Cisco APIC and NetFlow

Revised: October 19, 2023

New and Changed Information

The following table provides an overview of the significant changes up to the current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Table 1: New Features and Changed Behavior

Cisco APIC Release	Feature	Description
5.2(8)	NetFlow Scale	The NetFlow scale increased. For more information, see NetFlow Scale, on page 4.
5.1(1)	NetFlow Exporter Policies	You can now associate a Layer 3 EPG from the in-band management tenant with a NetFlow exporter.
4.0(1)	Remote Leaf Switches	NetFlow is now supported on remote leaf switches.
2.3(1)	FX-platform Switches	NetFlow is now supported on the FXplatform switches.
2.2(1)	Cisco APIC and NetFlow.	This guide is first released.

About NetFlow

The NetFlow technology provides the metering base for a key set of applications, including network traffic accounting, usage-based network billing, network planning, as well as denial of services monitoring, network monitoring, outbound marketing, and data mining for both service providers and enterprise customers. Cisco provides a set of NetFlow applications to collect NetFlow export data, perform data volume reduction, perform post-processing, and provide end-user applications with easy access to NetFlow data. If you have enabled NetFlow monitoring of the traffic flowing through your datacenters, this feature enables you to perform the same level of monitoring of the traffic flowing through the Cisco Application Centric Infrastructure (Cisco ACI) fabric.

Instead of hardware directly exporting the records to a collector, the records are processed in the supervisor engine and are exported to standard NetFlow collectors in the required format.

For information about configuring NetFlow with virtual machine networking, see the *Cisco ACI Virtualization Guide*.

NetFlow Monitor Policies

NetFlow policies can be deployed on a per-interface basis. Depending on the traffic-type or address family to be monitored (IPv4, IPv6, or Layer 2), you can enable different NetFlow monitor policies. A monitor policy (`netflowMonitorPol`) acts as a container to hold relationships to the record policy and exporter policy. A monitor policy identifies packet flows for ingress IP packets and provides statistics based on these packet flows. NetFlow does not require any change to either the packets themselves or to any networking device.

This policy can be configured under Fabric for deployment on physical interfaces or for a Tenant to be applied to bridge domains and L3Outs.

NetFlow can be deployed on the entire fabric or on a portion of the fabric to monitor packet statistics of different interface types.

NetFlow statistics are collected on the ingress packet prior to any policy enforcement. NetFlow statistics are recorded even if the packet is not permitted by policy (contract).

NetFlow Record Policies

A record policy (`netflowRecordPol`) lets you define a flow and what statistics to collect for each flow. This is achieved by defining the keys that NetFlow uses to identify packets in the flow as well as other fields of interest that NetFlow gathers for the flow. You can define a flow record with any combination of keys and fields of interest. A flow record also defines the types of counters gathered per flow, and you can configure 32-bit or 64-bit packet or byte counters.

A record policy has the following properties:

- `RecordPol.match`—A flow can be defined using the `match` property, which can be a combination of the following values:
  - `src-ipv4`, `dst-ipv4`, `src-port`, `dst-port`, `proto`, `vlan`, `tos`
  - `src-ipv6`, `dst-ipv6`, `src-port`, `dst-port`, `proto`, `vlan`, `tos`
  - `ethertype`, `src-mac`, `dst-mac`, `vlan`
  - `src-ip`, `dst-ip`, `src-port`, `dst-port`, `proto`, `vlan`, `tos`

  

  Note The `src-ip` and `dst-ip` parameters qualify both IPv4 and IPv6.

- `RecordPol.collect`—The `collect` property can be used to specify what information to collect for a given flow.

NetFlow Exporter Policies

An exporter policy (`netflowExporterPol`) specifies where the data collected for a flow must be sent. A NetFlow collector is an external entity that supports the standard NetFlow protocol and accepts packets marked with valid NetFlow headers.

An exporter policy has the following properties:

- **Destination IP Address**: This mandatory property specifies the IPv4 or IPv6 address of the NetFlow exporter that accepts the NetFlow flow packets. This must be in the host format (that is, `/32` or `/128`).
- **Destination Port**: This mandatory property specifies the port on which the exporter application is listening on, which enables the exporter to accept incoming connections.
- **Source IP Address Type**: This property populates the source IP address in the NetFlow record packets sent from the switch. The source IP address is populated in the NetFlow packets based on one of the configuration options below. No switch interface will be configured with this address.
  - **Custom Src IP**: When the source IP address type is **Custom Src IP**, the property is used similar to a tag to distinguish flows from different sections or nodes in the fabric. The address will be a prefix with at least 12 host bits. That is, the mask must be less than or equal to 20 for IPv4, or less than or equal to 116 for IPv6. The switch uses the configured prefix and host bits to populate the source IP address in the Netflow packet. The host portion will be equal to the node-id of the leaf sending the packet.
  - **Inband Management IP**: The source IP address in the NetFlow packets will be the configured switch inband management IP address.

- **OutOfband Management IP**: The source IP address in the NetFlow packets will be the configured switch out-of-band management IP address.
- **PTEP address**: The source IP address in the NetFlow packets will be the physical TEP (tunnel endpoint) address of the leaf switch.

Note You can use the "show flow exporter" leaf switch CLI command to display the source IP address for NetFlow records sent by that switch.

- **Version**: This property is used to specify the NetFlow version for the exporter to understand the packet. The only supported value is `v9`.
- **EPG Type**: App EPG or Layer 3 EPG

A NetFlow exporter can send data to a NetFlow collector directly connected to the fabric via an EPG or a remote collector reachable via an L3Oout. Choose the EPG type accordingly and complete the associated tenant/EPG as required.

The switch will send NetFlow packets from the VRF instance associated to the selected EPG or Layer 3 EPG. The VRF instance associated to the EPG or Layer EPG must be present on all leaf switches where NetFlow monitors are configured.

Beginning in the 5.1(1) release, you can associate an EPG or L3Out from the in-band VRF instance under the management tenant with a NetFlow exporter.

About NetFlow Node Policies

A node policy (netflowNodePol) deploys NetFlow timers that specify the rate at which flow records are sent to the external exporter. The timers are as follows:

- Collection interval—The time interval after which the leaf switch sends a NetFlow packet to the collector. The default value is 1 minute.
- Template interval—The time interval after which the leaf switch sends a record template to the collector. This template specifies the format of the records being sent to the collector. The default value 5 minutes.

NetFlow Scale

In the Cisco Application Policy Infrastructure Controller (APIC) 5.2(7) and earlier releases, some of the scale numbers for configurable NetFlow options are as follows:

Table 2: NetFlow scale in the Cisco APIC 5.2(7) and earlier releases

Configurable Options	Scale
Number of records per collect interval	20,000[1]
NetFlow monitor policies under the bridge domains per leaf switch	100

[1] NetFlow can collect a much higher number of flow records per minute, but only 20,000 flows are qualified.

Beginning with the Cisco APIC 5.2(8) release, the scale numbers of these configurable NetFlow options are as follows:

Table 3: NetFlow scale beginning in the Cisco APIC 5.2(8) release

Configurable Options	Scale
Number of records per collect interval	1,000,000
NetFlow monitor policies under the bridge domains per leaf switch	350

We increased the NetFlow scale due to compliance and monitoring use cases in telco and enterprise data centers, which have more flows. There is also the containerization of applications, in which one application is delivered by hundreds of containers as opposed to a single appliance or a small number of virtual machines. The multitude of containers causes a large number of flows to go to the data center.

For more NetFlow scale information, see the Verified Scalability Guide for your Cisco APIC release:

https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html#Verified_Scalability_Guides

NetFlow Deployment Considerations

Keep in mind the following NetFlow deployment considerations:

- Change the NetFlow node policy MTU to 9000 as follows:

  1. In the Menu bar, go to **Fabric** > **Access Policies**
  2. In the Navigation pane, go to **Policies** > **Switch** > **NetFlow Node** > **default**.
  3. In the Work pane, set the MTU property to 9000.

  This value ensures that more flow records get exported into a single packet, thereby reducing the CPU utilization when records are exported to the Netflow collector. This is especially important when you deploy NetFlow at a high scale. For example, with 1 million flow records, the CPU utilization can spike to 100% for a few seconds in each collection interval (during export to the external collector) when the NetFlow MTU is set to 1500. In contrast, those CPU utilization spikes should remain under 80% and with reduced duration after you set the NetFlow MTU to 9000.

- NetFlow captures flows in a hardware table based on a hash, then exports these flows from the hardware table at a regular interval to a software cache that is built in a switch's CPU. Because NetFlow captures flow records based on the hash, there can be flow collisions that prevent NetFlow from capturing a flow even though there is space in the flow table. Due to this, NetFlow might not capture all flows.
- The capacity of the hardware table is a fraction of the software cache. Because of this:
  - Some flows can be missed not only because of the hash collisions in the hardware table, but also because each update from the hardware table into the software cache might contain flows that are duplicated (the flows already exist in the software cache).
  - The granularity of the flow duration carried in the flow record degrades as the total number of flows increases.
- The software cache in a switch's CPU can hold up to 1 million IPv6 or IPv4 flow records. However, in the hardware table, an IPv6 flow record requires double the space as an IPv4 flow record.

NetFlow Support and Limitations

The following list provides information about the available support for NetFlow and the limitations of that support:

- EX, FX, FX2, and newer switches support NetFlow. For a full list of switch models supported on a specific release, see the *Cisco Nexus 9000 ACI-Mode Switches Release Notes* for that release.
- NetFlow on remote leaf switches is supported starting with Cisco Application Policy Infrastructure Controller (APIC) release 4.0(1).
- Cisco Application Centric Infrastructure (ACI) supports only ingress and not egress NetFlow. On a bridge domain, NetFlow cannot reliably capture packets entering from a spine switch.
- Spine switches do not support NetFlow, and tenant-level information cannot be derived locally from the packet on the spine switch.
- The hardware does not support any active/inactive timers. The flow table records get aggregated as the table gets flushed, and the records get exported every minute.
- At every export interval, the software cache gets flushed and the records that are exported in the next interval will have a reset packet/byte count and other statistics, even if the flow was long-lived.
- The filter TCAM has no labels for bridge domain or interfaces. If you add a NetFlow monitor to two bridge domains, the NetFlow monitor uses two rules for IPv4, or eight rules for IPv6. As such, the scale is limited with the 1K filter TCAM.
- ARP/ND are handled as IP packets and their target protocol addresses are put in the IP fields with some special protocol numbers from 249 through 255 as protocol ranges. NetFlow collectors might not understand this handling.
- The ICMP checksum is part of the Layer 4 src port in the flow record, so for ICMP records, many flow entries will be created if this is not masked, as is similar for other non-TCP/UDP packets.
- Cisco ACI-mode switches support only two active exporters.
- NetFlow traffic from leaf switches sometimes is unable to reach the collector due to the switch being unable to perform inter-VRF instance routing of the CPU-generated packet. As a workaround, create a fake static path for the EPG that is already configured under the same VRF instance as the L3Out that is used for the NetFlow collector. The fake path enables the traffic to reach the collector.
- When configuring a NetFlow exporters policy in mixed mode, you can configure a subnet for a specific VRF instance. Flow Telemetry will track all tenants that are associated with the EPG. You do not need to configure a separate policy for each subnet.

  For example, if you specify `0.0.0.0/0` as the subnet for the `t1:ctx2` VRF instance, Flow Telemetry tracks all IPv4 flows irrespective with which VRF instance they are associated.
- When a NetFlow exporter endpoint is behind a bridge domain, you must enable the unicast routing knob on the bridge domain to install the URIB routes for the bridge domain subnet. If the knob is disabled, then packets will not be forwarded to the collector and the operSt will be disabled for the collector policy.
- You cannot enable NetFlow and Flow Telemetry simultaneously.

NetFlow on EX Platform Switches

In addition to the generic support information, the following limitations apply to EX platform switches:

- NetFlow can be supported on a bridge domain; however, NetFlow cannot distinguish between bridged and routed packets. If you configure NetFlow on an interface VLAN (SVI) to capture only routed packets, NetFlow cannot limit collection to this type in EX switches.
- EX switches cannot provide an encapsulation VLAN in the flow record.
- EX switches do not have a MAC address packet classify feature, so the configuration engine flow record will contain only non-IP address flows (ARP is already treated as IP).
- EX switches do not support regularly-deployed and understood NetFlow sampling, such as packet-based sampling (M out of N).
- Having a type of service or source interface as part of the flow hash is not supported. Source interface information is collected in the record, but no type of service information is collected in EX switches.
- EX switches have fixed flow collection parameters.
- EX switches support only two flow records of each type. The exception is that four configuration engine flow records are supported.
- EX switches assign the following protocol numbers to identify the ARP and ND packets:
    - ARP Req 249
    - ARP Res 250
    - RARP Req 247
    - RARP Res 248
    - Nd Sol 249
    - Nd Adv 250

  All other ARP and ND packets are set to 255.

NetFlow Supported Interfaces

The following interfaces are supported for NetFlow:

- Physical Ethernet (Layer 2 and Layer 3)
- Port channel (PC)
- Virtual port channel (vPC)
- Fabric Extenders (FEX), FEX PC, and FEX VPC
- Layer 3 sub-interface
- SVI
- Bridge domains

Unlike other interface policies, NetFlow policies are not applied by default on interfaces. NetFlow must be explicitly enabled on a given interface.

For each interface, the address family (or filter) must be specified while enabling NetFlow monitoring. The address family can be one of the following types:

- IPv4
- IPv6
- CE (classical ethernet/Layer 2)

The address family causes the hardware to monitor packets only based on the address family that is provided. Different monitoring policies can be enabled per address family on the same interface.

NetFlow and Cisco Tetration Analytics Priority

As far the Cisco Application Centric Infrastructure (Cisco ACI) hardware is concerned, NetFlow and Cisco Tetration Analytics use the same ASIC building blocks to collect data. You cannot enable both features at the same time. NetFlow or Tetration Analytics must be explicitly enabled before configuring and deploying the related policies. The default is Tetration Analytics.

If the Cisco APIC pushes both Cisco Tetration Analytics and NetFlow configurations to a particular node, the chosen priority flag alerts the switch as to which feature should be given priority. The other feature's configuration is ignored.

Configuring NetFlow at the Fabric Level Using the GUI

Configuring a Fabric NetFlow Monitor Policy Using the GUI

The following procedure configures a fabric NetFlow monitor policy using the Cisco APIC GUI.

Procedure

Step 1 From the menu bar, choose **Fabric** > **Access Policies**.

Step 2 In the **Navigation** pane, choose **Policies** > **Interface** > **NetFlow** > **NetFlow Monitors**.

Note In earlier releases, the NetFlow Monitor policy configuration may be located under **Interface Policies** > **Policies** > **Analytics** > **NetFlow Monitors** instead.

Step 3 Right-click **NetFlow Monitors** and select **Create NetFlow Monitor**

Step 4 In the **Create NetFlow Monitor** dialog box, fill in the fields as required.

You can create new or add existing Flow Records and Exporters.

Creating **Associated Flow Record** is described in Configuring a Fabric NetFlow Record Policy Using the GUI, on page 8.

Creating **Associated Flow Exporters** is described in Configuring a Fabric NetFlow Exporter Policy Using the GUI, on page 9.

You can associate a maximum of two flow exporters with the monitor policy.

Configuring a Fabric NetFlow Record Policy Using the GUI

The following procedure configures a fabric NetFlow record policy using the Cisco APIC GUI.

Procedure

Step 1 From the menu bar, choose **Fabric** > **Access Policies**.

Step 2 In the **Navigation** pane, choose **Policies** > **Interface** > **NetFlow** > **NetFlow Records**.

Note In earlier releases, the NetFlow Record policy configuration may be located under **Interface Policies** > **Policies** > **Analytics** > **NetFlow Records** instead.

Step 3 Right-click **NetFlow Records** and choose **Create NetFlow Record**.

Step 4 In the **Create NetFlow Record** dialog box, fill in the fields as required, except as specified below:

a) For the **Collect Parameters** drop-down list, you can choose multiple parameters.

b) For the **Match Parameters** drop-down list, you can choose multiple parameters.

If you choose multiple parameters, your choices must be one of the following combinations or a subset of one of the combinations:

- Source IPv4, Destination IPv4, Source Port, Destination Port, IP Protocol, VLAN, IP TOS
- Source IPv6, Destination IPv6, Source Port, Destination Port, IP Protocol, VLAN, IP TOS
- Ethertype, Source MAC, Destination MAC, VLAN
- Source IP, Destination IP, Source Port, Destination Port, IP Protocol, VLAN, IP TOS, where Source IP/Destination IP qualifies both IPv4 and IPv6.

Configuring a Fabric NetFlow Exporter Policy Using the GUI

The following procedure configures a fabric NetFlow exporter policy using the Cisco APIC GUI.

Procedure

Step 1 From the menu bar, choose **Fabric** > **Access Policies**.

Step 2 In the **Navigation** pane, choose **Policies** > **Interface** > **NetFlow** > **NetFlow Exporters**.

Note In earlier releases, the NetFlow Monitor policy configuration may be located under **Interface Policies** > **Policies** > **Analytics** > **NetFlow Exporters** instead.

Step 3 Right-click **NetFlow Exporters** and choose **Create External Collector Reachability**

Step 4 In the **Create External Collector Reachability** dialog box, fill in the fields as required, except as specified below:

a) For the **NetFlow Exporter Version Format** buttons, **Version 9** is the only valid choice. Even if you click one of the other buttons, the version defaults to 9.

b) For the **EPG Type** check boxes, you can leave the boxes unchecked, or you can put a check in one box. You cannot put a check in multiple boxes.

Deploying NetFlow Monitor Policy Through a Selector Using Cisco APIC GUI

The following procedure deploys a NetFlow monitor policy through a selector using the Cisco APIC GUI.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Access Policies**.

Step 2 In the **Navigation** pane, choose **Interfaces** > **Leaf Interfaces** > **Policy Groups**.

In earlier releases, the configuration may be located under **Interface Policies** > **Policy Groups** > **Leaf Policy Groups** instead.

Step 3 You can deploy the NetFlow monitor policy when you create a new leaf policy group, or you can deploy the NetFlow monitor policy on an existing leaf policy group.

To deploy the NetFlow monitor policy when you create a new leaf policy group, use the following steps:

a) Right-click the type of interface group you want to create and choose **Create Leaf Access Port Policy Group**.
b) In the dialog box, fill in the fields as required

On the **NetFlow Monitor Policies** table, click **+** to add a policy, and choose the IP filter type and monitor policy.

To deploy the NetFlow monitor policy on an existing leaf policy group, use the following steps:

a) In the **Navigation** pane, choose one of the existing leaf access port policy groups, PC interface policy groups, or VPC interface policy groups.
b) In the **Work** pane, on the **NetFlow Monitor Policies** table, click **+** to add a policy, and choose the IP filter type and monitor policy.
c) Click **Submit**.

Configuring the Telemetry Method Using Cisco APIC GUI

This procedure uses the Cisco APIC GUI to create a fabric node profile and specify the telemetry method, then associate the fabric node profile to a fabric policy group.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Fabric Policies**.

Step 2 In the **Navigation** pane, right-click **Policies** > **Monitoring** > **Fabric Node Controls** and choose **Create Fabric Node Control**.

In earlier releases, in the **Navigation** pane, you instead directly right-click **Fabric Node Controls** (without needing to expand any navigation branches) and choose **Create Fabric Node Control**.

Step 3 In the **Create Fabric Node Control** dialog, fill in the fields as required, except as specified below:

a) For **Feature Selection**, choose a telemetry method:

- **Analytics Priority**: Specifies Cisco Tetration Analytics. This is the default value in the 3.1(1) release and earlier.
- **NetFlow Priority**: Specifies NetFlow.

- **Telemetry Priority**: Specifies Cisco Nexus Dashboard Insights flow telemetry. This became available beginning in the 3.1(2) release and became the default value in that release.

Step 4 Click **Submit**.

Step 5 Associate the fabric node control policy to the appropriate fabric policy group and profile. For example, to associate the fabric node control policy to a leaf switch policy group, perform the following substeps:

a) In the **Navigation** pane, choose **Switches** > **Leaf Switches** > **Policy Groups** > ***policy_group_name***.

b) In the **Work** pane, for **Node Control Policy**, choose the fabric node control policy that you created.

Configuring NetFlow at the Tenant Level Using the GUI

Configuring a Tenant NetFlow Monitor Policy Using the GUI

The following procedure configures a tenant NetFlow monitor policy using the Cisco APIC GUI.

Procedure

Step 1 From the menu bar, choose **Tenants** > **All Tenants**.

Step 2 In the **Work** pane, double-click the tenant's name.

Step 3 In the **Navigation** pane, choose **Tenant <*tenant-name*>** > **Policies** > **NetFlow** > **NetFlow Monitors**.

Note In earlier releases, the NetFlow Monitor policy configuration may be located under **Tenant <*tenant-name*>** > **Application Profiles** > **<*application-profile-name*>** instead.

Step 4 Right-click **NetFlow Monitors** and choose **Create NetFlow Monitor**.

Step 5 In the **Create NetFlow Monitor** dialog box, fill in the fields as required.

You can create new or add existing Flow Records and Exporters.

Creating **Associated Flow Record** is described in Configuring a Tenant NetFlow Record Policy Using the GUI, on page 11.

Creating **Associated Flow Exporters** is described in Configuring a Tenant NetFlow Exporter Policy Using the GUI, on page 12.

You can associate a maximum of two flow exporters with the monitor policy.

Configuring a Tenant NetFlow Record Policy Using the GUI

The following procedure configures a tenant NetFlow record policy using the Cisco APIC GUI.

Procedure

Step 1 From the menu bar, choose **Tenants** > **All Tenants**.

Step 2 In the Work pane, double-click the tenant's name.

Step 3 In the Navigation pane, choose **Tenant <*tenant-name*>** > **Policies** > **NetFlow** > **NetFlow Records**.

Note In earlier releases, the NetFlow Exporter policy configuration may be located under **Tenant <*tenant-name*>** > **Analytics** > **NetFlow Records** instead.

Step 4 Right-click **NetFlow Records** and choose **Create Flow Record**.

Step 5 In the **Create NetFlow Record** dialog box, fill in the fields as required, except as specified below:

a) For the **Collect Parameters** drop-down list, you can choose multiple parameters.

b) For the **Match Parameters** drop-down list, you can choose multiple parameters.

If you choose multiple parameters, your choices must be one of the following combinations or a subset of one of the combinations:

- Source IPv4, Destination IPv4, Source Port, Destination Port, IP Protocol, VLAN, IP TOS
- Source IPv6, Destination IPv6, Source Port, Destination Port, IP Protocol, VLAN, IP TOS
- Ethertype, Source MAC, Destination MAC, VLAN
- Source IP, Destination IP, Source Port, Destination Port, IP Protocol, VLAN, IP TOS, where Source IP/Destination IP qualifies both IPv4 and IPv6.

Configuring a Tenant NetFlow Exporter Policy Using the GUI

The following procedure configures a tenant NetFlow exporter policy using the Cisco APIC GUI.

Procedure

Step 1 From the menu bar, choose **Tenants** > **All Tenants**.

Step 2 In the Work pane, double-click the tenant's name.

Step 3 In the Navigation pane, choose **Tenant <*tenant-name*>** > **Policies** > **NetFlow** > **NetFlow Exporters**.

Note In earlier releases, the NetFlow Exporter policy configuration may be located under **Tenant <*tenant-name*>** > **Analytics** > **NetFlow Exporters** instead.

Step 4 Right-click **NetFlow Exporters** and choose **Create External Collector Reachability**.

Step 5 In the **Create External Collector Reachability** dialog box, fill in the fields as required, except as specified below:

a) For the **NetFlow Exporter Version Format** buttons, **Version 9** is the only supported choice.

b) For the **EPG Type** check boxes, you can leave the boxes unchecked, or you can put a check in one box. You cannot put a check in multiple boxes.

Deploying NetFlow Monitor Policy Through an L3Out Using Cisco APIC GUI

The following procedure deploys a NetFlow monitor policy through an L3Out using the Cisco APIC GUI.

Procedure

Step 1 From the menu bar, choose **Tenants** > **All Tenants**.

Step 2 In the Work pane, double-click the tenant's name.

Step 3 In the Navigation pane, choose **Tenant <*tenant-name*>** > **Networking** > **External Routed Networks** > **<*network-name*>** > **Logical Node Profiles** > **<*node-profile-name*>** > **Logical Interface Profile** > **<*interface-profile-name*>**.

Step 4 Select the **General** tab

Step 5 Under **NetFlow Monitor Policies**, click **+** to add a NetFlow policy.

Step 6 Click **Update** to add the NetFlow policy.

Deploying NetFlow Monitor Policy Through a Bridge Domain Using Cisco APIC GUI

The following procedure deploys a NetFlow monitor policy through a bridge domain using Cisco APIC GUI.

Procedure

Step 1 On the menu bar, choose **Tenants** > **All Tenants**.

Step 2 In the Work pane, double-click the tenant's name.

Step 3 In the Navigation pane, choose **Tenant *tenant_name*** > **Networking** > **Bridge Domains**.

Step 4 You can deploy the NetFlow monitor policy when you create a new bridge domain, or you can deploy the NetFlow monitor policy on an existing bridge domain.

To deploy the NetFlow monitor policy when you create a new bridge domain, use the following steps:

a) In the **Work** pane, choose **Actions** > **Create Bridge Domain**.

b) In the **Create Bridge Domain** dialog box, fill in the fields as required, except as specified below:

1. On the **Advanced Troubleshooting** step, on the **NetFlow Monitor Policies** table, click **+**, choose a NetFlow IP filter type, choose a NetFlow monitor policy, and click **Update**.
2. Click **Finish**.

To deploy the NetFlow monitor policy on an existing bridge domain, use the following steps:

a) In the **Navigation** pane, choose one of the existing bridge domains.

b) In the **Work** pane, choose **Policy** > **Advanced Troubleshooting**.

c) On the **NetFlow Monitor Policies** table, click **+**, choose a NetFlow IP filter type, choose a NetFlow monitor policy, and click **Update**.

d) Click **Submit**.

Configuring the Telemetry Method Using Cisco APIC GUI

This procedure uses the Cisco APIC GUI to create a fabric node profile and specify the telemetry method, then associate the fabric node profile to a fabric policy group.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Fabric Policies**.

Step 2 In the **Navigation** pane, right-click **Policies** > **Monitoring** > **Fabric Node Controls** and choose **Create Fabric Node Control**.

In earlier releases, in the **Navigation** pane, you instead directly right-click **Fabric Node Controls** (without needing to expand any navigation branches) and choose **Create Fabric Node Control**.

Step 3 In the **Create Fabric Node Control** dialog, fill in the fields as required, except as specified below:

a) For **Feature Selection**, choose a telemetry method:

- **Analytics Priority**: Specifies Cisco Tetration Analytics. This is the default value in the 3.1(1) release and earlier.
- **NetFlow Priority**: Specifies NetFlow.
- **Telemetry Priority**: Specifies Cisco Nexus Dashboard Insights flow telemetry. This became available beginning in the 3.1(2) release and became the default value in that release.

Step 4 Click **Submit**.

Step 5 Associate the fabric node control policy to the appropriate fabric policy group and profile. For example, to associate the fabric node control policy to a leaf switch policy group, perform the following substeps:

a) In the **Navigation** pane, choose **Switches** > **Leaf Switches** > **Policy Groups** > ***policy_group_name***.

b) In the **Work** pane, for **Node Control Policy**, choose the fabric node control policy that you created.

Configuring NetFlow Using the NX-OS-Style CLI

Configuring NetFlow Node Policy Using the NX-OS-Style CLI

The following example procedure uses the NX-OS-style CLI to configure a NetFlow node policy:

Procedure

Step 1 Enter the configuration mode.

Example:

```
apic1# config
```

Step 2 Configure the node policy.

Example:

```
apic1(config)# flow node-policy nodePol
apic1(config-flow-node-pol)# flow timeout collection 100
apic1(config-flow-node-pol)# flow timeout template 123
apic1(config-flow-node-pol)# exit
```

Configuring NetFlow Infra Selectors Using the NX-OS-Style CLI

You can use the NX-OS-style CLI to configure NetFlow infra selectors. The infra selectors are used for attaching a Netflow monitor to a PHY, port channel, virtual port channel, fabric extender (FEX), or port channel fabric extender (FEXPC) interface.

The following example CLI commands show how to configure NetFlow infra selectors using the NX-OS-style CLI:

Procedure

Step 1 Enter the configuration mode.

Example:

```
apic1# config
```

Step 2 Create a NetFlow exporter policy.

Example:

In the following commands, the destination endpoint group is the endpoint group that the exporter sits behind. This endpoint group can also be an external Layer 3 endpoint group.

```
apic1(config)# flow exporter infraExporter1 destination address 1.2.3.4 transpo udp 1234
apic1(config-flow-exporter)# destination epg tenant tn2 application ap2 epg epg2
apic1(config-flow-exporter)# vrf member tenant tn2 vrf vrf2
apic1(config-flow-exporter)# version v9
apic1(config-flow-exporter)# source address 1.1.1.1
apic1(config-flow-exporter)# exit
```

Step 3 Create a second NetFlow exporter policy.

Example:

In the following commands, the destination endpoint group is the endpoint group that the exporter sits behind, which in this case is an external Layer 3 endpoint group.

```
apic1(config)# flow exporter infraExporter2
apic1(config-flow-exporter)# transport udp 9990
apic1(config-flow-exporter)# destination address 2001:db5:a0c:1f0::2
apic1(config-flow-exporter)# destination external-l3 epg tenant tn2 vrf v2 epg accounting-inst
apic1(config-flow-exporter)# vrf member tenant tn2 vrf vrf2
apic1(config-flow-exporter)# version v5
apic1(config-flow-exporter)# source address 2001:db8:a0b:12f0::1
apic1(config-flow-exporter)# exit
```

Step 4 Create a NetFlow record policy.

Example:

```
apic1(config)# flow record infraRecord1
apic1(config-flow-record)# match dst-ip
apic1(config-flow-record)# match dst-ipv4
apic1(config-flow-record)# match dst-ipv6
apic1(config-flow-record)# match dst-mac
apic1(config-flow-record)# match dst-port
apic1(config-flow-record)# match ethertype
apic1(config-flow-record)# match proto
apic1(config-flow-record)# match src-ip
apic1(config-flow-record)# match src-ipv4
apic1(config-flow-record)# match src-ipv6
apic1(config-flow-record)# match src-mac
apic1(config-flow-record)# match src-port
```

```
apic1(config-flow-record)# match tos
apic1(config-flow-record)# match vlan
apic1(config-flow-record)# collect count-bytes
apic1(config-flow-record)# collect count-pkts
apic1(config-flow-record)# collect pkt-disp
apic1(config-flow-record)# collect sampler-id
apic1(config-flow-record)# collect src-intf
apic1(config-flow-record)# collect tcp-flags
apic1(config-flow-record)# collect ts-first
apic1(config-flow-record)# collect ts-recent
apic1(config-flow-record)# exit
```

Step 5 Create a NetFlow monitor policy.

Example:

```
apic1(config)# flow monitor infraMonitor1
apic1(config-flow-monitor)# record infraRecord1
apic1(config-flow-monitor)# exporter infraExporter1
apic1(config-flow-monitor)# exporter infraExporter2
apic1(config-flow-monitor)# exit
```

You can attach a maximum of two exporters.

Step 6 Create an interface policy group (AccPortGrp).

Example:

```
apic1(config)# template policy-group pg1
apic1(config-pol-grp-if)# ip flow monitor infraMonitor1
apic1(config-pol-grp-if)# ipv6 flow monitor infraMonitor2
apic1(config-pol-grp-if)# exit
```

You can have one monitor policy per address family (IPv4 and IPv6).

Step 7 Create a node profile and infra selectors.

Example:

```
apic1(config)# leaf-profile lp1
apic1(config-leaf-profile)# leaf-group lg1
apic1(config-leaf-group)# leaf 101
apic1(config-leaf-profile)# exit
apic1(config)# leaf-interface-profile lip1
apic1(config-leaf-if-profile)# exit
apic1(config)# leaf-interface-profile lip1
apic1(config-leaf-if-profile)# leaf-interface-group lig1
apic1(config-leaf-if-group)# interface ethernet 1/5
apic1(config-leaf-if-profile)# policy-group pg1
apic1(config-leaf-if-profile)# exit
apic1(config-leaf-profile)# exit
```

Step 8 Create a port channel policy group (AccBndlGrp).

Example:

```
apic1(config)# template port-channel po6
apic1(config-if)# ip flow monitor infraMonitor1
apic1(config-if)# ipv6 flow monitor infraMonitor1
apic1(config-if)# exit
apic1(config-leaf-profile)# leaf-profile lp2
apic1(config-leaf-group)# leaf-group lg2
apic1(config-leaf-profile)# leaf 101
apic1(config-leaf-profile)# exit
apic1(config)# leaf-interface-profile lip2
```

```
apic1(config-leaf-if-profile)# exit
apic1(config)# leaf-interface-profile lip2
apic1(config-leaf-if-profile)# leaf-interface-group lig2
apic1(config-leaf-if-group)# interface ethernet 1/6
apic1(config-leaf-if-profile)# channel-group po6
apic1(config-leaf-if-profile)# exit
```

You can have one monitor policy per address family (IPv4 and IPv6). The interfaces can also be vPCs.

Configuring NetFlow Overrides Using the NX-OS-Style CLI

The following procudure configures NetFlow overrides using the NX-OS-Style CLI:

Procedure

Step 1 Enter the configuration mode.

Example:

```
apic1# config
```

Step 2 Create the override.

Example:

```
apic1(config)# leaf 101
apic1(config-leaf)# vrf context tenant tn2 vrf vrf2
apic1(config-leaf)# exit
apic1(config)# interface ethernet 1/15
apic1(config-if)# ip flow monitor infraMonitor1
apic1(config-if)# ipv6 flow monitor infraMonitor2
apic1(config-if)# exit
apic1(config)# exit
apic1# exit
```

You can have one monitor policy per address family (IPv4 and IPv6). The interfaces can also be vPCs.

Configuring NetFlow Tenant Hierarchy Using the NX-OS-Style CLI

The following example procedure uses the NX-OS-style CLI to configure the NetFlow tenant hierarchy:

Procedure

Step 1 Enter the configuration mode.

Example:

```
apic1# config
```

Step 2 Create a tenant and bridge domain, and add them to a VRF.

Example:

```
apic1(config)# tenant tn2
apic1(config-tenant)# vrf context vrf2
apic1(config-tenant-vrf)# exit
apic1(config-tenant)# bridge-domain bd2
apic1(config-tenant-bridge-domain)# vrf member vrf2
apic1(config-tenant-bridge-domain)# exit
apic1(config-tenant)# bridge-domain bd3
apic1(config-tenant-bridge-domain)# vrf member vrf2
apic1(config-tenant-bridge-domain)# exit
```

Step 3 Create an application endpoint group behind which the exporter resides.

Example:

```
apic1(config-tenant)# application ap2
apic1(config-tenant-app)# epg epg2
apic1(config-tenant-app)# bridge-domain member bd2
apic1(config-tenant-app-bridge-domain)# exit
apic1(config-tenant-app)# exit
```

Step 4 Create a second application endpoint group behind which the exporter resides.

Example:

```
apic1(config-tenant)# application ap3
apic1(config-tenant-app)# epg epg3
apic1(config-tenant-app)# bridge-domain member bd3
apic1(config-tenant-app-bridge-domain)# exit
apic1(config-tenant-app)# exit
```

Step 5 Attach a NetFlow monitor policy on the bridge domains.

Example:

```
apic1(config)# interface bridge-domain bd2
apic1(config-if)# ipv6 flow monitor tnMonitor1
apic1(config-if)# ip flow monitor tnMonitor1
apic1(config-if)# layer2-switched flow monitor tnMonitor1
apic1(config-if)# exit
apic1(config)# interface bridge-domain bd3
apic1(config-if)# ipv6 flow monitor tnMonitor1
apic1(config-if)# ip flow monitor tnMonitor1
apic1(config-if)# exit
```

You can have one monitor policy per address family (IPv4 and IPv6). The interfaces can also be vPCs.

Step 6 Create the Netflow exporter policy.

Example:

In the following commands, the destination endpoint group is the endpoint group that the exporter sits behind. This endpoint group can also be an external Layer 3 endpoint group.

```
apic1(config)# flow exporter tnExporter1
apic1(config-flow-exporter)# transport udp 1234
apic1(config-flow-exporter)# destination address 2.2.2.2
apic1(config-flow-exporter)# destination epg tenant tn2 application ap2 epg epg2
apic1(config-flow-exporter)# vrf member tenant tn2 vrf vrf2
apic1(config-flow-exporter)# version v9
apic1(config-flow-exporter)# source address 1.1.1.1
apic1(config-flow-exporter)# exit
```

Step 7 Create a second Netflow exporter policy.

Example:

In the following commands, the destination endpoint group is the endpoint group that the exporter sits behind, which in this case is an external Layer 3 endpoint group.

```
apic1(config)# flow exporter tnExporter2
apic1(config-flow-exporter)# transport udp 9990
apic1(config-flow-exporter)# destination address 2001:db5:a0c:1f0::2
apic1(config-flow-exporter)# destination external-l3 epg tenant tn2 vrf v2 epg accounting-inst
apic1(config-flow-exporter)# vrf member tenant tn2 vrf vrf2
apic1(config-flow-exporter)# version v5
apic1(config-flow-exporter)# source address 2001:db8:a0b:12f0::1
apic1(config-flow-exporter)# exit
```

Step 8 Create a NetFlow record policy.

Example:

```
apic1(config)# flow record tnRecord1
apic1(config-flow-record)# match dst-ip
apic1(config-flow-record)# match dst-ipv4
apic1(config-flow-record)# match dst-ipv6
apic1(config-flow-record)# match dst-mac
apic1(config-flow-record)# match dst-port
apic1(config-flow-record)# match ethertype
apic1(config-flow-record)# match proto
apic1(config-flow-record)# match src-ip
apic1(config-flow-record)# match src-ipv4
apic1(config-flow-record)# match src-ipv6
apic1(config-flow-record)# match src-mac
apic1(config-flow-record)# match src-port
apic1(config-flow-record)# match tos
apic1(config-flow-record)# match vlan
apic1(config-flow-record)# collect count-bytes
apic1(config-flow-record)# collect count-pkts
apic1(config-flow-record)# collect pkt-disp
apic1(config-flow-record)# collect sampler-id
apic1(config-flow-record)# collect src-intf
apic1(config-flow-record)# collect tcp-flags
apic1(config-flow-record)# collect ts-first
apic1(config-flow-record)# collect ts-recent
apic1(config-flow-record)# exit
```

Step 9 Create a NetFlow monitor policy.

Example:

```
apic1(config)# flow monitor tnMonitor1
apic1(config-flow-monitor)# record tnRecord1
apic1(config-flow-monitor)# exporter tnExporter1
apic1(config-flow-monitor)# exporter tnExporter2
apic1(config-flow-monitor)# exit
```

You can attach a maximum of two exporters.

Step 10 Add VLANs to the VLAN domain and configure a VRF for a leaf node.

Example:

```
apic1(config)# vlan-domain dom1
apic1(config-vlan)# vlan 5-100
apic1(config-vlan)# exit
apic1(config)# leaf 101
apic1(config-leaf)# vrf context tenant tn2 vrf vrf2
apic1(config-leaf-vrf)# exit
```

Step 11 Deploy an endpoint group on an interface to deploy the bridge domain.

Example:

```
apic1(config-leaf)# interface ethernet 1/10
apic1(config-leaf-if)# vlan-domain member dom1
apic1(config-leaf-if)# switchport trunk allowed vlan 10 tenant tn2 application ap2 epg epg2
apic1(config-leaf-if)# exit
```

Step 12 Deploy another endpoint group on an interface.

Example:

```
apic1(config-leaf)# interface ethernet 1/11
apic1(config-leaf-if)# vlan-domain member dom1
apic1(config-leaf-if)# switchport trunk allowed vlan 11 tenant tn2 application ap3 epg epg3
apic1(config-leaf-if)# exit
```

Step 13 Attach the monitor policy to the sub-interface.

Example:

```
apic1(config-leaf)# interface ethernet 1/20
apic1(config-leaf-if)# vlan-domain member dom1
apic1(config-leaf-if)# no switchport
apic1(config-leaf-if)# exit
apic1(config-leaf)# interface ethernet 1/20.20
apic1(config-leaf-if)# vrf member tenant tn2 vrf vrf2
apic1(config-leaf-if)# ipv6 address 20::1/64 preferred
apic1(config-leaf-if)# ipv6 flow monitor tnMonitor1
apic1(config-leaf-if)# ip flow monitor tnMonitor2
apic1(config-leaf-if)# exit
```

Step 14 Attach the monitor policy to a switched virtual interface (SVI).

Example:

```
apic1(config-leaf)# interface vlan 30
apic1(config-leaf-if)# vrf member tenant tn2 vrf vrf2
apic1(config-leaf-if)# ipv6 address 64::1/64 preferred
apic1(config-leaf-if)# ip flow monitor tnMonitor1
apic1(config-leaf-if)# ip6 flow monitor tnMonitor1
apic1(config-leaf-if)# exit
```

Step 15 Associate the SVI to a Layer 2 interface.

Example:

```
apic1(config-leaf)# interface ethernet 1/30
apic1(config-leaf-if)# vlan-domain member dom1
apic1(config-leaf-if)# switchport trunk allowed vlan 30 tenant tn2 external-svi
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
apic1(config)# exit
```

Configuring NetFlow and Tetration Analytics Feature Priority Through Node Control Policy Using NX-OS-Style CLI

The following example procedure uses the NX-OS-style CLI to configure the NetFlow and Tetration Analytics feature priority through a node control policy:

Procedure

Step 1 Enter the configuration mode.

Example:

```
apic1# config
```

Step 2 Create a node control policy.

Example:

```
apic1(config)# node-control policy pol1
```

Step 3 Set NetFlow as the priority feature.

Example:

```
apic1(config-node)# feature netflow
```

Step 4 Exit the node control policy configuration.

Example:

```
apic1(config-node)# end
```

Step 5 Deploy the policy to node 101 and node 102.

Example:

```
ifav-isim15-ifc1(config)# fabric-internal
ifav-isim15-ifc1(config-fabric-internal)# template leaf-policy-group lpg1
ifav-isim15-ifc1(config-leaf-policy-group)# inherit node-control-policy pol1
ifav-isim15-ifc1(config-leaf-policy-group)# exit
ifav-isim15-ifc1(config-fabric-internal)# leaf-profile leafProfile1
ifav-isim15-ifc1(config-leaf-profile)# leaf-group leafgrp1
ifav-isim15-ifc1(config-leaf-group)# leaf 101
ifav-isim15-ifc1(config-leaf-group)# leaf 102
ifav-isim15-ifc1(config-leaf-group)# leaf-policy-group lpg1
ifav-isim15-ifc1(config-leaf-group)# end
```

Verifying the NetFlow Configuration Using the NX-OS-Style CLI

The following procedure verifies the NetFlow configuration using the Cisco Application Policy Infrastructure Controller (Cisco APIC) NX-OS-Style CLI and the NX-OS CLI of a leaf switch:

Procedure

Step 1 In the Cisco APIC NX-OS-Style CLI, show the NetFlow monitor information for the infra tenant or the specified tenant, as appropriate:

show flow monitor {infra *policy_name* detail | tenant *tenant_name*}

Example:

```
apic1# show flow monitor infra default detail
```

Step 2 Using the CLI one of the leaf switches, run the following commands:

Example:

```
leaf# show flow exporter
leaf# show flow record
leaf# show flow monitor
leaf# show flow timers
leaf# show flow interface
leaf# show flow vlan
```

Configuring NetFlow Using the REST API

Configuring NetFlow Infra Selectors Using REST API

You can use the REST API to configure NetFlow infra selectors. The infra selectors are used for attaching a Netflow monitor to a PHY, port channel, virtual port channel, fabric extender (FEX), or port channel fabric extender (FEXPC) interface.

The following example XML shows how to configure NetFlow infra selectors using the REST API:

```
<infraInfra>
    <!--Create Monitor Policy /-->
    <netflowMonitorPol name='monitor_policy1' descr='This is a monitor policy.'>
        <netflowRsMonitorToRecord tnNetflowRecordPolName='record_policy1' />
        <!-- A Max of 2 exporters allowed per Monitor Policy /-->
        <netflowRsMonitorToExporter tnNetflowExporterPolName='exporter_policy1' />
        <netflowRsMonitorToExporter tnNetflowExporterPolName='exporter_policy2' />
    </netflowMonitorPol>

    <!--Create Record Policy /-->
    <netflowRecordPol name='record_policy1' descr='This is a record policy.' match='src-ipv4,src-port'/>

    <!--Create Exporter Policy /-->
    <netflowExporterPol name='exporter_policy1' dstAddr='10.10.1.1' srcAddr='10.10.1.10' ver='v9' descr='This
 is an exporter policy.'>
        <!--Exporter can be behind app EPG or external L3 EPG (InstP) /-->
        <netflowRsExporterToEPg tDn='uni/tn-t1/ap-app1/epg-epg1'/>
        <!--This Ctx needs to be the same Ctx that EPG1's BD is part of /-->
        <netflowRsExporterToCtx tDn='uni/tn-t1/ctx-ctx1'/>
    </netflowExporterPol>

    <!--Node-level Policy for collection Interval /-->
    <netflowNodePol name='node_policy1' collectIntvl='500' />

    <!-- Node Selectors - usual config /-->
    <infraNodeP name="infraNodeP-17" >
        <infraLeafS name="infraLeafS-17" type="range">
            <!-- NOTE: The nodes can also be fex nodes /-->
            <infraNodeBlk name="infraNodeBlk-17" from_="101" to_="101"/>
            <infraRsAccNodePGrp tDn='uni/infra/funcprof/accnodepgrp-nodePGrp1' />
        </infraLeafS>
        <infraRsAccPortP tDn="uni/infra/accportprof-infraAccPortP"/>
    </infraNodeP>

    <!-- Port Selectors - usual config /-->
    <infraAccPortP name="infraAccPortP" >
        <infraHPortS name="infraHPortS" type="range">
            <!-- NOTE: The interfaces can also be Port-channels, fex interfaces or fex PCs /-->
            <infraPortBlk name="infraPortBlk" fromCard="1" toCard="1" fromPort="8" toPort="8"/>
            <infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-infraAccPortGrp"/>
```

```
        </infraHPortS>
    </infraAccPortP>

    <!-- Policy Groups - usual config /-->
    <infraFuncP>
        <!-- Node Policy Group - to setup Netflow Node Policy /-->
        <infraAccNodePGrp name='nodePGrp1' >
            <infraRsNetflowNodePol tnNetflowNodePolName='node_policy1' />
        </infraAccNodePGrp>

        <!-- Access Port Policy Group - to setup Netflow Monitor Policy /-->
        <infraAccPortGrp name="infraAccPortGrp" >
            <!--One Monitor Policy per address family (ipv4, ipv6, ce) /-->
            <infraRsNetflowMonitorPol tnNetflowMonitorPolName='monitor_policy1' fltType='ipv4'/>
            <infraRsNetflowMonitorPol tnNetflowMonitorPolName='monitor_policy2' fltType='ipv6'/>
            <infraRsNetflowMonitorPol tnNetflowMonitorPolName=‘monitor_policy2' fltType=‘ce'/>
        </infraAccPortGrp>
    </infraFuncP>
</infraInfra>
```

Configuring NetFlow Tenant Hierarchy Using REST API

You can use the REST API to configure the NetFlow tenant hierarchy. The tenant hierarchy is used for attaching a NetFlow monitor to a bridge domain, Layer 3 sub-interface, or Layer 3 switched virtual interface (SVI).

The following example XML shows how to configure the NetFlow tenant hierarchy using the REST API:

```
<?xml version="1.0" encoding="UTF-8"?>

<!-- api/policymgr/mo/.xml -->
<polUni>
    <fvTenant name="t1">

        <!--Create Monitor Policy /-->
        <netflowMonitorPol name='monitor_policy1' descr='This is a monitor policy.'>
            <netflowRsMonitorToRecord tnNetflowRecordPolName='record_policy1' />
            <!-- A Max of 2 exporters allowed per Monitor Policy /-->
            <netflowRsMonitorToExporter tnNetflowExporterPolName='exporter_policy1' />
            <netflowRsMonitorToExporter tnNetflowExporterPolName='exporter_policy2' />
        </netflowMonitorPol>
        <!--Create Record Policy /-->
        <netflowRecordPol name='record_policy1' descr='This is a record policy.'/>

        <!--Create Exporter Policy /→
        <netflowExporterPol name='exporter_policy1' dstAddr='10.0.0.1' srcAddr='10.0.0.4'>

            <!--Exporter can be behind app EPG or external L3 EPG (InstP) /-->
            <netflowRsExporterToEPg tDn='uni/tn-t1/ap-app1/epg-epg2'/>
            <!--netflowRsExporterToEPg tDn='uni/tn-t1/out-out1/instP-accountingInst' /-->
            <!--This Ctx needs to be the same Ctx that EPG2’s BD is part of /-->
            <netflowRsExporterToCtx tDn='uni/tn-t1/ctx-ctx1' />
        </netflowExporterPol>

        <!--Create 2nd Exporter Policy /-->
        <netflowExporterPol name='exporter_policy2' dstAddr='11.0.0.1' srcAddr='11.0.0.4'>
            <netflowRsExporterToEPg tDn='uni/tn-t1/ap-app1/epg-epg2'/>
            <netflowRsExporterToCtx tDn='uni/tn-t1/ctx-ctx1' />
        </netflowExporterPol>

        <fvCtx name="ctx1" />

        <fvBD name="bd1" unkMacUcastAct="proxy" >
            <fvSubnet descr="" ip="11.0.0.0/24"\>
```

```
            <fvRsCtx tnFvCtxName="ctx1" />

            <!--One Monitor Policy per address family (ipv4, ipv6, ce) /-->
            <fvRsBDToNetflowMonitorPol tnNetflowMonitorPolName='monitor_policy1' fltType='ipv4'/>
            <fvRsBDToNetflowMonitorPol tnNetflowMonitorPolName='monitor_policy2' fltType='ipv6'/>
            <fvRsBDToNetflowMonitorPol tnNetflowMonitorPolName=‘monitor_policy2' fltType='ce'/>
        </fvBD>

        <!--Create App EPG /-->
        <fvAp name="app1">
            <fvAEPg name="epg2" >
                <fvRsBd tnFvBDName="bd1" />
                <fvRsPathAtt encap="vlan-20" instrImedcy="lazy" mode="regular"
tDn="topology/pod-1/paths-101/pathep-[eth1/20]"/>
            </fvAEPg>
        </fvAp>

        <!--L3 Netflow Config for sub-intf and SVI /-->
        <l3extOut name="out1">
            <l3extLNodeP name="lnodep1" >
                <l3extRsNodeL3OutAtt tDn="topology/pod-1/node-101" rtrId="1.2.3.4" />
                <l3extLIfP name='lifp1'>
                    <!--One Monitor Policy per address family (ipv4, ipv6, ce) /-->
                    <l3extRsLIfPToNetflowMonitorPol tnNetflowMonitorPolName='monitor_policy1' fltType='ipv4'
/>
                    <l3extRsLIfPToNetflowMonitorPol tnNetflowMonitorPolName='monitor_policy2' fltType='ipv6'
/>
                    <l3extRsLIfPToNetflowMonitorPol tnNetflowMonitorPolName=‘monitor_policy2' fltType=‘ce' />

                    <!--Sub-interface 1/40.40 on node 101 /-->
                  <l3extRsPathL3OutAtt tDn="topology/pod-1/paths-101/pathep-[eth1/40]" ifInstT='sub-interface'
 encap='vlan-40' />

                    <!--SVI 50 attached to eth1/25 on node 101 /-->
                   <l3extRsPathL3OutAtt tDn="topology/pod-1/paths-101/pathep-[eth1/25]" ifInstT='external-svi'
 encap='vlan-50' />
                </l3extLIfP>
            </l3extLNodeP>

            <!--External L3 EPG for Exporter behind external L3 Network /-->
            <l3extInstP name="accountingInst">
                <l3extSubnet ip="11.0.0.0/24" />
            </l3extInstP>
            <l3extRsEctx tnFvCtxName="ctx1"/>
        </l3extOut>
    </fvTenant>
</polUni>
```

Configuring NetFlow or Tetration Analytics Priority Using REST API

You can specify whether to use the NetFlow or Cisco Tetration Analytics feature by setting the `FeatureSel` attribute of the `<fabricNodeControl>` element. The `FeatureSel` attribute can have one of the following values:

- `analytics`—Specifies Cisco Tetration Analytics. This is the default value.
- `netflow`—Specifies NetFlow.

The following example REST API post specifies for the switch "test1" to use the NetFlow feature:

```
http://192.168.10.1/api/node/mo/uni/fabric.xml
<fabricNodeControl name="test1" FeatureSel="netflow" />
```

Addendum

About NetFlow Match Criteria

The filter ternary content-addressable memory (TCAM) in the FT block matches which flows must be installed in the flow table. This TCAM supports IPv4 and IPv6, as well as Layer 2 keys. For IPv4, the TCAM can hold 1k match criteria. IPv6 requires 4 entries and can only hold 256 match criteria.

Following keys are supported in the TCAM:

IP:

- Src TEP / VIF
- Dst TEP
- IP Flags
- TCP Flags
- Src IP
- Dst IP
- Tenant = VNI for infra transit or BD.
- Protocol
- Src L4 Port
- Dst L4 Port

CE:

- Src TEP
- Dst TEP
- Tenant
- Mac SA
- Mac DA
- Ethertype

Once a packet matches the criteria that is programmed in the TCAM and the TCAM action says to collect the flow with a certain mask, the packet is installed in the flow table.

About NetFlow Flow Masks

The EX switches provide 4 masks for each type of flow: IPv4, IPv6, and CE. This mask defines what constitutes the same flow from a set of packets, and one flow occupies one entry in the flow table. For example, you can configure a 5-tuple (SIP, DIP, Protocol, Sport, and Dport) and a bridge domain as a flow so that any packet that differs in these fields from any other packet is part of a different flow. If Sport is masked out, then all packets that match all the rest of the fields, but differ in this field, still constitute the same flow and statistics are collected in one entry in the table.

The following example packets illustrate how flow a mask works:

```
Pkt 1: BD1, 10.1.1.12 > 10.1.1.13, TCP, Sport 10000, Dport 80 Bytes = 100
Pkt 2: BD1, 10.1.1.12 > 10.1.1.13, TCP, Sport 20000, Dport 80 Bytes = 200
```

If the mask for these packets is set to mask off the Layer 4 Sport, the mask will create one entry in the flow table as follows:

```
Flow 1: BD1, 10.1.1.12 > 10.1.1.13, TCP, Sport = 0, Dport 80, Bytes = 300
```

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.