC1, which appears to be offline. As a result, APIC1a does not accept messages from APIC2.</p> <p>To resolve the issue, decommission APIC1 on APIC2, and commission APIC1 again.</p>
<p>A decommissioned APIC joins a cluster.</p>	<p>The issue can occur after the following sequence of events:</p> <ul style="list-style-type: none"> • A member of the cluster becomes unavailable or the cluster splits. • An APIC is decommissioned. • After the cluster recovers, the decommissioned APIC is automatically commissioned. <p>To resolve the issue, decommission the APIC after the cluster recovers.</p>

Problem	Solution
Mismatched ChassisID following reboot.	<p>The issue occurs when an APIC boots with a ChassisID different from the ChassisID registered in the cluster. As a result, messages from this APIC are discarded.</p> <p>To resolve the issue, ensure that you decommission the APIC before rebooting.</p>
The APIC displays faults during changes to cluster size.	<p>A variety of conditions can prevent a cluster from extending the OperationalClusterSize to meet the AdministrativeClusterSize. For more information, inspect the fault and review the "Cluster Faults" section in the <i>Cisco APIC Basic Configuration Guide</i>.</p>
An APIC is unable to join a cluster.	<p>The issue occurs when two APICs are configured with the same ClusterID when a cluster expands. As a result, one of the two APICs cannot join the cluster and displays an expansion-contender-chassis-id-mismatch fault.</p> <p>To resolve the issue, configure the APIC outside the cluster with a new cluster ID.</p>
APIC unreachable in cluster.	<p>Check the following settings to diagnose the issue:</p> <ul style="list-style-type: none"> • Verify that fabric discovery is complete. • Identify the switch that is missing from the fabric. • Check whether the switch has requested and received an IP address from an APIC. • Verify that the switch has loaded a software image. • Verify how long the switch has been active. • Verify that all processes are running on the switch. For more information, see the "acidiag Command" section in the <i>Cisco APIC Basic Configuration Guide</i>. • Confirm that the missing switch has the correct date and time. • Confirm that the switch can communicate with other APICs.

Problem	Solution
Cluster does not expand.	<p>The issue occurs under the following circumstances:</p> <ul style="list-style-type: none"> • The OperationalClusterSize is smaller than the number of APICs. • No expansion contender (for example, the admin size is 5 and there is not an APIC with a clusterID of 4). • There is no connectivity between the cluster and a new APIC • Heartbeat messages are rejected by the new APIC • System is not healthy. • An unavailable appliance is carrying a data subset that is related to relocation. • Service is down on an appliance with a data subset that is related to relocation. • Unhealthy data subset related to relocation.
An APIC is down.	<p>Check the following:</p> <ul style="list-style-type: none"> • Connectivity issue—Verify connectivity using ping. • Interface type mismatch—Confirm that all APICs are set to in-band communication. • Fabric connectivity—Confirm that fabric connectivity is normal and that fabric discovery is complete. • Heartbeat rejected—Check the fltInfraIICIMsgSrcOutsider fault. Common errors include operational cluster size, mismatched ChassisID, source ID outside of the operational cluster size, source not commissioned, and fabric domain mismatch.

Cluster Faults

The APIC supports a variety of faults to help diagnose cluster problems. The following sections describe the two major cluster fault types.

Discard Faults

The APIC discards cluster messages that are not from a current cluster peer or cluster expansion candidate. If the APIC discards a message, it raises a fault that contains the originating APIC's serial number, cluster ID, and a timestamp. The following table summarizes the faults for discarded messages:

Fault	Meaning
expansion-contender-chassis-id-mismatch	The ChassisID of the transmitting APIC does not match the ChassisID learned by the cluster for expansion.
expansion-contender-fabric-domain-mismatch	The FabricID of the transmitting APIC does not match the FabricID learned by the cluster for expansion.

Fault	Meaning
expansion-contender-id-is-not-next-to-oper-cluster-size	The transmitting APIC has an inappropriate cluster ID for expansion. The value should be one greater than the current OperationalClusterSize.
expansion-contender-message-is-not-heartbeat	The transmitting APIC does not transmit continuous heartbeat messages.
fabric-domain-mismatch	The FabricID of the transmitting APIC does not match the FabricID of the cluster.
operational-cluster-size-distance-cannot-be-bridged	The transmitting APIC has an OperationalClusterSize that is different from that of the receiving APIC by more than 1. The receiving APIC rejects the request.
source-chassis-id-mismatch	The ChassisID of the transmitting APIC does not match the ChassisID registered with the cluster.
source-cluster-id-illegal	The transmitting APIC has a clusterID value that is not permitted.
source-has-mismatched-target-chassis-id	The target ChassisID of the transmitting APIC does not match the Chassis ID of the receiving APIC.
source-id-is-outside-operational-cluster-size	The transmitting APIC has a cluster ID that is outside of the OperationalClusterSize for the cluster.
source-is-not-commissioned	The transmitting APIC has a cluster ID that is currently decommissioned in the cluster.

Cluster Change Faults

The following faults apply when there is an error during a change to the APIC cluster size.

Fault	Meaning
cluster-is-stuck-at-size-2	This fault is issued if the OperationalClusterSize remains at 2 for an extended period. To resolve the issue, restore the cluster target size.
most-right-appliance-remains-commissioned	The last APIC within a cluster is still in service, which prevents the cluster from shrinking.
no-expansion-contender	The cluster cannot detect an APIC with a higher cluster ID, preventing the cluster from expanding.
service-down-on-appliance-carrying-replica-related-to-relocation	The data subset to be relocated has a copy on a service that is experiencing a failure. Indicates that there are multiple such failures on the APIC.
unavailable-appliance-carrying-replica-related-to-relocation	The data subset to be relocated has a copy on an unavailable APIC. To resolve the fault, restore the unavailable APIC.
unhealthy-replica-related-to-relocation	The data subset to be relocated has a copy on an APIC that is not healthy. To resolve the fault, determine the root cause of the failure.

APIC Unavailable

The following cluster faults can apply when an APIC is unavailable:

Fault	Meaning
fltInfraReplicaReplicaState	The cluster is unable to bring up a data subset.
fltInfraReplicaDatabaseState	Indicates a corruption in the data store service.
fltInfraServiceHealth	Indicates that a data subset is not fully functional.
fltInfraWiNodeHealth	Indicates that an APIC is not fully functional.