Data Plane Policing

This chapter contains the following sections:

Overview of Data Plane Policing

Use data plane policing (DPP) to manage bandwidth consumption on Cisco Application Centric Infrastructure (ACI) fabric access interfaces. DPP policies can apply to egress traffic, ingress traffic, or both. DPP monitors the data rates for a particular interface. When the data rate exceeds user-configured values, marking or dropping of packets occurs immediately. Policing does not buffer the traffic; therefore, the transmission delay is not affected. When traffic exceeds the data rate, the Cisco ACI fabric can either drop the packets or mark QoS fields in them.

Before the 3.2 release, the standard behavior for the policer was to be per-EPG member in the case of DPP policy being applied to the EPG, while the same policer was allocated on the leaf switch for the Layer 2 and Layer 3 case. This distinction was done because the DPP policer for Layer 2/Layer 3 case was assumed to be per-interface already, hence it was assumed different interfaces might get different ones. While the per-EPG DPP policy was introduced, it was clear that on a given leaf switch, several members could be present and therefore the policer it made sense to be per-member in order to avoid unwanted drops.

Starting with release 3.2, a clear semantic is given to the Data Plane Policer policy itself, as well as a new flag introducing the sharing-mode setting as presented in the CLI. Essentially, there is no longer an implicit behavior, which is different if the Data Plane Policer is applied to Layer 2/Layer 3 or to per-EPG case. Now the user has the control of the behavior. If the sharing-mode is set to **shared**, then all the entities on the leaf switch referring to the same Data Plane Policer, will share the same hardware policer. If the sharing-mode is set to **dedicated** then there would be a different HW policer allocated for each Layer 2 or Layer 3 or EPG member on the leaf switch. The policer is then dedicated to the entity that needs to be policed.

DPP policies can be single-rate, dual-rate, and color-aware. Single-rate policies monitor the committed information rate (CIR) of traffic. Dual-rate policers monitor both CIR and peak information rate (PIR) of traffic. In addition, the system monitors associated burst sizes. Three colors, or conditions, are determined by

the policer for each packet depending on the data rate parameters supplied: conform (green), exceed (yellow), or violate (red).

Typically, DPP policies are applied to physical or virtual layer 2 connections for virtual or physical devices such as servers or hypervisors, and on layer 3 connections for routers. DPP policies applied to leaf switch access ports are configured in the fabric access (infra) portion of the Cisco ACI fabric, and must be configured by a fabric administrator. DPP policies applied to interfaces on border leaf switch access ports (l3extOut or l2extOut) are configured in the tenant (fvTenant) portion of the Cisco ACI fabric, and can be configured by a tenant administrator.

The data plane policer can also be applied on an EPG so that traffic that enters the Cisco ACI fabric from a group of endpoints are limited per member access interface of the EPG. This is useful to prevent monopolization of any single EPG where access links are shared by various EPGs.

Only one action can be configured for each condition. For example, a DPP policy can to conform to the data rate of 256000 bits per second, with up to 200 millisecond bursts. The system applies the conform action to traffic that falls within this rate, and it would apply the violate action to traffic that exceeds this rate. Color-aware policies assume that traffic has been previously marked with a color. This information is then used in the actions taken by this type of policer.

For information about traffic storm control, see the *Cisco APIC Layer 2 Networking Configuration Guide*.

Guidelines and Limitations

The following are the guidelines and limitations for configuring data plane policing:

- The Data plane policer (DPP) does not police the packets transmitted from CPU and CPU bound packets on ACI fabric access interfaces.
- The **Dedicated Policer** sharing mode is only supported for EPG level policer, not supported for physical interfaces, Layer 2 interfaces nor Layer 3 interfaces.
- Egress traffic policing is not supported on Fabric Extender (FEX) ports

The following are the guidelines and limitations for the policer mode options: Bit policer mode (Bits-Per-Seconds: BPS) and Packet Policer mode (Packet-Per-Seconds: PPS)

- In egress direction, both BPS and PPS policer modes are supported. (PPS policer mode requires Cisco ACI 6.1(2) or later)
- In egress direction, PPS policer mode is supported only from FX switch onwards.
- In ingress direction, only BPS policer mode is supported.
- Policer statistics are available only in the policer mode that it is currently configured in. When you change the policer mode, the previously accumulated statistics are deleted.

The following are the guidelines and limitations for configuring Conform Action and Violate Action:

- In ingress direction, the following actions are supported:
  - Conform Action: **Drop and Transmit**
  - Violate Action: Drop, **Mark and Transmit**
- In egress direction, the following actions are supported:

- Conform Action: **Drop and Transmit**
- Violate Action: **Drop**

The following are guidelines and limitations for EPG policing:

- Feature support begins with switch models ending in EX/FX (example: N9K-C93180YC-EX) and subsequent models.
- Egress traffic policing is not supported on the EPG level policer.
- Policer type 2R3C is not supported.
- Policer is not supported when **intra-EPG isolation** is enforced in EPG.
- Statistics and considerations for **tuning** include:
  - Awareness of packets that are dropped/allowed is important to know to mitigate issues or for overuse of resources.
  - Statistics are provided in the GUI using the statistics infrastructure. Statistics are exported through the REST API as for any statistic in the Cisco ACI fabric.
  - Statistics are available on per-EPG member, and are useful if the Data Plane Policer policy is of type **dedicated**, otherwise the statistics reflect the statistics of all the ports using it on the leaf switch.
- In certain cases, such as when frames goes through FCoE supported devices, these get classified into the no drop FCoE class. In FCoE devices, this can cause drop off packets when the packet length is higher than the allowed 2184 bytes.

Configuring Data Plane Policing for Layer 2 Interface Using the GUI

Before you begin

The tenant, VRF, and external routed network where you configure the Data Plane Policing policy must be already created.

To apply the Layer 2 Data Plane Policing policy, the policy must be added to a policy group and the policy group must be mapped to an interface profile.

Procedure

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

In only the 3.2(1) release, the menu bar path is **Fabric** > **External Access Policies**

Step 2 In the Navigation pane, choose **Policies** > **Interface** > **Data Plane Policing**.

Step 3 Right-click **Data Plane Policing Policing**, and click **Create a Data Plane Policing Policy**.

Step 4 In the **Create a Data Plane Policing Policy** dialog box, in the **Name** field, enter a name for the policy.

Step 5 For **Administrative State**, choose **enabled**.

Step 6 For **BGP Domain Policer Mode**, choose either **Bit Policer** or **Packet Policer**.

Step 7 For **Type**, choose **1 Rate 2 Color** or **2 Rate 3 Color**.

Switch models ending in EX/FX (for example: N9K-C93180YC-EX) and subsequent models do not support **2 Rate 3 Color**.

Step 8 For **Conform Action**, choose an action.

This choice defines an actions for traffic that conforms with certain conditions.

- **Drop**: Drops the packets if the conditions are met.
- **Mark**: Marks the packets if the conditions are met.
- **Transmit**: Transmits the packets if the conditions are met.

Step 9 If for **Conform Action** you chose **Mark**, perform the following substeps:

a) For **Conform mark CoS**, enter the class of service for packets that conformed with the conditions.
b) For **Conform mark dscp**, enter the differentiated services code point (DSCP) for packets that conformed with the conditions.

Step 10 The administrator can configure the CoS and DSCP values in the **Conform** and **Violate** fields.

Step 11 If for **Type** you chose **2 Rate 3 Color**, then for **Exceed Action**, choose an action.

This choice defines an actions for traffic that exceeds certain conditions.

- **Drop**: Drops the packets if the conditions are met.
- **Mark**: Marks the packets if the conditions are met.
- **Transmit**: Transmits the packets if the conditions are met.

Step 12 If for **Exceed Action** you chose **Mark**, perform the following substeps:

a) For **Exceed mark CoS**, enter the class of service for packets that exceeded the conditions.
b) For **Exceed mark dscp**, enter the differentiated services code point (DSCP) for packets that exceeded the conditions.

Step 13 For **Violate Action**, choose an action.

This choice defines an actions for traffic that violates to certain conditions.

- **Drop**: Drops the packets if the conditions are met.
- **Mark**: Marks the packets if the conditions are met.
- **Transmit**: Transmits the packets if the conditions are met.

Step 14 If for **Violate Action** you chose **Mark**, perform the following substeps:

a) For **Violate mark CoS**, enter the class of service for packets that violated the conditions.
b) For **Violate mark dscp**, enter the differentiated services code point (DSCP) for packets that violated the conditions.

Step 15 For **Sharing Mode**, choose **Shared Policer**.

Shared Policer mode allows you to apply the same policing parameters to several interfaces simultaneously. The **Dedicated Policer** mode is not supported for Layer 2 interfaces.

Step 16 For **Rate**, enter the rate at which to allow packets are allowed into the system and choose the unit per packet.

Step 17 For **Burst**, enter the number of packets allowed at the line rate during a burst and choose the unit per packet.

Step 18 If for **Type** you chose **2 Rate 3 Color**, perform the following substeps:

a) For **Peak Rate**, enter the peak information rate, which is the rate above which data traffic is negatively affected, and choose the unit per packet.

b) For **Excessive Burst**, enter the size that a traffic burst can reach before all traffic exceeds the peak information rate, and choose the unit per packet.

Step 19 Click **Submit**.

Configuring Data Plane Policing for Layer 3 Interface Using the APIC GUI

Before you begin

The tenant, VRF, and external routed network where you configure the Data Plane Policing policy is already created.

The Data Plane Policing policy must be added to a policy group and the policy group mapped to an interface profile to apply the L3 DPP policy.

Procedure

Step 1 In the **Navigation** pane, click on **Tenant_name** > **Networking** > **External Routed Network** > **Network_name** > **Logical Node Profiles** > **Logical Node Profile_name** > **Logical Interface Profiles**, and perform the following actions.

a) Right-click on **Logical Interface Profiles**, and select **Create Interface Profile**.

b) In the **Create Interface Profile** dialog box, in the **Name** field, enter a name for the profile.

c) Next to **Ingress Data Plane Policing Policy**, select **Create Data Plane Policing Policy**.

d) In the **Name** field, enter a name for the policy.

e) In the **Administrative State** field, click **enabled**.

f) Next to **Policer Mode**, select a button for either **Bit Policer** or **Packet Policer**.

g) Next to **Type**, select a button for **1 Rate 2 Color** or **2 Rate 3 Color**.

Switch models ending in EX/FX (for example: N9K-C93180YC-EX) and subsequent models don't support 2 Rate 3 Color).

a) The administrator can configure the CoS and DSCP values in the **Conform** and **Violate** fields.

b) In the **Sharing Mode** field, select the policer mode.

Note
Shared Policer Mode allows you to apply the same policing parameters to several interfaces simultaneously.

c) Next to the **Burst**, **Excessive Burst** and **Rate** fields, select the drop down arrow to set the per packet rate for **1 Rate 2 Color** policy type.

Note
For **2 Rate 3 Color** policy type, the **Peak Rate** field is added.

d) Click **Submit**.

Step 2 Expand the **Routed Interfaces** table, in the **Path** field navigate to the interface to apply the policy and perform the following actions:

a) Next to **IPv4/Ipv6 Preferred Address**, enter a subnet IP address.
b) Click **OK**.
c) Click on the **SVI** tab and expand, in the **Path** field navigate to the interface to apply the policy.
d) Next to **Encap**, enter the VLAN name.
e) Next to **IPv4/Ipv6 Preferred Address**, enter a subnet IP address.
f) Click **OK**.
g) Expand the **Routed Sub-Interfaces** tab, and follow the same configuration steps as for the Routed Interfaces.
h) Click **OK**. This completes DPP configuration for L3.

Configuring Data Plane Policing Using the REST API

To police the Layer 2 traffic coming in to the leaf switch:

```
<!--  api/node/mo/uni/.xml  -->
<infraInfra>
<qosDppPol name="infradpp5" burst="2000" rate="2000" be="400" sharingMode="shared"/>
<!--
 List of nodes. Contains leaf selectors. Each leaf selector contains list of node blocks
-->
<infraNodeP name="leaf1">
<infraLeafS name="leaf1" type="range">
<infraNodeBlk name="leaf1" from_="101" to_="101"/>
</infraLeafS>
<infraRsAccPortP tDn="uni/infra/accportprof-portselector1"/>
</infraNodeP>
<!--
 PortP contains port selectors. Each port selector contains list of ports. It
       also has association to port group policies
-->
<infraAccPortP name="portselector1">
<infraHPortS name="pselc" type="range">
<infraPortBlk name="blk" fromCard="1" toCard="1" fromPort="48" toPort="49"></infraPortBlk>
<infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-portSet2"/>
</infraHPortS>
</infraAccPortP>
<!--  FuncP contains access bundle group policies  -->
<infraFuncP>
<infraAccPortGrp name="portSet2">
<infraRsQosIngressDppIfPol tnQosDppPolName="infradpp5"/>
</infraAccPortGrp>
</infraFuncP>
</infraInfra>
```

To police the Layer 2 traffic going out of the leaf switch:

```
<!--  api/node/mo/uni/.xml  -->
<infraInfra>
<qosDppPol name="infradpp2" burst="4000" rate="4000"/>
<!--
 List of nodes. Contains leaf selectors. Each leaf selector contains list of node blocks
-->
<infraNodeP name="leaf1">
```

```
<infraLeafS name="leaf1" type="range">
<infraNodeBlk name="leaf1" from_="101" to_="101"/>
</infraLeafS>
<infraRsAccPortP tDn="uni/infra/accportprof-portselector2"/>
</infraNodeP>
<!--
 PortP contains port selectors. Each port selector contains list of ports. It
       also has association to port group policies
-->
<infraAccPortP name="portselector2">
<infraHPortS name="pselc" type="range">
<infraPortBlk name="blk" fromCard="1" toCard="1" fromPort="37" toPort="38"></infraPortBlk>
<infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-portSet2"/>
</infraHPortS>
</infraAccPortP>
<!--  FuncP contains access bundle group policies  -->
<infraFuncP>
<infraAccPortGrp name="portSet2">
<infraRsQosEgressDppIfPol tnQosDppPolName="infradpp2"/>
</infraAccPortGrp>
</infraFuncP>
</infraInfra>
```

To police the Layer 3 traffic coming in to the leaf switch:

```
<!--  api/node/mo/uni/.xml  -->
<fvTenant name="dppTenant">
<qosDppPol name="gmeo" burst="2000" rate="2000"/>
<l3extOut name="Outside">
<l3extInstP name="extroute"/>
<l3extLNodeP name="borderLeaf">
<l3extRsNodeL3OutAtt tDn="topology/pod-1/node-101" rtrId="10.0.0.1">
<ipRouteP ip="0.0.0.0">
<ipNexthopP nhAddr="192.168.62.2"/>
</ipRouteP>
</l3extRsNodeL3OutAtt>
<l3extLIfP name="portProfile">
<l3extRsPathL3OutAtt addr="192.168.40.1/30" ifInstT="l3-port"
tDn="topology/pod-1/paths-101/pathep-[eth1/40]"/>
<l3extRsPathL3OutAtt addr="192.168.41.1/30" ifInstT="l3-port"
tDn="topology/pod-1/paths-101/pathep-[eth1/41]"/>
<l3extRsIngressQosDppPol tnQosDppPolName="gmeo"/>
</l3extLIfP>
</l3extLNodeP>
</l3extOut>
</fvTenant>
```

To police the Layer 3 traffic going out of the leaf switch:

```
<!--  api/node/mo/uni/.xml  -->
<fvTenant name="dppTenant">
<qosDppPol name="gmeo" burst="2000" rate="2000"/>
<l3extOut name="Outside">
<l3extInstP name="extroute"/>
<l3extLNodeP name="borderLeaf">
<l3extRsNodeL3OutAtt tDn="topology/pod-1/node-101" rtrId="10.0.0.1">
<ipRouteP ip="0.0.0.0">
<ipNexthopP nhAddr="192.168.62.2"/>
</ipRouteP>
</l3extRsNodeL3OutAtt>
<l3extLIfP name="portProfile">
<l3extRsPathL3OutAtt addr="192.168.40.1/30" ifInstT="l3-port"
tDn="topology/pod-1/paths-101/pathep-[eth1/40]"/>
<l3extRsPathL3OutAtt addr="192.168.41.1/30" ifInstT="l3-port"
tDn="topology/pod-1/paths-101/pathep-[eth1/41]"/>
```

```
<l3extRsEgressQosDppPol tnQosDppPolName="gmeo"/>
</l3extLIfP>
</l3extLNodeP>
</l3extOut>
</fvTenant>
```

Configuring Data Plane Policing Using NX-OS Style CLI

Procedure

Step 1 Configure a Layer 2 port to carry one EPG.

Example:

```
apic1# conf t
apic1(config)# vlan-domain test
apic1(config-vlan)# vlan 1000-2000
apic1(config-vlan)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/10
apic1(config-leaf-if)# vlan-domain member test
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
apic1(config)# tenant test1
apic1(config-tenant)# vrf context v1
apic1(config-tenant-vrf)# exit
apic1(config-tenant)# bridge-domain bd1
apic1(config-tenant-bd)# vrf member v1
apic1(config-tenant-bd)# exit
apic1(config-tenant)# application ap1
apic1(config-tenant-app)# epg e1
apic1(config-tenant-app-epg)# bridge-domain member bd1
apic1(config-tenant-app-epg)# exit
apic1(config-tenant-app)# exit
apic1(config-tenant)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/10
apic1(config-leaf-if)# switchport trunk allowed vlan 1001 tenant test1 application ap1 epg e1
apic1(config-leaf-if)# switchport trunk allowed vlan 1501 tenant test1 application ap1 epg e1
# Now the port leaf 101 ethernet 1/10 carries two vlan mapped both to the same Tenant/Application/EPG
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
```

a) Create a policy-map to apply to the interface.

Example:

```
apic1(config)# policy-map type data-plane qosTest
apic1(config-pmap-dpp)# set burst 2400 mega
apic1(config-pmap-dpp)# set cir 70 mega

apic1(config-pmap-dpp)# set sharing-mode shared
apic1(config-pmap-dpp)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/10
apic1(config-leaf-if)# service-policy type data-plane input qosTest
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
```

```
apic1(config)# policy-map type data-plane qosTest2
apic1(config-pmap-dpp)# set cir 78 mega
apic1(config-pmap-dpp)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/10
apic1(config-leaf-if)# service-policy type data-plane output qosTest2
apic1(config-leaf-if)# end
```

b) Visualize the policy configured.

Example:

```
apic1# show policy-map type data-plane infra
Type data-plane policy-maps
===================
Global Policy
policy-map type data-plane default
    set burst unspecified
    set conform-cos-transmit unspecified
    set conform-dscp-transmit unspecified
    set conform transmit
    set excessive-burst unspecified
    set exceed-cos-transmit unspecified
    set exceed-dscp-transmit unspecified
    set exceed drop
    set mode byte
    set pir 0
    set cir 78 mega
    set type 1R2C
    set violate-cos-transmit unspecified
    set violate-dscp-transmit unspecified
    set violate drop
Global Policy
policy-map type data-plane qosTest
    set burst 2400 mega
    set cir 78 mega
    set conform-cos-transmit unspecified
    set conform-dscp-transmit unspecified
    set conform transmit
    set excessive-burst unspecified
    set exceed-cos-transmit unspecified
    set exceed-dscp-transmit unspecified
    set exceed drop
    set mode byte
    set pir 0
    set type 1R2C
    set violate-cos-transmit unspecified
    set violate-dscp-transmit unspecified
    set violate drop
Global Policy
policy-map type data-plane qosTest2
    set burst unspecified
    set conform-cos-transmit unspecified
    set conform-dscp-transmit unspecified
    set conform transmit
    set excessive-burst unspecified
    set exceed-cos-transmit unspecified
    set exceed-dscp-transmit unspecified
    set exceed drop
    set mode byte
    set pir 0
    set cir 78 mega
    set type 1R2C
    set violate-cos-transmit unspecified
```

```
    set violate-dscp-transmit unspecified
    set violate drop
```

c) Show running-config.

Example:

```
apic1# show runn policy-map
# Command: show running-config policy-map
# Time: Fri Jan 29 19:26:18 2016
  policy-map type data-plane default
    exit
  policy-map type data-plane qosTest
    set burst 2400 mega
    set cir 78 mega
    no shutdown
    exit
  policy-map type data-plane qosTest2
    set cir 78 mega
    no shutdown
    exit
apic1# show runn leaf 101
# Command: show running-config leaf 101
# Time: Fri Jan 29 19:26:29 2016
  leaf 101
    interface ethernet 1/10
      vlan-domain member test
      switchport trunk allowed vlan 1501 tenant test1 application ap1 epg e1
      service-policy type data-plane input qosTest
      service-policy type data-plane output qosTest2
      exit
    exit
```

Step 2 Preparation to configure Layer 3 ports.

Example:

```
apic1# conf t
apic1(config)# vlan-domain l3ports
apic1(config-vlan)# vlan 3000-3001
apic1(config-vlan)# exit
apic1(config)# tenant l3test1
apic1(config-tenant)# vrf context v1
apic1(config-tenant-vrf)# exit
apic1(config-tenant)# exit
apic1(config)# leaf 102
apic1(config-leaf)# vrf context tenant l3test1 vrf v1
apic1(config-leaf-vrf)# exit
# Configure a physical Layer 3 port
apic1(config-leaf)# interface ethernet 1/20
apic1(config-leaf-if)# no switchport
apic1(config-leaf-if)# vlan-domain member l3ports
apic1(config-leaf-if)# vrf member tenant l3test1 vrf v1
apic1(config-leaf-if)# ip address 56.1.1.1/24
apic1(config-leaf-if)# ipv6 address 2000::1/64  preferred
apic1(config-leaf-if)# exit
# Configure base interface for L3 subinterfaces
apic1(config-leaf)# interface ethernet 1/21
apic1(config-leaf-if)# vlan-domain member l3ports
apic1(config-leaf-if)# no switchport
apic1(config-leaf-if)# exit
# Configure a Layer 3 subinterface
apic1(config-leaf)# interface ethernet 1/21.3001
apic1(config-leaf-if)# vrf member tenant l3test1 vrf v1
```

```
apic1(config-leaf-if)# ip address 60.1.1.1/24
apic1(config-leaf-if)# ipv6 address 2001::1/64 preferred
apic1(config-leaf-if)# exit
# Configure a Switched Vlan Interface
apic1(config-leaf)# interface vlan 3000
apic1(config-leaf-if)# vrf member tenant l3test1 vrf v1
apic1(config-leaf-if)# ip address 70.1.1.1/24
apic1(config-leaf-if)# ipv6 address 3000::1/64 preferred
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
```

a) Configure the policer in the tenant for Layer 3 usage.

Example:

```
apic1(config)# tenant l3test1
apic1(config-tenant)# policy-map type data-plane iPol
apic1(config-tenant-pmap-dpp)# set cir 56 mega
apic1(config-tenant-pmap-dpp)# set burst 2000 kilo
apic1(config-tenant-pmap-dpp)# exit
apic1(config-tenant)# policy-map type data-plane ePol
apic1(config-tenant-pmap-dpp)# set burst 2000 kilo
apic1(config-tenant-pmap-dpp)# set cir 56 mega
apic1(config-tenant-pmap-dpp)# exit
apic1(config-tenant)# exit
```

b) Apply policer on a Layer 3 interface

Example:

```
apic1(config)# leaf 102
apic1(config-leaf)# interface ethernet 1/20
apic1(config-leaf-if)# service-policy type data-plane input iPol
apic1(config-leaf-if)# service-policy type data-plane output ePol
apic1(config-leaf-if)# exit
apic1(config-leaf)# interface ethernet 1/21.3001
apic1(config-leaf-if)# service-policy type data-plane input iPol
apic1(config-leaf-if)# service-policy type data-plane output ePol
apic1(config-leaf-if)# exit
apic1(config-leaf)# interface vlan 3000
apic1(config-leaf-if)# service-policy type data-plane input iPol
apic1(config-leaf-if)# service-policy type data-plane output ePol
apic1(config-leaf-if)# end
```

c) Show commands for policers used on a Layer 3 interface.

Example:

```
apic1# show tenant l3test1 policy-map type data-plane
Type data-plane policy-maps
====================
Policy in Tenant: l3test1
policy-map type data-plane ePol
    set burst 2000 kilo
    set conform-cos-transmit unspecified
    set conform-dscp-transmit unspecified
    set conform transmit
    set excessive-burst unspecified
    set exceed-cos-transmit unspecified
    set exceed-dscp-transmit unspecified
    set exceed drop
    set mode byte
    set pir 0
    set cir 56 mega
    set type 1R2C
```

```
    set violate-cos-transmit unspecified
    set violate-dscp-transmit unspecified
    set violate drop
Policy in Tenant: l3test1
policy-map type data-plane iPol
    set burst 2000 kilo
    set burst unspecified
    set conform-cos-transmit unspecified
    set conform-dscp-transmit unspecified
    set conform transmit
    set excessive-burst unspecified
    set exceed-cos-transmit unspecified
    set exceed-dscp-transmit unspecified
    set exceed drop
    set mode byte
    set pir 0
    set cir 56 mega
    set type 1R2C
    set violate-cos-transmit unspecified
    set violate-dscp-transmit unspecified
    set violate drop
```

d) Show running-config for policers used for Layer 3.

Example:

```
apic1# show runn tenant l3test1
# Command: show running-config tenant l3test1
# Time: Fri Jan 29 19:48:20 2016
  tenant l3test1
    vrf context v1
      exit
    policy-map type data-plane ePol
      set burst 2000 kilo
      set cir 56 mega
      no shutdown
      exit
    policy-map type data-plane iPol
      set burst 2000 kilo
      set cir 56 mega
      no shutdown
      exit
    exit
apic1# show running-config leaf 102
# Command: show running-config leaf 102
# Time: Fri Jan 29 19:48:33 2016
  leaf 102
    vrf context tenant l3test1 vrf v1
      exit
    interface vlan 3000
      vrf member tenant l3test1 vrf v1
      ip address 70.1.1.1/24
      ipv6 address 3000::1/64 preferred
      bfd ip tenant mode
      bfd ipv6 tenant mode
      service-policy type data-plane input iPol
      service-policy type data-plane output ePol
      exit
    interface ethernet 1/20
      vlan-domain member l3ports
      no switchport
      vrf member tenant l3test1 vrf v1
      ip address 56.1.1.1/24
      ipv6 address 2000::1/64 preferred
      bfd ip tenant mode
```

```
      bfd ipv6 tenant mode
      service-policy type data-plane input iPol
      service-policy type data-plane output ePol
      exit
    interface ethernet 1/21
      vlan-domain member l3ports
      no switchport
      bfd ip tenant mode
      bfd ipv6 tenant mode
      exit
    interface ethernet 1/21.3001
      vrf member tenant l3test1 vrf v1
      ip address 60.1.1.1/24
      ipv6 address 2001::1/64 preferred
      bfd ip tenant mode
      bfd ipv6 tenant mode
      service-policy type data-plane input iPol
      service-policy type data-plane output ePol
    exit
    exit
apic1#
```

Data Plane Policing at the Endpoint Group Level

Data Plane Policing (DPP) can be applied to an endpoint group (EPG). The policing of the traffic is applied to all the EPG members on every leaf switch where the EPG is deployed.

Prior to the 3.2(1) release, each EPG member had its own policer. Beginning in the 3.2(1) release, the behavior is dependent on the sharing-mode property (if configured through the CLI or GUI) on the Data Plane Policer. If that is set to **dedicated**, then the situation is similar to before the 3.2(1) release. If the sharing-mode is set to **shared**, then all the members in the same slice using the same Data Plane Policer policy use the hardware policer on the leaf switch.

For example, an EPG has the following members:

- Leaf 101, Eth1/1, vlan-300
- Leaf 101, Eth1/2, vlan-301
- Leaf 102, Eth1/2, vlan-500

In this case, each member will limit the traffic according to the policer, independent from the other members. If the Data Plane Policer has the sharing-mode set to **shared**, then all the members in the same slice above use only one policer on the leaf switch.

The Data Plane Policer works independently on Leaf 101 and Leaf 102 if the sharing-mode is set to **dedicated**. For example:

- Policer-A (100Mbps policing) is applied to EPG1 (Leaf101 e1/1 vlan-300 and e1/2 vlan-301. Leaf 102 e1/2 vlan-500)
- Leaf 101: police traffic at the EPG1 level, which is applied to traffic through E1/1 vlan-300 and E1/2 vlan-301 (100Mbps for each interface).
- Leaf 102: police traffic at the EPG1 level, which is applied to traffic through E1/2 vlan-500 (another 100Mbps for each interface).

The total is up to 300Mbps for EPG1.

If the sharing-mode is set to **shared**, 100Mbps is shared across EPGs using the same policer if the interfaces are in the same slice. For example:

- Policer-A (100Mbps policing) applied to EPG1 and EPG2.
- Leaf 101: police traffic at EPG1 and EPG2 in total.
- Leaf 102: police traffic at EPG1 and EPG2 in total.

The total is up to 200Mbps for EPG1 and EPG2 if the interfaces are in the same slice.

The following are limitations for Data Plane Policing at the EPG level:

- EPG policer feature is supported with switch models that have -EX, -FX, or later suffixes in the product ID.
- Egress traffic policing is not supported for the EPG level policer.
- Policer mode **Packet-per-second** is not supported.
- Policer type 2R3C is not supported in EPG policer.
- Policer is not supported when **intra-EPG isolation-enforced** is applied to the EPG.
- The scale limit allows for 128 EPG policers supported per node.

Configuring Data Plane Policing at the Endpoint Group Level Using CLI

SUMMARY STEPS

1. Define the policer:

DETAILED STEPS

Procedure

Define the policer:

Example:

```
apic1# conf t
apic1(config)# vlan-domain test
apic1(config-vlan)# vlan 1000-2000
apic1(config-vlan)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/10
apic1(config-leaf-if)# vlan-domain member test
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
apic1(config)# tenant test1
apic1(config-tenant)# vrf context v1
apic1(config-tenant-vrf)# exit
apic1(config-tenant)# bridge-domain bd1
apic1(config-tenant-bd)# vrf member v1
```

```
apic1(config-tenant-bd)# exit
apic1(config)# policy-map type data-plane pol1
apic1(config-pmap-dpp)# set burst 2400 mega
apic1(config-pmap-dpp)# set cir 78 mega
apic1(config-pmap-dpp)# exit
apic1(config-tenant)# application ap1
apic1(config-tenant-app)# epg e1
apic1(config-tenant-app-epg)# bridge-domain member db1
apic1(config-tenant-app-epg)# service-policy type data-plane pol1
apic1(config-tenant-app-epg)# exit
apic1(config-tenant-app)# exit
apic1(config-tenant)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/10
apic1(config-leaf-if)# switchport trunk allowed vlan 1001 tenant test1 application ap1 epg e1
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
```

Configuring Data Plane Policing at the Endpoint Group Level Using the APIC GUI

Procedure

In the **Tenants** pane, click on **Tenant_name** > **Policies** > **Protocol** > **Data Plane Policing**. Right-click on **Data Plane Policing** to **Create Data Plane Policing Policy**.

a) In the **Name** field, enter a name for the policy.
b) In the **Administrative State** field, click **enabled**.
c) Next to **Policer Mode**, select a button for either **Bit Policer** or **Packet Policer**.
d) Next to **Type**, select a button for **1 Rate 2 Color**.
e) For **Conform Action**, select **Drop**, **Mark**, or **Transmit**.
f) The administrator can configure the CoS and DSCP values in the **Conform** and **Violate** fields.
g) Next to the **Burst**, **Excessive Burst** and **Rate** fields, click the drop down arrow to select from the following:

- **Bytes/Packets**
- **Kilo Bytes/Packets**
- **Mega Bytes/Packets**
- **Giga Bytes/Packets**
- **Milli Seconds**
- **Micro Seconds**

Configuring Data Plane Policing at the Endpoint Group Level Using Rest API

To police the traffic coming into the leaf switch:

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

    <qosDppPol name="gmeo" burst="2000" rate="2000"/>
    <fvAp name="ap1">
      <fvAEPg name="ep1">
        <fvRsDppPol tnQosDppPolName="gmeo"/>
      </fvAEPg>
    </fvAp>
  </fvTenant>
</polUni>
```

Accessing Statistics for the Data Plane Policer at the Endpoint Group Level in the GUI

DPP at the EPG level is used to police traffic at the EPG member level. As such, statistics are integral in ensuring the policer is dropping substantial traffic. Statistics are reported at the EPG member level for fine granularity.

Procedure

Step 1 In the **Tenants** pane, click on **Tenant_name** > **Application EPGs** > **EPG Members** > **Static EPG Members** .

Step 2 Select a node.

Step 3 Click **Select Stats**.

a) Select a **Sampling Interval** unit of time.
b) From the **Available** policer attributes, use the arrows to choose the attributes. You can select up to two attributes.
c) Click **Submit**.

What to do next

You will see a graphical representation of the DPP statistics.