

Cisco APIC M1/M2/M3/L1/L2/L3 to M4/L4 Cluster Migration, Release 5.3(1)

Version 1.0



Contents

- Goals of This Document..... 3
- Software Release Requirements 3
- Hardware Compatibility..... 3
- Guidelines and Limitations for Migrating Cisco APIC Servers 4
- Replacing the In-service APIC Servers..... 4
- Decommissioning the Standby Cisco APIC Servers to be Replaced by a Normal Cluster 11
- Troubleshooting the New Cluster 12

Goals of This Document

This document provides details on how to perform an in-service replacement of older generation Cisco APIC servers with the M4/L4 model. As announced on [cisco.com](https://www.cisco.com)¹ both APIC L1/M1 and APIC L2/M2 servers have reached their end-of-sales and end-of-life date. At the time of this writing, the suggested Cisco APIC server replacement is APIC M4/L4.

Note: This document is for the Cisco APIC 5.3 releases. For cluster migration information for the 6.0(2) and later releases, see [Cisco APIC M1/M2/M3/L1/L2/L3 to M4/L4 Cluster Migration, Release 6.0\(2\)](#).

Software Release Requirements

The APIC M4/L4 requires the Cisco APIC software 5.3(1) release or later or the 6.0(2) release or later. This document uses the Cisco APIC 5.3(1d) release as an example. Cisco APIC servers forming a cluster must all run the same software release. You cannot have different software releases inside one cluster; doing so will result in the cluster not converging. There is one exception to this rule: during a software upgrade process, there will be a temporary divergence in software releases within the cluster. This means that before you attempt to replace the existing Cisco APIC M1/L1, M2/L2, or M3/L3 server with a Cisco APIC M4/L4 server, you must bring the running cluster to a supported release.

To determine which release version you are currently running on the Cisco APIC M4/L4 server:

- Step 1.** Power on your Cisco APIC M4/L4 and determine which release you are currently running. If the APIC is already running release 5.3(1), skip to Step 3.
- Step 2.** If the Cisco APIC M4/L4 is not running release 5.3(1), install the 5.3(1) release. For the procedure, see [Installing Cisco APIC Software Using CIMC Virtual Media](#) in the *Cisco APIC Installation and ACI Upgrade and Downgrade Guide*. Follow the procedure up through step 8.
- Step 3.** Upgrade or downgrade (as appropriate) each Cisco APIC in the cluster to the same release before continuing any further.

Hardware Compatibility

You can mix Cisco APIC servers using any possible combination. There are no restrictions other than the minimum software release mentioned in the Software Release Requirements.

Table 1. Table Caption

	APIC-M1/L1	APIC-M2/L2	APIC-M3/L3	APIC-M4/L4
APIC-M1/L1	X	X	X	X
APIC-M2/L2	X	X	X	x

¹ <https://www.cisco.com/c/en/us/products/cloud-systems-management/application-policy-infrastructure-controller-apic/eos-eol-notice-listing.html>

	APIC-M1/L1	APIC-M2/L2	APIC-M3/L3	APIC-M4/L4
APIC-M3/L3	X	X	X	x
APIC-M4/L4	X	X	X	x

When a cluster mixes different hardware models, its performance aligns to the lowest common denominator. For example, an APIC-M2 cluster scales up to 1000 edge ports while an APIC-M3 cluster increases that number to 1200².

Guidelines and Limitations for Migrating Cisco APIC Servers

- The Cisco APIC L1/M1 server is no longer supported. However, you can still use the procedures in this document to migrate Cisco APIC L1/M1 servers to a newer server model.
- When you decommission a Cisco APIC, the 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.
- Do not decommission more than one Cisco APIC at a time.
- Wait until the cluster reaches the fully fit state before proceeding with a new replacement.
- Do not leave a decommissioned Cisco APIC powered on.

Replacing the In-service APIC Servers

This section describes how to replace every server with an M4/L4 server model in service with no impact to the data plane nor the control plane. The procedure is fully supported by Cisco. This procedure focuses on a 3 node Cisco APIC cluster and the process is similar for larger clusters.

Step 1. Validate the existing cluster is fully-fit.

Ensure your existing cluster is fully fit before attempting this procedure. You must not upgrade or modify a Cisco APIC cluster that is not fully fit. To verify your existing cluster is fully fit:

² <https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/application-policy-infrastructure-controller-apic/datasheet-c78-739715.html>

- a. In the menu bar, choose **System > Controllers**.
- b. In the Navigation pane, expand **Controllers** and choose any Cisco APIC.
- c. Expand the Cisco APIC and choose **Cluster as seen by node**.



Figure 1: Cluster as Seen by Node

- d. Check the operational state of all nodes. The nodes must be "Available" and the health state must be "Fully Fit."

Step 2. Record the name and infra VLAN of your existing fabric.

You can obtain fabric name from the Cluster as Seen by Node screen as shown in step 1c, Figure 1.

- a. If you do not know the infra VLAN and fabric ID of the Cisco APIC, use the Cisco APIC GUI to get it. In the menu bar, go to **System > Controllers**. In the Navigation pane, go to **Controllers > apic_name**. In the Work pane, go to **General > Controllers** and find the **Infra VLAN** property.



Figure 2: Infra VLAN and Fabric ID of the Cisco APIC

- b. Obtain the TEP pool that you used when you first brought up your fabric. In the menu bar, go to **Fabric > Inventory**. In the Navigation pane, go to **Pod Fabric Setup Policy**. In the Work pane, see the **TEP Pool** column.
- c. Obtain the group IP outer (GIPo) pool address (multicast pool address) that you used when you first brought up your fabric. In the menu bar, go to **System > Controllers**. In the Navigation pane, go to **Controllers > apic_name**. In the Work pane, go to **General > IP Settings** and see **Multicast Pool Address**.

- d. Obtain the pod ID using the CLI:

```
apic1# moquery -d "topology/pod-1/node-1/av/node-3" | grep -e podId
podId          : 1
```

- e. Obtain the out-of-band management IP address. In the menu bar, go to **System > Controllers**. In the Navigation pane, go to **Controllers > apic_name**. In the Work pane, go to **General > IP Settings** and see **Out-of-Band Management**.

Step 3. Only for a standalone APIC (APIC over a Layer 3 network), obtain the following information:

- VLAN ID for interface
- Cisco APIC IPv4 address
- IPv4 address of the Cisco APIC default gateway
- IPv4 address of an active Cisco APIC

For an active Cisco APIC, use the APIC GUI to obtain the IP address of an APIC that you are not planning to decommission:

- In the menu bar, choose **System > Controllers**.
- In the Navigation pane, expand **Controllers** and choose any Cisco APIC.
- Expand the Cisco APIC and choose **Cluster as seen by node**.
- In the Work pane, get the IP address from the **IP** column.

Step 4. Decommission the last Cisco APIC.

From Cisco APIC number 1 or 2, within the 'cluster as seen by node' view (Figure 1), decommission the last Cisco APIC by right-clicking that APIC and choosing 'Decommission' as shown in Figure 3.

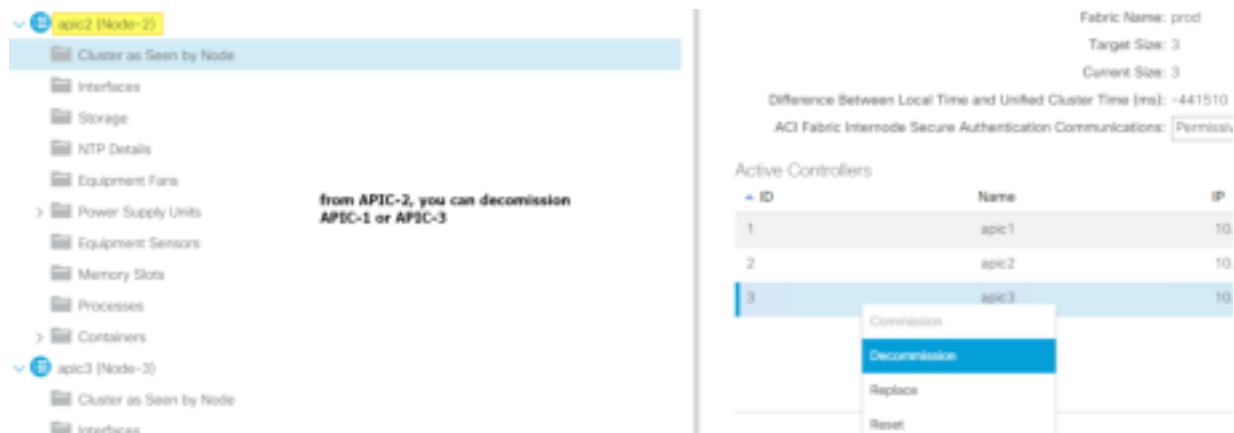


Figure 3: Decommissioning the Last APIC

Wait roughly 5 minutes, then log into that Cisco APIC's CIMC or attach a physical keyboard and monitor to its back so you can initiate a power off sequence after having decommissioned the Cisco APIC server. You will see the admin status change from "In Service" to "Out of Service" and the operational status change to "Unregistered":

Active Controllers						
ID	Name	IP	Admin State	Operational State	Health State	
1	apic1	10.8.8.1	In Service	Online	Full	
2	apic2	10.8.8.2	In Service	Online	Full	
3	apic3	10.8.8.3	Out of Service	Unregistered	Unknown	

Figure 4: The APIC Becomes Out of Service and Unregistered

When the old Cisco APIC is out of service, power it off:

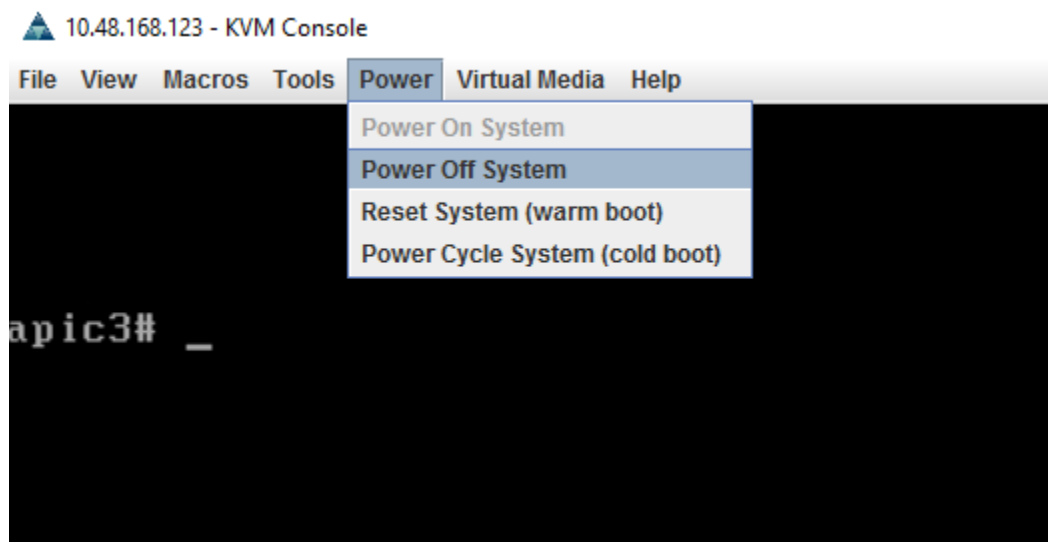


Figure 5: Powering Off the APIC

Step 5. Cable the replacement Cisco APIC M4/L4 servers.

Physically install the replacement the servers in the data center and cable them up to the existing Cisco ACI fabric as you would with any server. If necessary, ensure LLDP is disabled at the CIMC NIC level. Cable the out-of-band (OOB) management connection. There is no need to set aside new IP addresses for the replacement Cisco APIC servers, because each Cisco APIC will simply take over the IP of the server it is replacing.

Step 6. Power up the replacement Cisco APIC M4/L4 servers.

Power up all Cisco APIC M4/L4 servers and bring up a virtual keyboard, video, mouse session, Serial over LAN (SoL), or physical VGA connection so you can monitor their boot process. After a few minutes, you will be prompted to press any key to continue. Do not press a key just yet. Leave the Cisco APIC M4/L4 servers in that stage for the time being. See Figure 6:




Figure 6: APIC M4/L4 Boot Sequence

Step 7. Bring in the replacement APIC.

For a Layer 2 mode Cisco APIC (an APIC that is directly connected to a leaf switch), pick one of the new Cisco APIC M4/L4 servers that is waiting at the "press any key to continue" prompt and press a key. You will be prompted to configure this Cisco APIC. Enter the details that you recorded into the new Cisco APIC as shown below:


```

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Note: setup is mainly used for configuring the system initially,
when no configuration is present. So setup always assumes system
defaults and not the current system configuration values.

Press Enter at anytime to assume the default values. Use ctrl-d
at anytime to restart from the beginning.

Cluster configuration ...
Enter the fabric name [ACI Fabric1]: prod
Enter the fabric ID (1-128) [1]:
Enter the number of active controllers in the fabric (1-9) [3]:
Is this a standby controller? [no]:
Enter the controller ID (1-3) [1]: 3
Standalone APIC Cluster ? yes/no [no]:
Enter the POD ID (1-254) [1]: 1
Enter the controller name [apic3]:
Enter address pool for TEIP addresses [10.0.0.0/16]: 10.0.0.0/16
Note: The infra VLAN ID should not be used elsewhere in your environment
and should not overlap with any other reserved VLANs on other platforms.
Enter the VLAN ID for infra network (1-4094) [1]: 4093

Out-of-band management configuration ...
Enable IPv6 for Out of Band Mgmt Interface? [N]:
Enter the IPv4 address [192.168.10.1/24]: 10.23.193.38/24
Enter the IPv4 address of the default gateway [None]: 10.23.193.1

```

Figure 7: Entering the Cisco APIC Information

Only for a standalone APIC (APIC over a Layer 3 network), you also need to enter the following data:

```

Standalone APIC Cluster ? yes/no [no]: yes
Enter the VLAN ID for interface (0-access) (0-4094) [0]: 0
Enter the APIC IPV4 address [A.B.C.D/NN]: 15.152.2.1/30
Enter the IPv4 address of the APIC default gateway [A.B.C.D]: 15.152.2.2
Enter the IPv4 address of an active APIC [A.B.C.D]: 15.150.2.1

```

After you have entered all parameters, you will be asked whether you want to modify them. Enter 'N' unless you made a mistake that you want to correct.

Step 8. Register the new Cisco APIC for the cluster membership.

After roughly 7 to 10 minutes, the new server appears as unregistered in the 'cluster as seen by node' tab in the GUI, as shown below. Right-click on the server and commission it. Wait till the health state goes to fully fit for the new and all servers before continuing any further. This usually takes 5 minutes.

ID	Name	IP	Admin State	Operational State
1	apic1	10.0.0.1	In Service	Available
2	apic2	10.0.0.2	In Service	Available
3	apic3	10.0.0.3	Out of Service	Unregistered

Figure 8: Commissioning the New Cisco APIC

In the case of strict mode, you must approve the controller.

Step 9. Validate the cluster membership.

After 5 minutes or so, you will observe transitions in operational state and health status. The new server first has a data layer partially diverged state before fully converging:

ID	Name	IP	Admin State	Operational State	Health Status	Database Status	Serial Number
1	apic1	10.0.0.1	In Service	Available	Full Fit	OK	XXXXXXXXXX
2	apic2	10.0.0.2	In Service	Available	Full Fit	OK	XXXXXXXXXX
3	apic3	10.0.0.3	In Service	Available	Data Layer Partially Diverged	OK	XXXXXXXXXX

Figure 9: Changes in the New Cisco APIC's Operational State and Health Status

Shortly after, the new server's database is fully in sync with the other members of the cluster. This is reflected in a fully fit health state.

ID	Name	IP	Admin State	Operational State	Health Status	Database Status	Serial Number
1	apic1	10.0.0.1	In Service	Available	Full Fit	OK	XXXXXXXXXX
2	apic2	10.0.0.2	In Service	Available	Full Fit	OK	XXXXXXXXXX
3	apic3	10.0.0.3	In Service	Available	Full Fit	OK	XXXXXXXXXX

Figure 10: The New Cisco APIC Becomes Fully Fit

If you zoom in on the new server's properties, you will see it is indeed an M4/L4 with a new serial number:

Hardware Usage	
Allocated Memory (KB):	24588768
Free Memory (KB):	72759008
APIC Model:	APIC-SERVER-M4
APIC Serial Number:	W2P21100U4P
CPU Architecture:	x86_64
Cores:	8
CPU Model:	Intel(R) Xeon(R) Bronze 3106 CPU @ 1.70GHz
Speed (GHz):	0.8
Vendor:	GenuineIntel
Locator LED:	Off
Locator LED Color:	No color

Figure 11: New Cisco APIC's Model and Serial Number

Step 10. Decommission the next Cisco APIC server.

To decommission the next server, repeat steps 4 through 9. Remember that to decommission a controller, you need to perform the operation from another server's standpoint. If you are logged into APIC-1 for instance, do not decommission APIC-1. Log into APIC-2, go to the "cluster as seen by node" view for APIC-2 and decommission APIC-1. This is shown below:

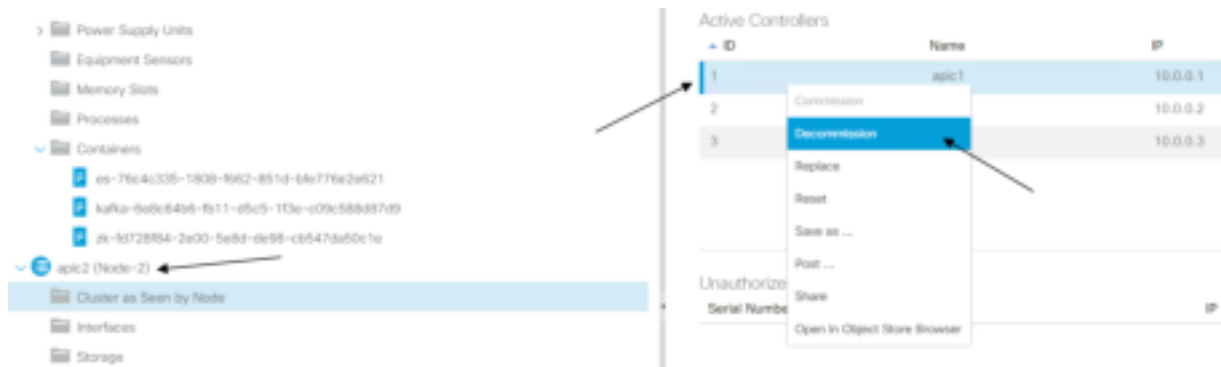


Figure 12: Decommissioning the Next Cisco APIC

Do not forget to power off the server that has been decommissioned before attempting to bring in a replacement.

Step 11. Verify the entire cluster.

After you have decommissioned and powered off the server, boot up, configure, and commission the M4, routing as many times as necessary. Validate that the entire cluster is fully fit:

ID	Name	IP	Admin State	Operational State	Health State	Redundancy State	Serial Number
1	apic1	10.0.0.1	In Standby	OK	Up/Up	OK	10000000000000000000
2	apic2	10.0.0.2	In Standby	OK	Up/Up	OK	10000000000000000000
3	apic3	10.0.0.3	In Standby	OK	Up/Up	OK	10000000000000000000

Figure 13: Validating That the Cluster is Fully Fit

The replacement for APIC-1 is also an M4 model:

```

Hardware Usage
  Allocated Memory (KB): 31056080
  Free Memory (KB): 68291780
  APIC Model: APIC-SERVER-M4
  APIC Serial Number: WZP23100T5F
  CPU Architecture: x86_64
  Cores: 8
  CPU Model: Intel(R) Xeon(R) Bronze 3106 CPU @ 1.70GHz
  Speed (GHz): 1.44
  Vendor: GenuineIntel
  Locator LED: Off
  Locator LED Color: No color
  
```

Figure 14: Verifying the Cisco APIC's Model

At this point, you have a fully operational fully-fit Cisco APIC cluster with new hardware.

Decommissioning the Standby Cisco APIC Servers to be Replaced by a Normal Cluster

If your cluster contains obsolete standby Cisco APIC servers, the same process applies. When you bring your existing cluster to a supported release, the standby Cisco APIC servers are automatically upgraded.

To decommission the standby Cisco APIC servers:

Step 1. Ensure the new M4 or L4 model is running the same software release as the rest of the cluster members.

Step 2. Decommission the standby Cisco APIC to be replaced to a normal cluster member. Power down the APIC and issue the following command for the controller to become unregistered:

```
acidiag cluster erase standby_node_id standby_serial_number
```

Step 3. Bring in the new M4 or L4 server and specify that server is a standby Cisco APIC during the setup. When you are prompted with "Is this a standby controller? [NO]", enter the following:

```
Is this a standby controller? [NO]: YES
```

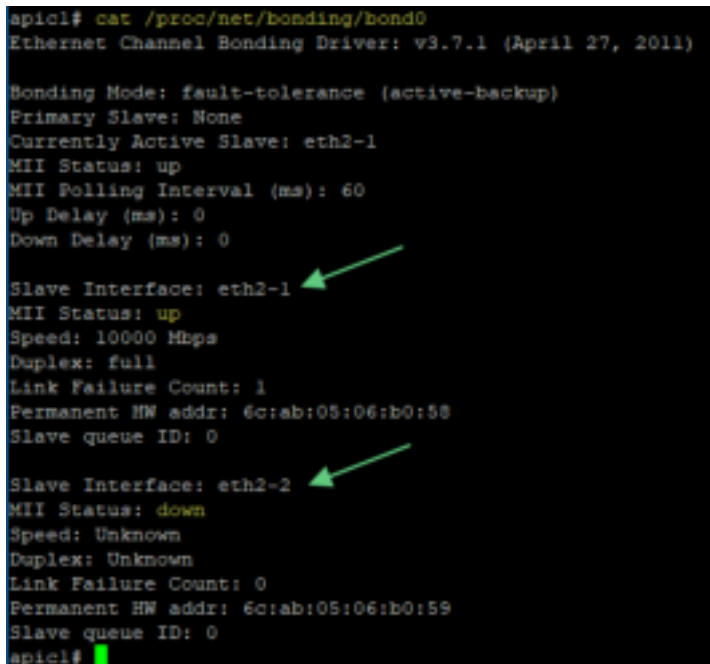
In the case of strict mode, you must approve the controller.

Troubleshooting the New Cluster

In most cases a new cluster member will not join the cluster due to incorrect configuration parameters with the infra VLAN, TEP pool, fabric name, and multicast pool or incorrect cabling. You will need to double-check these. Keep in mind it takes a bit of time for a new controller to fully converge, wait at least 10 minutes. You can always log into a non-ready cluster member using the rescue-user account. No password will be required if the cluster is in discovery mode. If a password is required, use the admin password.

Step 1. Verify the physical interfaces toward the fabric.

Ensure interfaces toward the fabric are up. You can enter the `cat /proc/net/bonding/bond0` command. At least one interface must be up. It is a necessary and sufficient condition to establish cluster membership. However, if a single interface is up then a major or critical fault will be raised in the Cisco APIC.



```
apicl# cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)

Bonding Mode: fault-tolerance (active-backup)
Primary Slave: None
Currently Active Slave: eth2-1
MII Status: up
MII Polling Interval (ms): 60
Up Delay (ms): 0
Down Delay (ms): 0

Slave Interface: eth2-1
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 1
Permanent HW addr: 6c:ab:05:06:b0:58
Slave queue ID: 0

Slave Interface: eth2-2
MII Status: down
Speed: Unknown
Duplex: Unknown
Link Failure Count: 0
Permanent HW addr: 6c:ab:05:06:b0:59
Slave queue ID: 0
apicl#
```

Figure 15: Verifying the Physical Interfaces Toward the Fabric

You can run the `acidiag bond0test` command to validate the cabling:

```

apic1# acidiag bond0test
Ping test via eth2-1
Setting eth2-2 down
Setting eth2-1 up
Checking whether currently active interface is eth2-1, loopCount: 0
Pinging 10.0.0.30 via bond0/eth2-1
PING 10.0.0.30 (10.0.0.30) 56(84) bytes of data:
64 bytes from 10.0.0.30: icmp_seq=1 ttl=64 time=0.215 ms
64 bytes from 10.0.0.30: icmp_seq=2 ttl=64 time=0.219 ms
64 bytes from 10.0.0.30: icmp_seq=3 ttl=64 time=0.226 ms
64 bytes from 10.0.0.30: icmp_seq=4 ttl=64 time=0.244 ms
64 bytes from 10.0.0.30: icmp_seq=5 ttl=64 time=0.199 ms

--- 10.0.0.30 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4089ms
rtt min/avg/max/mdev = 0.199/0.220/0.244/0.021 ms
Ping via eth2-1 passed
Setting eth2-2 up
Ping test via eth2-2

```

Figure 16: Validating the Cabling

Step 2. Check the cluster health from the new Cisco APIC.

At the prompt of the new Cisco APIC using either the console, VGA output, or SSH, use the "acidiag avread" command to examine this Cisco APIC's view of the cluster. If you do not see the other Cisco APIC servers, there is probably a configuration parameter mismatch, a cabling problem or a software release problem. A healthy 3-node cluster shows exactly three active servers in the output of the acidiag avread command:

```

apic1# acidiag avread | grep -i active | wc -l
3
apic1#

```

Figure 17: Checking the Cluster Health

Step 3. Verify the database consistency.

Cisco APIC stores all configuration and runtime data in a distributed database that is broken down into units called shards. Shards are triplicated within a cluster for resiliency purposes. The command allows you inspect whether the database is fully synchronized across the cluster with a consistent data layer. Use the acidiag rvread command and ensure no forward or backward slashes appear anywhere in the shard or service ID matrix:

```

spit# ended error
^~ unexpected state: /~unexpected state:
0- 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34
0- 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
Explores are in expected states and are mutated by proper apis's

```

Figure 18: Verifying the Database Consistency