Cisco ACI In-Band Management Configuration for Hardware Flow Telemetry Export

Revised: March 2, 2023

Configuring ACI In-Band Management for Hardware Flow Telemetry Export

This document provides procedures for configuring Cisco ACI in-band management for Cisco Tetration hardware sensors.

Prerequisites for Configuring Cisco ACI In-Band Management for Hardware Flow Telemetry Export

These are the prerequisites for configuring Cisco Application Centric Infrastructure (ACI) in-band management for hardware flow telemetry.

Supported Hardware and Software

For supported ACI hardware switches, ACI software version and Tetration software version refer to the Tetration platform datasheet https://www.cisco.com/c/en/us/products/collateral/data-center-analytics/tetration-analytics/datasheet-c78-737256.html

Required `inb` VRF Under `mgmt` Tenant

You must use the `inb` VRF under the `mgmt` tenant because it is hardcoded in the Hardware Agent.

1. In your Cisco APIC system, navigate to the bridge domain page under the `mgmt` tenant:

   Tenants > **mgmt** > **Networking** > **Bridge Domains** > **inb**

2. On the **Bridge Domain - inb** page, click the **Policy** tab.

   The **General** subtab under the **Policy** tab should be selected automatically.

3. Locate the **VRF** field and verify that the `inb` VRF is selected for the `mgmt` tenant.

Required Use of Bridge Domain Subnet as Gateway for Hardware Agent For Spine Hardware Sensor

For the spine hardware sensor only, you must use the bridge domain subnet as the gateway for the Hardware Agent, which means that you will also need an L3Out to reach the telemetry collector. This is because the spine switch doesn?t apply the ARP for the node in-band management IP address.

The following figure shows an example of this configuration.

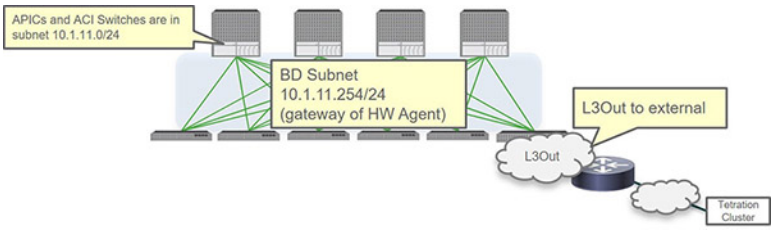

In-Band or Out-of-Band Considerations

- If you use both out-of-band and in-band for external connections, in-band is preferred for packets sourced from the APIC by default (for example, VMM integration).
- Cisco APIC uses the following forwarding logic:
  - Packets that come in from an interface go out from that same interface
  - Packets sourced from a Cisco APIC that are destined to a directly connected network go out from the directly connected interface
  - Packets sourced from a Cisco APIC that are destined to a remote network prefer in-band primarily, followed by out-of-band
- If you prefer to use out-of-band for external connections, navigate to:

  System > **System Settings** > **APIC Connectivity Preferences**

  Then select **ooband** in the **Interface to use for external connections** field.

Configure the Pod Policy

Before you can configure in-band management, you must first configure the pod policy. Configuring the pod policy consists of these tasks:

- Configuring the BGP route reflector
- Configuring NTP
- Enabling HTTP on the Cisco APIC

Procedure

Step 1 Determine which pod policy group is being used by your APIC system.

a) Navigate to the **Pod Selector** page:

Fabric > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default** > **default**

The **Pod Selector - default** page is displayed.

b) Locate the **Fabric Policy Group** field and note the name of the name of the pod policy group displayed in that field.

Step 2 Locate the BGP route reflector, Date and Time, and Management Access policies used by the pod policy group.

a) Navigate to the **Pod Policy Group** page:

Fabric > **Fabric Policies** > **Pods** > **Policy Groups** > ***pod_policy_group_name***

The **Pod Policy Group** page is displayed.

b) Locate the following fields in the **Pod Policy Group** page and note the policies used for each field:

- **BGP Route Reflector**
- **Date and Time**
- **Management Access**

Step 3 Configure the BGP route reflector.

The ACI fabric route reflectors use multiprotocol BGP (MP-BGP) to distribute external routes within the fabric. To enable route reflectors in the ACI fabric, the fabric administrator must select the spine switches that will be the route reflectors, and provide the autonomous system (AS) number. Once route reflectors are enabled in the ACI fabric, administrators can configure connectivity to external networks.

a) Navigate to the **BGP Route Reflector** page:

System > **System Settings** > **BGP Route Reflector**

b) Configure the following fields for the BGP route reflector, if they are not already configured:

- **Autonomous System Number**: The autonomous system number must match the leaf switch-connected router configuration if Border Gateway Protocol (BGP) is configured on the router. If you are using routes learned using static or Open Shortest Path First (OSPF), the autonomous system number value can be any valid value.

  The Autonomous System Number can be in 4-byte asplain format from 1 to 4294967295.

- **Route Reflector Nodes**: Configure up to two spine nodes as route reflectors. For redundancy, configure primary and secondary route reflectors.

Step 4 Configure NTP.

a) Navigate to the **Date and Time Policy** page:

Fabric > **Fabric Policies** > **Policies** > **Pod** > **Date and Time**

b) Verify that NTP servers have been configured under the **NTP Servers** field in the **Date and Time Policy** page.

For more information on configuring NTP, see the "Time Synchronization and NTP" section in the *Cisco APIC Basic Configuration Guide* on the APIC documentation page.

Step 5 Enable HTTP.

You must have HTTP enabled because the switches download the Hardware Agent from the APIC through HTTP.

In the Cisco APIC 6.0(1) release and earlier:

a) Navigate to the **Management Access** page:

Fabric > **Fabric Policies** > **Policies** > **Pod** > **Management Access**

b) Locate the **HTTP** area in the **Management Access** page and verify that the entry in the **Admin State** field is set to **Enabled**.

If the field is set to **Disabled**, change the setting to **Enabled** and click **Submit**.

In the Cisco APIC 6.0(2) release and later:

a) Navigate to the **Management Access** page:

Fabric > **Fabric Policies** > **Policies** > **Pod** > **Management Access** > **policy_name**

b) In the **Work** pane, choose the **Policy** > **Web Access** tab.

c) Locate the **HTTP** area in the **Management Access** page and verify that the entry in the **Admin State** field is set to **Enabled**.

If the field is set to **Disabled**, change the setting to **Enabled** and click **Submit**.

What to do next

Go to Configuring Cisco ACI In-Band Management for Hardware Flow Telemetry, on page 5.

Configuring Cisco ACI In-Band Management for Hardware Flow Telemetry

The following topology is used as an example configuration for these procedures.

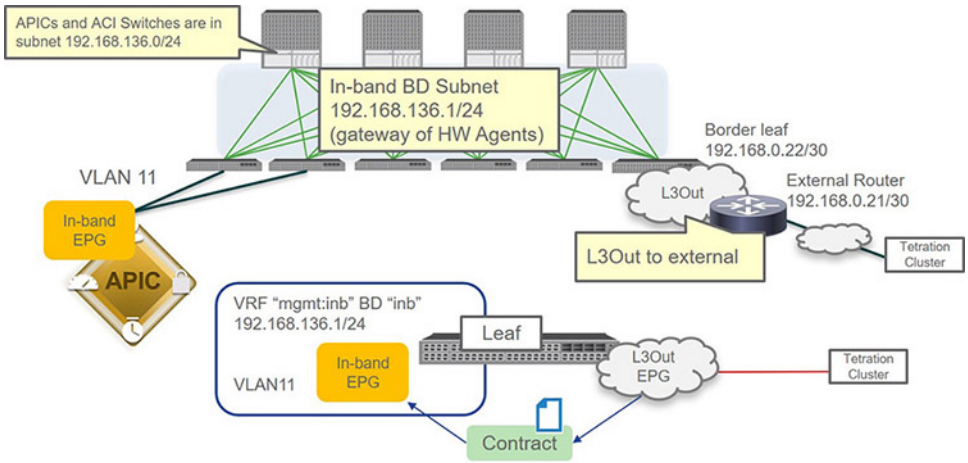

Before you begin

- Verify that you have reviewed and followed the information provided in Prerequisites for Configuring Cisco ACI In-Band Management for Hardware Flow Telemetry Export, on page 2.
- Verify that the pod policy is configured correctly using the information provided in Configure the Pod Policy, on page 3.

Procedure

Step 1 Configure the VLAN pool.

a) Navigate to the **Create VLAN Pool** page:

Fabric > **Access Policies** > **Pools** > **VLAN**, then right-click and choose **Create VLAN Pool**.

b) On the **Create VLAN Pool** page, perform the following actions:

1. Enter a name for the VLAN pool.
2. (Optional) Enter a description for the VLAN pool.
3. In the **Allocation Mode** field, select **Static Allocation**.

   This is typically used when the pool will be referenced from a static source, such as a static path binding for an EPG for use with new-deployment servers.
4. Click Add (+) in the **Encap Blocks** area to add an encapsulation block.

   The encapsulation blocks define the range of VLANs in the VLAN pool.
5. On the **Create Ranges** page, enter the following information:

- **Range**: Enter a value in this field. For this use case, we would enter **11** in this field.
- **Allocation Mode**: Choose **Static Allocation**.
- Role: Choose **External or On the wire encapsulations**.

Then click **OK** to save the values entered on the **Create Ranges** page.

c) On the **Create VLAN Pool** page, click **Submit** to save the values entered on this page.

Step 2 Configure the physical domain and AEP.

a) Navigate to the **Create Physical Domain** page:

Fabric > **Access Policies** > **Physical and External Domains** > **Physical Domains**, then right-click and choose **Create Physical Domain**.

b) On the **Create Physical Domain** page, perform the following actions:

1. Enter a name for the physical domain.
2. In the **Associated Attachable Entity Profile** field, select the AEP that is used for APIC connectivity.

   This could be the **default** AEP or some other AEP. Select the AEP that is used for APIC connectivity so that you are deploying the management EPG for the APIC in-band management interface.
3. In the **VLAN Pool** field, choose the VLAN pool that you configured in the previous step.
4. Click **Submit** to save the values entered on this page.

Step 3 Apply the access policy to the interface connecting to the Cisco APIC.

a) Create a leaf access port policy group by navigating to the **Create Leaf Access Port Policy Group** page:

Fabric > **Access Policies** > **Interfaces** > **Leaf Interfaces** > **Policy Groups** > **Leaf Access Port**, then right-click and choose **Create Leaf Access Port Policy Group**.

b) Enter a name for the leaf access port policy group, then, in the **Attached Entity Profile** field, choose the AEP that has the domain with the VLAN pool for the in-band management, and then click **Submit**.

c) Create a leaf interface policy by navigating to the **Create Leaf Interface Profile** page:

Fabric > **Access Policies** > **Interfaces** > **Leaf Interfaces** > **Profiles**, then right-click and choose **Create Leaf Interface Profile**.

d) Enter a name for the leaf interface profile, then click Add (+) in the **Interface Selectors** area.

e) In the **Create Access Port Selector** page, enter the necessary information, then, in the **Interface Policy Group** field, choose the leaf access port policy group you created in the previous steps.

f) Click **OK** to complete the configuration in the **Create Access Port Selector** page, then click **Submit** to complete the configuration in the **Create Leaf Interface Profile** page.

g) Create a leaf profile by navigating to the **Create Leaf Profile** page:

Fabric > **Access Policies** > **Switches** > **Leaf Switches** > **Profiles**, then right-click and choose **Create Leaf Profile**.

h) On the **Create Leaf Profile** page, enter the necessary information:

- In the **Leaf Selectors** area, select the necessary leaf switches and configure the switch selector information for those leaf switches.

- In the **Interface Selector Profiles** area, choose the leaf interface profile that you created in the previous set of steps.

i) Click **Finish** in the **Create Leaf Profile** page.

Step 4 Configure the in-band management EPG.

a) Navigate to the **Create In-Band Management EPG** page:

Tenant > **mgmt** > **Node Management EPGs**, then right-click and choose **Create In-Band Management EPG**.

b) Enter the necessary information on the **Create In-Band Management EPG** page, specifically these fields:

- **Name**: Leave `default` as the name for the in-band management EPG.
- **Encap**: Enter the access encapsulation. For example, for this use case, you would enter vlan-11 to match the information that you entered in Step 1.b, on page 5.
- **Bridge Domain**: The **inb** bridge domain under the `mgmt` tenant.

  This **inb** bridge domain is the bridge domain that was mentioned in the Required inb VRF Under mgmt Tenant, on page 2 section earlier in this document. Technically, this could be a different bridge domain, as long as it is in the `mgmt` VRF.

c) Click **Submit**.

Click the new in-band management EPG to be displayed under the **Node Management EPGs** area in the left navigation tree and verify that no fault is displayed for the new EPG.

Step 5 Assign in-band management IP addresses to the leaf and spine switches.

a) Navigate to the **Create Node Management Addresses** page:

Tenant > **mgmt** > **Node Management Addresses**, then right-click and choose **Create Node Management Addresses**.

b) Enter the necessary info on the **Create Node Management Addresses** page:

1. In the **Select Nodes By** field, choose **Specific**.
2. In the **Nodes** area, select the specific nodes for the leaf and spine switches.
3. In the **Config** area, choose **In-Band Addresses**.
4. In the **In-Band Management EPG** field, select the `default` in-band management EPG that you configured in the previous step.
5. In the **In-Band Gateway** and **In-Band IP Addresses** fields, set the in-band gateway and IP address range for the switches.

c) Click **Submit**.

Step 6 Assign in-band management IP addresses to the APICs.

a) Navigate to the **Create Node Management Addresses** page:

Tenant > **mgmt** > **Node Management Addresses**, then right-click and choose **Create Node Management Addresses**.

b) Enter the necessary information on the **Create Node Management Addresses** page:

1. In the **Select Nodes By** field, choose **Specific**.
2. In the **Nodes** area, select the specific nodes for the APICs (shown as `controller` under the **Role** column).
3. In the **Config** area, choose **In-Band Addresses**.
4. In the **In-Band Management EPG** field, select the `default` in-band management EPG that you configured in the previous step.
5. In the **In-Band Gateway** and **In-Band IP Addresses** fields, set the in-band gateway and IP address range for the switches.

c) Click **Submit**.

Step 7 Configure the `inb` bridge domain subnet.

a) Navigate to the **Create Subnet** page:

Tenant > **mgmt** > **Networking** > **Bridge Domains** > **inb**, then right-click and choose **Create Subnet**.

b) Enter the necessary info on the **Create Subnet** page:

1. In the **Gateway IP** field, enter the bridge domain subnet that will be used as the gateway for the hardware agent.
2. In the **Scope** area, click **Advertise Externally**.

c) Click **Submit**.

Step 8 Verify that the configurations have been completed successfully thus far.

a) Navigate to the **Node Management Addresses** page for the leaf and spine switches:

Tenant > **mgmt** > **Node Management Addresses** > **Switches**.

b) In the `Nodes Within the Policy` area, verify that the leaf and spine switches are listed correctly in the **In-Band Management IP** and **In-Band Management Gateway** column.

c) Verify that the APIC and the switches can ping each other:

```
Leaf1# show ip route vrf mgmt:inb
<snip>
192.51.100.0/24, ubest/mbest: 1/0, attached, direct, pervasive
    *via 10.0.200.66%overlay-1, [1/0], 00:00:40, static
192.51.100.1/32, ubest/mbest: 1/0, attached, pervasive
    *via 192.51.100.1, vlan6, [1/0], 00:00:40, local, local
192.51.100.24/32, ubest/mbest: 1/0, attached
    *via 192.51.100.24, vlan6, [1/0], 00:00:40, local, local

Leaf1# iping -V mgmt:inb 192.51.136.21
PING 192.51.100.21 (192.51.100.21) from 192.51.100.24: 56 data bytes
64 bytes from 192.51.100.21: icmp_seq=0 ttl=64 time=0.461 ms
```

Step 9 Configure the L3Out EPG.

a) Navigate to the **L3 Outside** page:

Tenant > **mgmt** > **Networking** > **L3Outs**, then right-click and choose **Create L3Out**.

b) Enter the necessary information in the **Create L3Out** wizard, specifically:

- In the **Name** field, enter a name for this L3Out (for example, `L3Out-mgmt`).

- In the **VRF** field, select `inb:mgmt`.
- In the **External EPG** pane, configure an external EPG for the L3Out.

Step 10 Create a contract between the L3Out EPG and the in-band management EPG.

a) Navigate to the **Create Contract** page for the L3Out EPG:

Tenant > **mgmt** > **Contracts** > **Standard**, then right-click and choose **Create Contract**.

b) Enter the necessary information in the **Create Contract** page, then click **Submit**.

c) Navigate to the **External Network Instance Profile** page for the L3Out EPG:

Tenant > **mgmt** > **Networking** > **L3Outs** > **L3Out-mgmt** > **Networks** > **Mgmt**

d) In the **External Network Instance Profile** page for the L3Out EPG, choose the contract that you just created in the **Provided Contracts** area.

e) Navigate to the **In-Band EPG** page for the in-band management EPG:

Tenant > **mgmt** > **Node Management EPGs** > ***in-band_management_EPG_name***

f) Choose the contract that you just created in the **Consumed Contracts** area.

Step 11 Configure the `inb` bridge domain for the L3Out.

a) Navigate to the **Bridge Domain - inb** page:

Tenant > **mgmt** > **Networking** > **Bridge Domains** > **inb**

b) Click the **Policy** tab, then the **L3 Configurations** subtab.

c) Click Add (+) in the **Associated L3Outs** area and choose the L3Out that you configured in Step 9, on page 8.

d) Click **Submit**.

Step 12 Verify that the switches can ping the telemetry collector IP address.

```
Leaf1# iping -V mgmt:inb 10.28.121.132

PING 10.28.121.132 (10.28.121.132) from 192.100.0.26: 56 data bytes
64 bytes from 10.28.121.132: icmp_seq=0 ttl=62 time=0.407 ms
64 bytes from 10.28.121.132: icmp_seq=1 ttl=62 time=0.455 ms
64 bytes from 10.28.121.132: icmp_seq=2 ttl=62 time=0.344 ms


Spine1# iping -V mgmt:inb 10.28.121.132

PING 10.28.121.132 (10.28.121.132): 56 data bytes
64 bytes from 10.28.121.132: icmp_seq=0 ttl=61 time=0.592 ms
64 bytes from 10.28.121.132: icmp_seq=1 ttl=61 time=0.433 ms
64 bytes from 10.28.121.132: icmp_seq=2 ttl=61 time=0.411 ms
```

Step 13 From the telemetry collector, download the hardware agent (RPM).

a) In the telemetry collector, click the **Action** icon, choose **Agent Config**, then click the **Hardware Agent Download** tab.

b) Locate the row with the latest version of the hardware agent and click the **Download** button in that row.

Step 14 Upload the hardware agent onto the APIC.

a) In the APIC GUI, navigate to:

Admin > **Firmware** > **Images**

b) Click the **Actions** button and choose **Add Firmware to APIC**.
c) In the **Firmware Image Location** field, select **Local**.
d) In the **File Name** field, click **Browse** and navigate to the location on your computer where you downloaded the hardware agent in the previous step.
e) Select that downloaded file, then click **Submit** on the **Add Firmware to APIC** page.

Step 15 Understand the upcoming steps on enabling leaf switches for analytics.

Before going through the next few steps in these procedures, it s helpful to understand what you will be doing and why.

- Telemetry collection is supported only on EX-model switches and later.
- EX/FX/FX2 model switches can run in one of these modes:
  - Analytics Priority
  - Netflow Priority
  - Telemetry Priority

  Analytics Priority is the mode used for telemetry collection, so this will be the mode that you will select in the upcoming steps.
- You will create a Node Control Policy to enable Analytics Priority for consistency.
- You will configure Node Control Policies under the Fabric Policies.

The following figure shows how the components that you will be configuring in the upcoming steps tie in with one another.

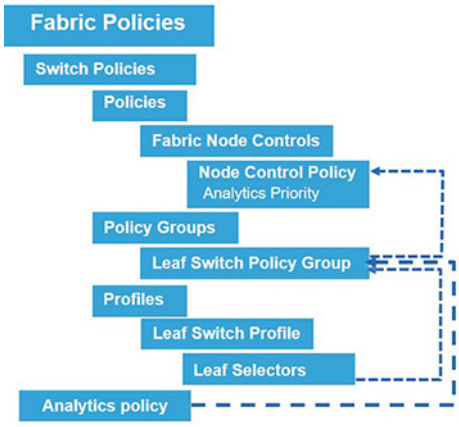

Step 16 Configure the fabric node control policy.

a) In the APIC interface, navigate to:

Admin > **Fabric** > **Fabric Policies** > **Policies** > **Monitoring** > **Fabric Node Controls** > **default**

b) In the **Feature Selection** area, click **Analytics Priority**.

Analytic priority downloads the telemetry sensor software for installation on the switches.

c) Click **Submit**.

Step 17 Create an analytics policy.

a) In the APIC interface, navigate to:

Admin > **Fabric** > **Fabric Policies** > **Policies** > **Analytics**, then right-click and choose **Create Analytics Policy**.

b) On the **Create Analytics Policy** page, enter the necessary information to create the analytics policy.

c) Click **Submit**.

Step 18 Create a leaf and spine switch policy group.

a) In the APIC interface, navigate to the **Create Leaf Switch Policy Group** page:

Admin > **Fabric** > **Fabric Policies** > **Switches** > **Leaf Switches** > **Policy Groups**, then right-click and choose **Create Leaf Switch Policy Group**.

b) On the **Create Leaf Switch Policy Group** page, enter the necessary information, specifically the following fields:

- **Analytics Policy**: Select the analytics policy that you created in the previous step.
- **Node Control Policy**: Select the `default` fabric node control policy that you configured in Step 16, on page 10.

c) Click **Submit**.

d) Navigate to the **Create Spiue Switch Policy Group** page:

Admin > **Fabric** > **Fabric Policies** > **Switches** > **Spine Switches** > **Policy Groups**, then right-click and choose **Create Spine Switch Policy Group**.

e) In the **Create Spine Switch Policy Group** page, enter the necessary information, specifically the following fields:

- **Analytics Policy**: Select the analytics policy that you created in the previous step.
- **Node Control Policy**: Select the `default` fabric node control policy that you configured in Step 16, on page 10.

f) Click **Submit**.

Step 19 Create the leaf and spine switch profiles.

a) In the APIC interface, navigate to the **Create Leaf Switch Profile** page:

Admin > **Fabric** > **Fabric Policies** > **Switches** > **Leaf Switches** > **Profiles**, then right-click and choose **Create Leaf Switch Profile**.

b) On the **Create Leaf Switch Profile** page, enter the necessary information, specifically the following fields:

- **Switch Associations**: Select the leaf switches and associate the leaf switch policy group that you created in the previous step.

c) Click **Submit**.

d) Navigate to the **Create Spine Switch Profile** page:

Admin > **Fabric** > **Fabric Policies** > **Switches** > **Spine Switches** > **Profiles**, then right-click and choose **Create Spine Switch Profile**.

e) On the **Create Spine Switch Profile** page, enter the necessary information, specifically the following fields:

- **Switch Associations**: Select the spine switches and associate the spine switch policy group that you created in the previous step.

f) Click **Submit**.

Step 20 Verify the configurations were set correctly.

a) Log into the APIC CLI and enter the following:

```
apic1# fabric 101 show flow monitor
------------------------------------------------------------------
 Node 101 (Leaf1)
------------------------------------------------------------------

Feature Prio: Analytics
```

b) Log into the leaf switch and enter the following:

```
Leaf1# ps -ef | grep ta_agent
root     19200 18286  0 04:42 pts/0    00:00:00 /usr/local/bin/node
/bootflash/tetration/ta_agent/ta_agent.js
admin    33433 32405  0 04:44 pts/2    00:00:00 grep ta_agent


Leaf1# cd /.aci/mitfs/sys/analytics/inst-analytics

Leaf1# cat summary
# Netflow Instance
mode           : analytics
adminSt        : enabled
childAction    :
ctrl           :
dn             : sys/analytics/inst-analytics
ipFiltAct      : deny
lcOwn          : local
modTs          : 2017-12-18T18:22:16.751+00:00
monPolDn       : uni/fabric/monfab-default
name           :
nwIssues       :
operErr        :
operSt         : enabled
operStQual     :
pltoperStQual  :
policyDn       : uni/fabric/analytics/cluster-<cluster_name>/cfgsrv-<analytics_policy_name>
rn             : inst-analytics
status         :


Leaf1# cd controller-<cluster_name>_<analytics_policy_name>

Leaf1# cat summary
# Controller Reachability
name               : <cluster_name>_<analytics_policy_name>
InstallOperSt      : success
InstallOperStQual  : installed
childAction        :
descr              :
dn               : sys/analytics/inst-analytics/controller-<cluster_name>_<analytics_policy_name>
dscp               : CS4
dstAddr            : 10.28.121.132
dstPort            : 5640
imageUri           : http://10.0.0.1:7777/fwrepo/aci-analyticsagent-dk9.default.bin
imageUri2          : https://10.0.0.1/fwrepo/aci-analyticsagent-dk9.default.bin
imageVer           : 2.2.1.31
```

```
lcOwn              : local
modTs              : 2017-12-18T18:22:18.432+00:00
monPolDn           : uni/fabric/monfab-default
nameAlias          :
rn                 : controller-<cluster_name>_<analytics_policy_name>
srcAddr            : 192.5100.100.24/24
srcIf              : unspecified
status             :
uid                : 0
vrfName            : mgmt:inb
```

Step 21 Verify the configuration.

a) In the telemetry collector, click the **Action** button, choose **Agent Config**, then click the **Hardware Agent Conig** tab.

The leaf and spine switches should be displayed in this screen.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
CiscoSystems(USA)Pte.Ltd.
Singapore

Europe Headquarters
CiscoSystemsInternationalBV
Amsterdam,TheNetherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.