

Troubleshoot ACI Management and Core Services - In-band and Out-of-band Management

Contents

[Introduction](#)

[Background Information](#)

[In-band and out-of-band management](#)

[APIC Connectivity Preferences](#)

[Scenario: Unable to reach management network](#)

[Out-of-Band Management Access](#)

[Out-of-band configuration verification](#)

[Static Node Management Addresses GUI verification](#)

[Out-of-band EPG - default](#)

[Out-of-band contract](#)

[External Management Network Instance Profile](#)

[In-band management configuration](#)

[Bridge Domain Subnet which will act as the in-band management gateway](#)

[Fault F0467 - inb EPG](#)

[In-band EPG](#)

[External EPG Instance Profile](#)

[Static Node Management Addresses](#)

Introduction

This document describes steps to troubleshoot ACI out-of-band (OOB) and in-band (INB) management.

Background Information

The material from this document was extracted from the [Troubleshooting Cisco Application Centric Infrastructure, Second Edition book](#) specifically the **Management and Core Services - In-band and out-of-band Management** chapter.

In-band and out-of-band management

ACI fabric nodes have two options for management connectivity; out-of-band (OOB), which governs the dedicated physical management port on the back of the device, or in-band (INB), which is provisioned using a specific EPG/BD/VRF in the management tenant with a degree of configurable parameters. There's an OOB EPG present in the management ('mgmt') tenant, but it's there by default and can't be modified. It only allows configuration of Provided OOB Contracts. On the APIC, the OOB interface is observed in the 'ifconfig' command output as 'oobmgmt' and

the in-band interface will be represented by the 'bond.x' interface, where is the encap VLAN configured for the in-band EPG.

```
apicl# ifconfig oobmgmt
oobmgmt: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
  inet 192.168.4.20 netmask 255.255.255.0 broadcast 192.168.4.255
  inet6 fe80::7269:5aff:feca:2986 prefixlen 64 scopeid 0x20
  ether 70:69:5a:ca:29:86 txqueuelen 1000 (Ethernet)
  RX packets 495815 bytes 852703636 (813.2 MiB)
  RX errors 0 dropped 0 overruns 0 frame 0
  TX packets 432927 bytes 110333594 (105.2 MiB)
  TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

```
apicl# ifconfig bond0.300
bond0.300: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1496
  inet 10.30.30.254 netmask 255.255.255.0 broadcast 10.30.30.255
  inet6 fe80::25d:73ff:fecl:8d9e prefixlen 64 scopeid 0x20
  ether 00:5d:73:c1:8d:9e txqueuelen 1000 (Ethernet)
  RX packets 545 bytes 25298 (24.7 KiB)
  RX errors 0 dropped 0 overruns 0 frame 0
  TX packets 6996 bytes 535314 (522.7 KiB)
  TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

On the leaf, the OOB interface is seen as 'eth0' in the 'ifconfig' command output and the INB is seen as a dedicated SVI. The user can view the interface with 'ifconfig' or with 'show ip interface vrf mgmt:' where is the name selected for the in-band VRF.

```
leaf101# show interface mgmt 0
mgmt0 is up
admin state is up,
  Hardware: GigabitEthernet, address: 00fc.baa8.2760 (bia 00fc.baa8.2760)
  Internet Address is 192.168.4.23/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec
  reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, medium is broadcast
  Port mode is routed
  full-duplex, 1000 Mb/s
  Beacon is turned off
  Auto-Negotiation is turned on
  Input flow-control is off, output flow-control is off
  Auto-mdix is turned off
  EtherType is 0x0000
  30 seconds input rate 3664 bits/sec, 4 packets/sec
  30 seconds output rate 4192 bits/sec, 4 packets/sec
Rx
  14114 input packets 8580 unicast packets 5058 multicast packets
  476 broadcast packets 2494768 bytes
Tx
  9701 output packets 9686 unicast packets 8 multicast packets
  7 broadcast packets 1648081 bytes
```

```
leaf101# show ip interface vrf mgmt:inb
IP Interface Status for VRF "mgmt:inb-vrf"
vlan16, Interface status: protocol-up/link-up/admin-up, iod: 4, mode: pervasive
  IP address: 10.30.30.1, IP subnet: 10.30.30.0/24
  secondary IP address: 10.30.30.3, IP subnet: 10.30.30.0/24
  IP broadcast address: 255.255.255.255
```

IP primary address route-preference: 0, tag: 0

The **'show ip interface vrf mgmt:'** will show the in-band management BD subnet IP as the secondary IP address; this is expected output.

On spine switches the in-band management IP address is added as a dedicated loopback interface in the 'mgmt:' VRF. This implementation is thus different from the in-band management IP implementation on leaf switches. Observe the **'show ip int vrf mgmt:'** command output below on a spine switch


```
spine201# show ip interface vrf mgmt:inb
IP Interface Status for VRF "mgmt:inb"
lo10, Interface status: protocol-up/link-up/admin-up, iod: 98, mode: pervasive
  IP address: 10.30.30.12, IP subnet: 10.30.30.12/32
  IP broadcast address: 255.255.255.255
IP primary address route-preference: 0, tag: 0
```

Under the System Settings, there is a setting to select either the in-band or out-of-band connectivity preference for the APICs.

Only the traffic sent from the APIC will use the management preference selected in the 'APIC Connectivity Preferences'. The APIC can still receive traffic on either in-band or out-of-band, assuming either is configured. APIC uses the following forwarding logic:

- Packets that come in an interface and go out that same interface.
- Packets sourced from the APIC, destined to a directly connected network, go out the directly connected interface.
- Packets sourced from the APIC, destined to a remote network, prefer in-band or out-of-band based on the APIC Connectivity Preferences.

APIC Connectivity Preferences



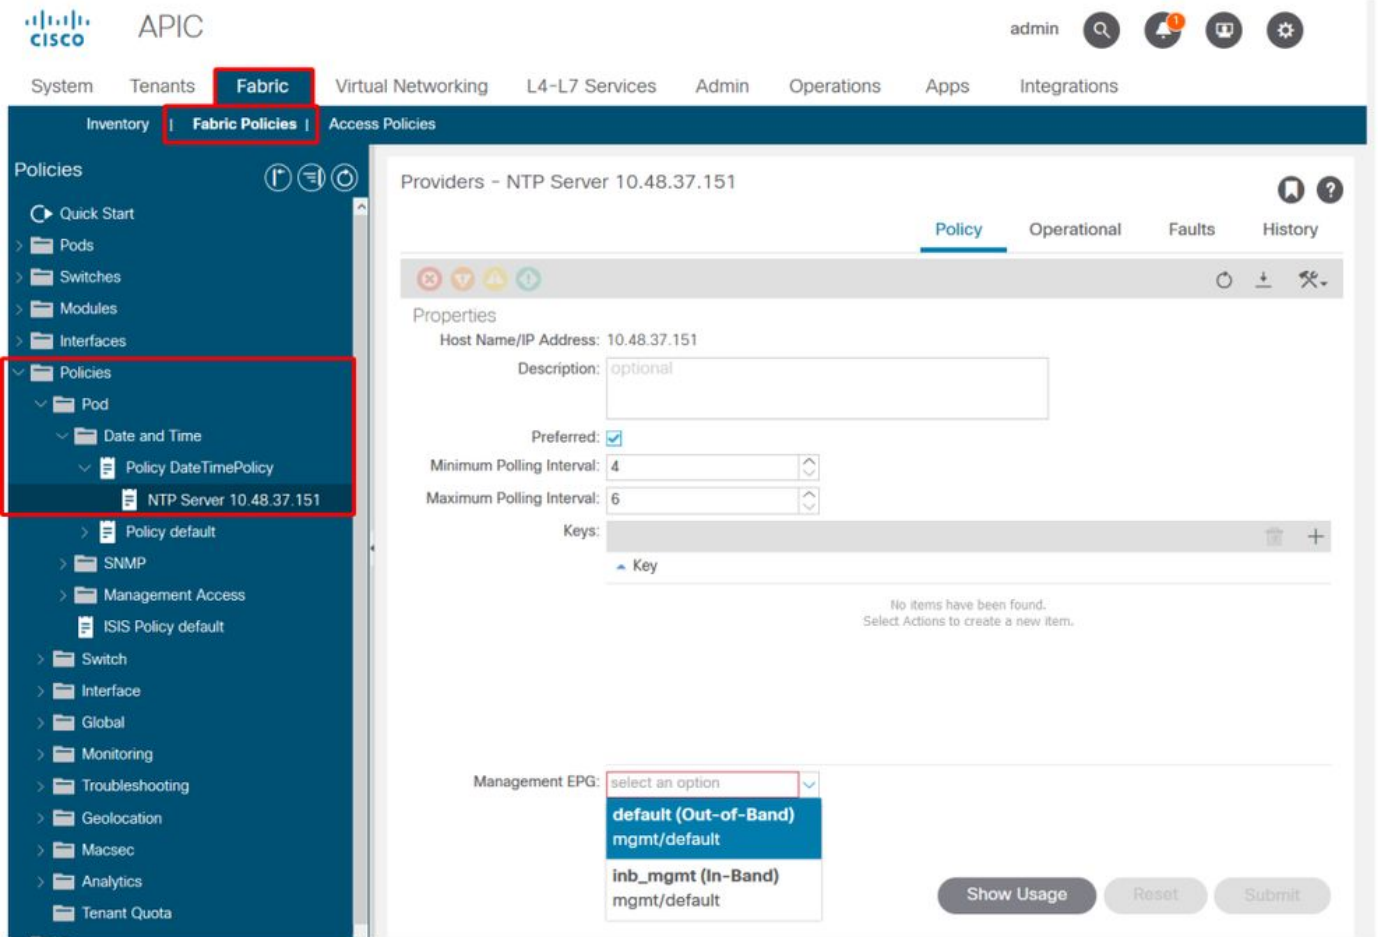
APIC routing table with OOB selected. Observe the metric value of 16 for the oobmgmt interface which is lower than the bond0.300 in-band management interface metric of 32. Meaning the oobmgmt out-of-band management interface will be used for outgoing management traffic.

```
apic1# bash
admin@apic1:~> route -n
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref    Use Iface
0.0.0.0          192.168.4.1    0.0.0.0        UG    16    0      0 oobmgmt
0.0.0.0          10.30.30.1     0.0.0.0        UG    32    0      0 bond0.300
```

APIC routing table with in-band selected. Observe the bond0.300 in-band management interface's metric if 8 which is now lower than the oobmgmt interface metric of 16. Meaning the bond0.300 in-band management interface will be used for outgoing management traffic.

```
admin@apic1:~> route -n
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref    Use Iface
0.0.0.0          10.30.30.1     0.0.0.0        UG    8     0      0 bond0.300
0.0.0.0          192.168.4.1    0.0.0.0        UG   16    0      0 oobmgmt
```

The leaf and spine node management preferences are not affected by this setting. These connectivity preferences are selected under the protocol policies. Below is an example for NTP.



If in-band is selected under the APIC Connectivity Preferences, but then out-of-band is selected under the protocol, which interface with the protocol packet use?

- The APIC Connectivity Preference will always take precedence over the protocol selection on the APIC.
- The leaf nodes are the opposite, they only reference the selection under the protocol.

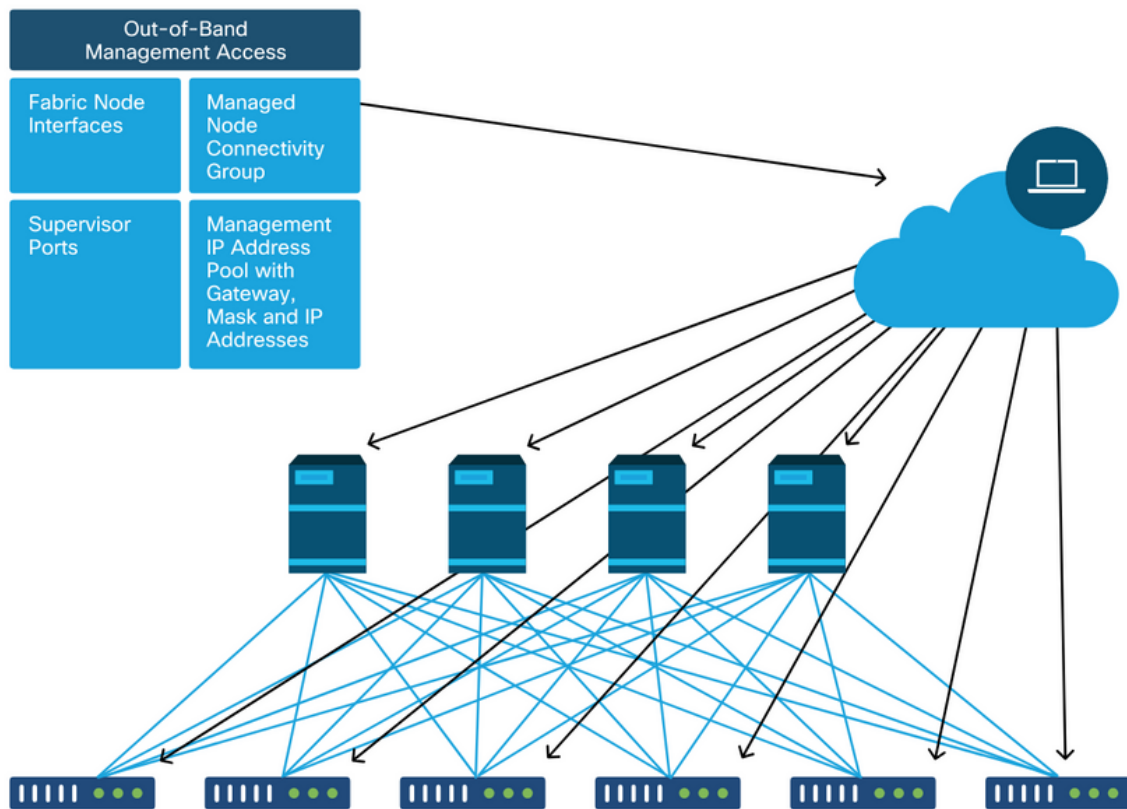
Scenario: Unable to reach management network

If the user is unable to reach the management network, it may be due to a number of different issues, but they can always use the same methodology to isolate the issue. The assumption in this scenario is that the user cannot reach any devices in the management network from behind their L3Out.

- Verify the APIC connectivity preference. This is outlined in figure 'APIC Connectivity Preferences', and the options are OOB or in-band.
- Depending on which preference is selected, verify the configuration is correct, the interfaces are up, the default gateway is reachable via the selected interface, and there are no drops on the path of the packet.

Do not forget to check for faults in each section of configuration in the GUI. However, some configuration mistakes can manifest in unexpected states, but a fault may be generated in another section than the one the user would initially consider.

Out-of-Band Management Access



Out-of-band configuration verification

For out-of-band configuration, there are four folders to verify under a special tenant called 'mgmt':

- Node Management Addresses.
- Node Management EPGs.
- Out-of-band Contracts (under Contracts).
- External Network Instance Profiles.

Node Management Addresses can either be assigned statically or from a pool. Below is an example of static address assignment. Verify that the type out-of-band IP addresses are assigned and that the default gateway is correct.

Static Node Management Addresses GUI verification

Static Node Management Addresses

Node ID	Name	Type	EPG	IPv4 Address	IPv4 Gateway	IPv6 Address	IPv6 Gateway
pod-1/node-1	bdsol-aci37-apic1	Out-Of-Band	default	10.48.176.57/24	10.48.176.1	::	::
pod-1/node-101	S1P1-Leaf101	Out-Of-Band	default	10.48.176.70/24	10.48.176.1	::	::
pod-1/node-102	S1P1-Leaf102	Out-Of-Band	default	10.48.176.71/24	10.48.176.1	::	::
pod-1/node-2	bdsol-aci37-apic2	Out-Of-Band	default	10.48.176.58/24	10.48.176.1	::	::
pod-1/node-201	S1P1-Spine201	Out-Of-Band	default	10.48.176.74/24	10.48.176.1	::	::
pod-1/node-202	S1P1-Spine202	Out-Of-Band	default	10.48.176.75/24	10.48.176.1	::	::
pod-1/node-301	S1P2-Leaf301	Out-Of-Band	default	10.48.176.72/24	10.48.176.1	::	::
pod-1/node-302	S1P2-Leaf302	Out-Of-Band	default	10.48.176.73/24	10.48.176.1	::	::
pod-1/node-401	S1P2-Spine401	Out-Of-Band	default	10.48.176.76/24	10.48.176.1	::	::
pod-1/node-402	S1P2-Spine402	Out-Of-Band	default	10.48.176.77/24	10.48.176.1	::	::
pod-2/node-3	bdsol-aci37-apic3	Out-Of-Band	default	10.48.176.59/24	10.48.176.1	::	::

The out-of-band EPG should be present under the Node Management EPGs folder.

Out-of-band EPG - default

Out-of-Band EPG - default

Policy | Faults | History

Properties

Name: default

Tags:

Configuration Issues:

Configuration State: applied

Class ID: 16387

QoS Class: Unspecified

Provided Out-of-Band Contracts:

OOB Contract	Tenant	Type	QoS Class	State
OOB-default	mgmt	oobrc-OOB-default	Unspecified	formed

Show Usage | Reset | Submit

The contracts which govern which management services are provided from the out-of-band EPG are special contracts that are configured in the out-of-band contracts folder.

Out-of-band contract

The screenshot shows the APIC interface with the 'Tenants' tab selected. In the left-hand navigation menu, 'Out-Of-Band Contracts' and its sub-item 'OOB-default' are highlighted with a red box. The main content area displays the configuration for 'Contract Subject - OOB-default'. The 'Policy' tab is active, showing a 'General' sub-tab. The configuration includes a 'Name' field set to 'OOB-default', a 'Description' field with the value 'optional', and a checked 'Reverse Filter Ports' option. Below this, a 'Filters' table is shown:

Name	Tenant	State	Action
default	common	formed	Permit

At the bottom of the configuration page, there are buttons for 'Show Usage', 'Reset', and 'Submit'.

Next, verify the External Management Network Instance Profile is created and that the correct out-of-band contract is configured as the 'Consumed Out-Of-Band Contract'.

External Management Network Instance Profile

The screenshot shows the APIC interface with the 'Tenants' tab selected. In the left-hand navigation menu, 'External Management Network Instance Profile' and its sub-item 'default' are highlighted with a red box. The main content area displays the configuration for 'External Management Network Instance Profile - default'. The 'Policy' tab is active, showing a 'General' sub-tab. The configuration includes a 'Name' field set to 'default', a 'Tags' field, and a 'QoS Class' dropdown menu set to 'Unspecified'. Below this, a 'Consumed Out-of-Band Contracts' table is shown:

Out-of-Band Contract	Tenant	Type	QoS Class	State
OOB-default	mgmt	oobrc-OOB-default	Unspecified	formed

At the bottom of the configuration page, there are buttons for 'Show Usage', 'Reset', and 'Submit'.

The next items to verify are the interface state and cabling, and then the connectivity to the gateway.

- To check if the oobmgmt interface is up, enter 'ifconfig oobmgmt' on the APIC CLI. Verify that

the interface flags are 'UP' and 'RUNNING', that the correct IP address is configured, and that packets are increasing in the RX and TX counters. If any checks are missing, then verify the correct cables are being used and that they are connected to the correct physical management ports on the APIC. The management ports will be labelled Eth1-1 and Eth1-2 and recent hardware have oobmgmt stickers to indicate the out-of-band interface. For more information about the physical out-of-band mgmt ports on the back of an APIC, please refer to the section "Initial fabric setup" in chapter "Fabric discovery".

```
apicl# ifconfig oobmgmt
oobmgmt: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 192.168.4.20 netmask 255.255.255.0 broadcast 192.168.4.255
inet6 fe80::7269:5aff:feca:2986 prefixlen 64 scopeid 0x20
ether 70:69:5a:ca:29:86 txqueuelen 1000 (Ethernet)
RX packets 295605 bytes 766226440 (730.7 MiB)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 253310 bytes 38954978 (37.1 MiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

- To check the network connectivity through the OOB, use ping to test the path of the packet through the out-of-band network.

```
apicl# ping 192.168.4.1
PING 192.168.4.1 (192.168.4.1) 56(84) bytes of data.
64 bytes from 192.168.4.1: icmp_seq=1 ttl=255 time=0.409 ms
64 bytes from 192.168.4.1: icmp_seq=2 ttl=255 time=0.393 ms
64 bytes from 192.168.4.1: icmp_seq=3 ttl=255 time=0.354 ms
```

Using traceroute in the bash shell on the APIC, trace the connectivity to the end user. If the traceroute is incomplete, login to this device (if accessible) and ping the oobmgmt interface and ping the host. Depending on which direction fails, troubleshoot the issue as a traditional networking problem.

Traceroute works by sending UDP packets with an increasing TTL, starting with 1. If a router receives the packet with TTL 1 and needs to route it, it drops the frame and sends back an ICMP unreachable message to the sender. Each hop is sent 3 UDP packets at the current TTL, and asterisks represent attempts where an ICMP unreachable / TTL Exceeded packet was not received. These 3 asterisk blocks are expected in most networks as some routing devices have ICMP unreachable / TTL Exceeded messages disabled, so when they receive TTL 1 packets that they need to route, they simply drop the packet and do not send the message back to the sender.

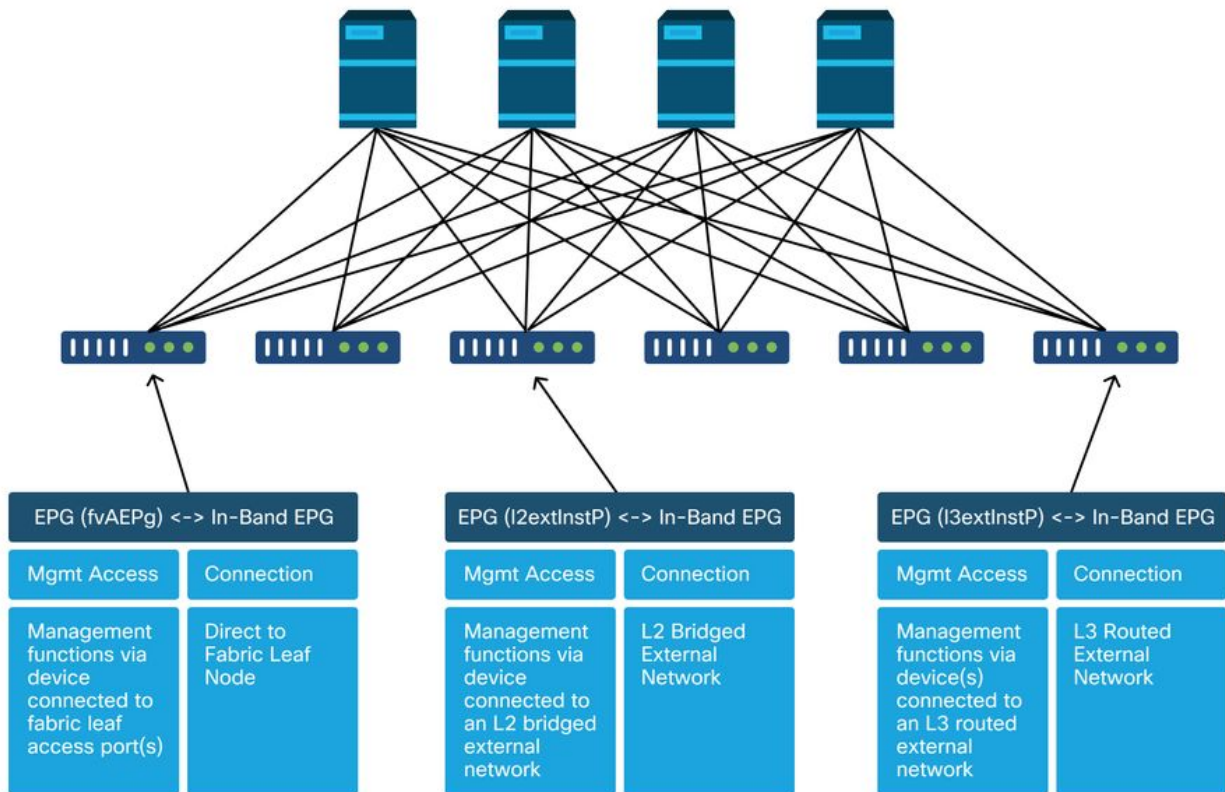
```
apicl# bash
admin@apicl:~> traceroute 10.55.0.16
traceroute to 10.55.0.16 (10.55.0.16), 30 hops max, 60 byte packets
 1  192.168.4.1 (192.168.4.1)  0.368 ms  0.355 ms  0.396 ms
 2  * * *
 3  * * *
 4  10.0.255.221 (10.0.255.221)  6.419 ms  10.0.255.225 (10.0.255.225)  6.447 ms  *
 5  * * *
 6  * * *
```

The leaf switches have access to the tcpdump command, which can be used to verify which packets are traversing the oobmgmt interface. The example below captures on 'eth0', which is the oobmgmt interface used on the leaf and spine switches, and uses '-n' option for tcpdump to give the IP addresses used instead of the DNS names, and then filtering specifically for NTP packets (UDP port 123). Recall that in the previous example the leaf is polling NTP server 172.18.108.14. Below, the user can verify that NTP packets are being transmitted via the out-of-band interface and also that the leaf is receiving a response from the server.

```
fab1-leaf101# tcpdump -n -i eth0 dst port 123
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
16:49:01.431624 IP 192.168.4.23.123 > 172.18.108.14.123: NTPv4, Client, length 48
16:49:01.440303 IP 172.18.108.14.123 > 192.168.4.23.123: NTPv4, Server, length 48
```

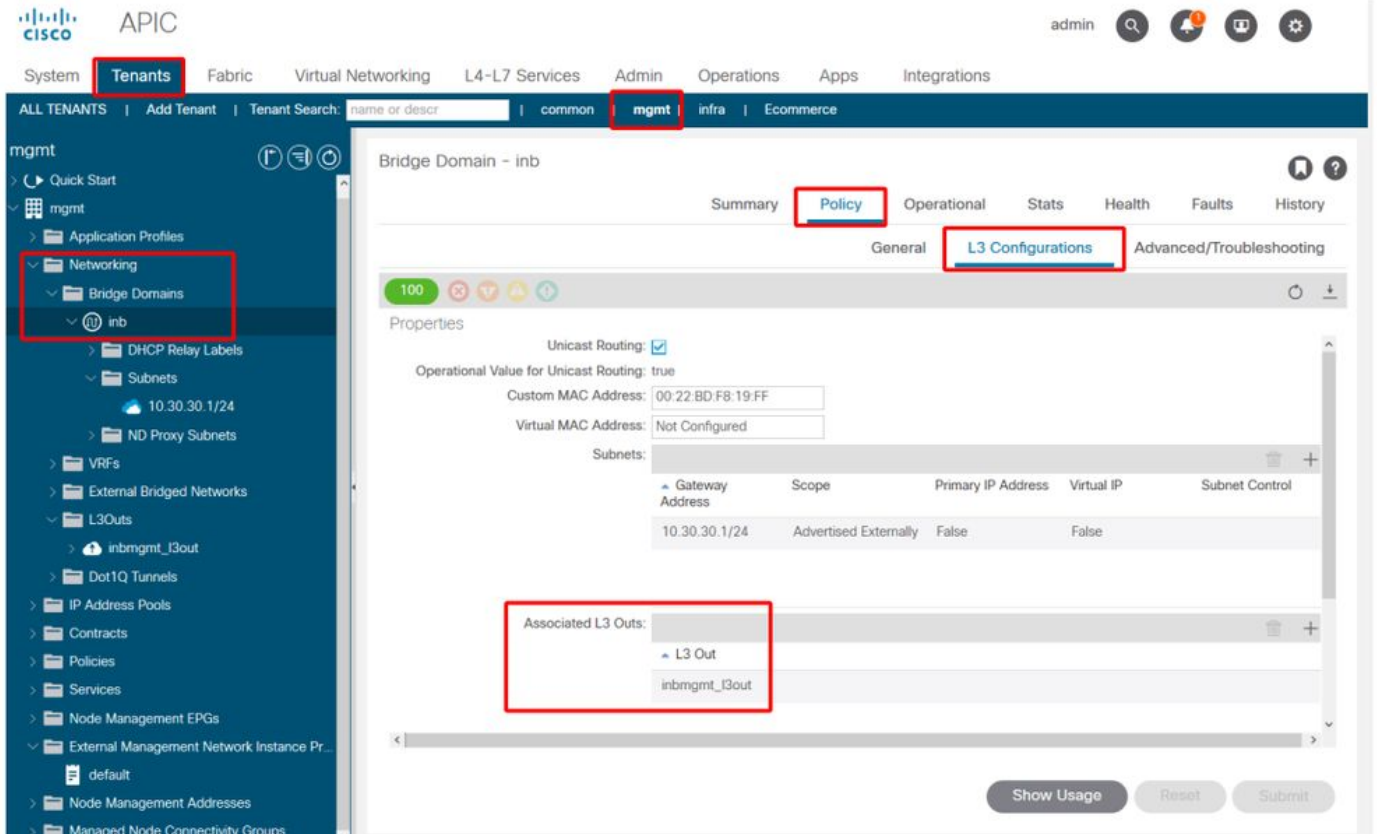
The in-band management configuration requires specific considerations for Layer 2 or Layer 3 deployments. This example will only cover the Layer 3 deployment and troubleshooting.

In-band management configuration



Verify that there is a BD in the mgmt tenant with a subnet from which in-band node mgmt addresses will be allocated to the fabric nodes for in-band connectivity, and make sure that the L3Out is associated under the in-band management BD.

Bridge Domain Subnet which will act as the in-band management gateway



Verify an in-band node management EPG is present. As per screenshot below, the in-band EPG names are denoted in the GUI with the prefix 'inb-'. Verify the in-band EPG encap VLAN is associated correctly with a VLAN pool.

The encapsulation VLAN configured in the in-band management EPG needs to be allowed by Access Policies: 'inb mgmt EPG encap VLAN > VLAN Pool > Domain > AEP > Interface Policy Group > Leaf Interface Profile > Switch Profile'. If the supporting access policies are not configured, a fault with code F0467 will be raised as per below screenshot.

Fault F0467 - inb EPG

8589935303

ID: 8589935303

Description: Fault delegate: Configuration failed for uni/tn-mgmt/mgmt-default/inb-inbmgmt due to Invalid VLAN Configuration, debug message: i
vlan-300STP Segment Id not present for Encap. Either the EpG is not associated with a domain or the domain does not have this vlan i

Severity: minor

Related Object: [uni/tn-mgmt/mgmt-default/inb-inbmgmt](#)

Generated From: topology/pod-1/node-101/local/svc-policyelem-id-0/uni/epp/inb-[uni/tn-mgmt/mgmt-default/inb-inbmgmt]/nwissues

Created: 2019-10-03T02:23:04.637+00:00

Code: F0467

Type: Config

Cause: configuration-failed

Change Set:

Action: deletion

Domain: Tenant

Life Cycle:

Time Occurred: 1

Event Status: false

Verify that the bridge domain is the same as the one created above for the in-band subnet. Lastly, verify that there is a Provided Contract configured on the in-band management EPG, which is consumed by the external EPG.

In-band EPG

The screenshot shows the Cisco APIC interface. The top navigation bar includes 'System', 'Tenants', 'Fabric', 'Virtual Networking', 'L4-L7 Services', 'Admin', 'Operations', 'Apps', and 'Integrations'. The 'Tenants' tab is active, and the 'mgmt' tenant is selected. The left sidebar shows a tree view of the tenant configuration, with 'Node Management EPGs' expanded and 'In-Band EPG - inb_mgmt' selected. The main content area displays the configuration for 'In-Band EPG - inb_mgmt'. The 'Policy' tab is active, showing a '100' status indicator and a 'Policy' sub-tab. The 'Properties' section includes:

- Name: inb_mgmt
- Tags: (empty)
- Encap: vlan-300 (e.g., vlan-1)
- Configuration Issues: (empty)
- Configuration State: applied
- Class ID: 32770
- QoS Class: Unspecified
- Bridge Domain: inb
- Resolved Bridge Domain: inb
- Provided Contracts: (table below)

Name	Tenant	Type	QoS Class	Match Type	State
default	common	Contract	Unspecified	AtleastOne	formed

At the bottom right, there are buttons for 'Show Usage', 'Reset', and 'Submit'.

External EPG Instance Profile

External EPG Instance Profile - Inband-Out

Policy | Operational | Stats | Health | Faults | History

General | Contracts | Subject Labels | EPG Labels

Provided Contracts | Consumed Contracts | Contract Interfaces | Taboo Contracts | Inherited Contra

Name	Tenant	Type	QoS Class	State
default	common	Contract	Unspecified	formed

Similar to out-of-band, fabric node in-band mgmt IP addresses can be statically assigned or dynamically assigned from a pre-selected range. Verify the addresses applied for type in-band match the previous BD subnet that was configured. Also verify that the default gateway is correct.

Static Node Management Addresses

Static Node Management Addresses

Node ID	Name	Type	EPG	IPv4 Address	IPv4 Gateway	IPv6 Address	IPv6 Gateway
pod-1/node-1	bdsol-aci37-apic1	Out-Of-Band	default	10.48.176.57/24	10.48.176.1	::	::
pod-1/node-101	S1P1-Leaf101	In-Band	inb_mg...	10.30.30.101/24	10.30.30.1	::	::
pod-1/node-101	S1P1-Leaf101	Out-Of-Band	default	10.48.176.70/24	10.48.176.1	::	::
pod-1/node-102	S1P1-Leaf102	Out-Of-Band	default	10.48.176.71/24	10.48.176.1	::	::
pod-1/node-2	bdsol-aci37-apic2	Out-Of-Band	default	10.48.176.58/24	10.48.176.1	::	::
pod-1/node-201	S1P1-Spine201	Out-Of-Band	default	10.48.176.74/24	10.48.176.1	::	::
pod-1/node-202	S1P1-Spine202	Out-Of-Band	default	10.48.176.75/24	10.48.176.1	::	::
pod-1/node-301	S1P2-Leaf301	Out-Of-Band	default	10.48.176.72/24	10.48.176.1	::	::
pod-1/node-302	S1P2-Leaf302	Out-Of-Band	default	10.48.176.73/24	10.48.176.1	::	::
pod-1/node-401	S1P2-Spine401	Out-Of-Band	default	10.48.176.76/24	10.48.176.1	::	::
pod-1/node-402	S1P2-Spine402	Out-Of-Band	default	10.48.176.77/24	10.48.176.1	::	::
pod-2/node-3	bdsol-aci37-apic3	Out-Of-Band	default	10.48.176.59/24	10.48.176.1	::	::

If everything has been configured correctly, and there are no faults in any above-mentioned section, the next step is to ping between the switches and/or APICs to verify that in-band connectivity is working correctly inside ACI.

The spine nodes will not respond to ping on the in-band as they use loopback interfaces for connectivity which do not respond to ARP.

The in-band interface used on the leaf switches is kpm_inb. Using a similar tcpdump capture, verify the packet is egressing the in-band CPU interface.

```
fab2-leaf101# tcpdump -n -i kpm_inb dst port 123
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on kpm_inb, link-type EN10MB (Ethernet), capture size 65535 bytes
16:46:50.431647 IP 10.30.30.3.123 > 172.18.108.14.123: NTPv4, Client, length 48
16:47:19.431650 IP 10.30.30.3.123 > 172.18.108.15.123: NTPv4, Client, length 48
```

Verify that the SVI used for in-band is **'protocol-up/link-up/admin-up'**.

```
fab1-leaf101# show ip interface vrf mgmt:inb-vrf
IP Interface Status for VRF "mgmt:inb-vrf"
vlan16, Interface status: protocol-up/link-up/admin-up, iod: 4, mode: pervasive
  IP address: 10.30.30.1, IP subnet: 10.30.30.0/24 secondary
  IP address: 10.30.30.3, IP subnet: 10.30.30.0/24
  IP broadcast address: 255.255.255.255
  IP primary address route-preference: 0, tag: 0
```