Configure Atomic Counter Policies

Contents

Introduction

This document describes how atomic counters policies work on the fabric. This feature allows you to monitor the traffic drops/excess packets on your fabric.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on these software and hardware versions:

- Cisco Application Centric Infrastructure (ACI)
- APIC version 1.0(3n)
- n9000-aci version 11.0(3n)

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Background Information

These acronyms are used in this article:

- APIC - Application Policy Infrastructure Controller
- TEP - Tunnel End Point
- VRF - Virtual Routing & Forwarding
- TCAM - Ternary Content-Addressable Memory
- EPG - End Point Group
- MO - Managed Object

There is some important information contained in the "Troubleshoot" section that helps to understand the topic. Most importantly, the traffic that is measured must traverse through the fabric (leaf > spine > leaf) in order to take advantage of all of the atomic counter policies. Creation of a policy for two endpoints attached to the same leaf will only allow for the transmit counter to increment.

Note that there is more than one type of atomic counter. This document identifies how to configure the on-demand atomic counter policies. These can be toggled on or off by the administrator. There are also "always-on" atomic counters that measure traffic between leafs. These are the TEP-to-TEP atomic counters. They can be seen in these items:

- dbgIngrTep (Ingress TEP Counters)
- dbgEgrTep (Egress TEP Counters)

They are counted for each of the TEPs on each of the leafs. It is possible to poll the Application Policy Infrastructure Controller (APIC) for these numbers, but it is not recommended. The best course of action for a customer interested in monitoring traffic on their network would be to configure on-demand counters.

In order to work, the atomic counters flip an "M" bit on or off in the eVXLAN header. They are not incremented with respect to time, but with respect to the "packet". The M bit tells the node which Bank (Odd or Even) to increment for the packet. Atomic counters work by polling the nodes for the counter on their respective Odd and Even banks. For example, the APIC might measure the Odd Bank on leaf 1 and the Even Bank on leaf 4 because of an on-demand policy that is configured. This allows the APIC to derive transmitted and received packets from each bank count, then the number of drops and excess packets based on the difference.

When the on-demand policy is configured, the counters increment if a TCAM entry is matched and the odd/even bit is set. This means you must have a policy set via contracts between the two endpoints/endpoint groups/IPs that you try to measure before the atomic counters will work.

Here are some caveats to consider when you configure atomic policy counters:

- Use of atomic counters is not supported when the endpoints are in different tenants or in different contexts (VRFs) within the same tenant.
- In pure Layer 2 configurations where the IP address is not learned (the IP address is 0.0.0.0), endpoint-to-EPG and EPG-to-endpoint atomic counter policies are not supported. In these cases, endpoint-to-endpoint and EPG-to-EPG policies are supported. External policies are virtual routing and forwarding (VRF)-based, which requires learned IP addresses, and are supported.
- When the atomic counter source or destination is an endpoint, the endpoint must be dynamic

and not static. Unlike a dynamic endpoint (fv:CEp), a static endpoint (fv:StCEp) does not have a child object (fv:RsCEpToPathEp) that is required by the atomic counter.
- In a transit topology, where leaf switches are not in full mesh with all spine switches, then leaf-to-leaf (TEP to TEP) counters do not work as expected.
- For leaf-to-leaf (TEP to TEP) atomic counters, once the number of tunnels increases the hardware limit, the system changes the mode from trail mode to path mode and the user is no longer presented with per-spine traffic.
- The atomic counter does not count spine proxy traffic.
- Packets dropped before they enter the fabric or before they are forwarded to a leaf port are ignored by atomic counters.
- Packets that are switched in the hypervisor (same Port Group and Host) are not counted.
- Atomic counters require an active fabric Network Time Protocol (NTP) policy.
- An atomic counter policy configured with fvCEp as the source and/or destination counts only the traffic that is from/to the MAC and IP addresses that are present in the fvCEp managed objects (MOs). If the fvCEp MO has an empty IP address field, then all traffic to/from that MAC address would be counted regardless of the IP address. If the APIC has learned multiple IP addresses for an fvCEp, then traffic from only the one IP address in the fvCEp MO itself is counted as previously stated. In order to configure an atomic counter policy to/from a specific IP address, use the fvIp MO as the source and/or destination.
- If there is an fvIp behind an fvCEp, you must add fvIP-based policies and not fvCEp-based policies.

See Cisco APIC Troubleshooting Guide - Atomic Counters Guidelines and Restrictions for more information.

Configure

In order to configure atomic counter policies, complete these steps:

1. Determine which type of atomic counter policy you want to configure.
2. Create the policy.
3. Add the filter you would like to use for the policy.

Determine Which Type of Atomic Policy You Want to Configure

These types of on-demand atomic counter policies can be configured:

- EP to EP
- EP to EPG
- EP to Ext
- EPG to EP
- EPG to EPG
- EPG to IP
- Ext to IP
- IP to EPG

The meaning of each acronym is as follows:

- EP - Endpoint

- EPG - Endpoint Group
- Ext - External Network
- IP - IP address

Note that for any of the EP-based policies, the endpoint must already be learned on the fabric before the policy can be configured.

The type of policy that you choose to configure will determine the parameters that must be configured in the next section.

Create the Policy

The screenshots used in this section are for an EPG to EPG policy. Your view might vary based upon which type you configure, but the core concepts will be the same.

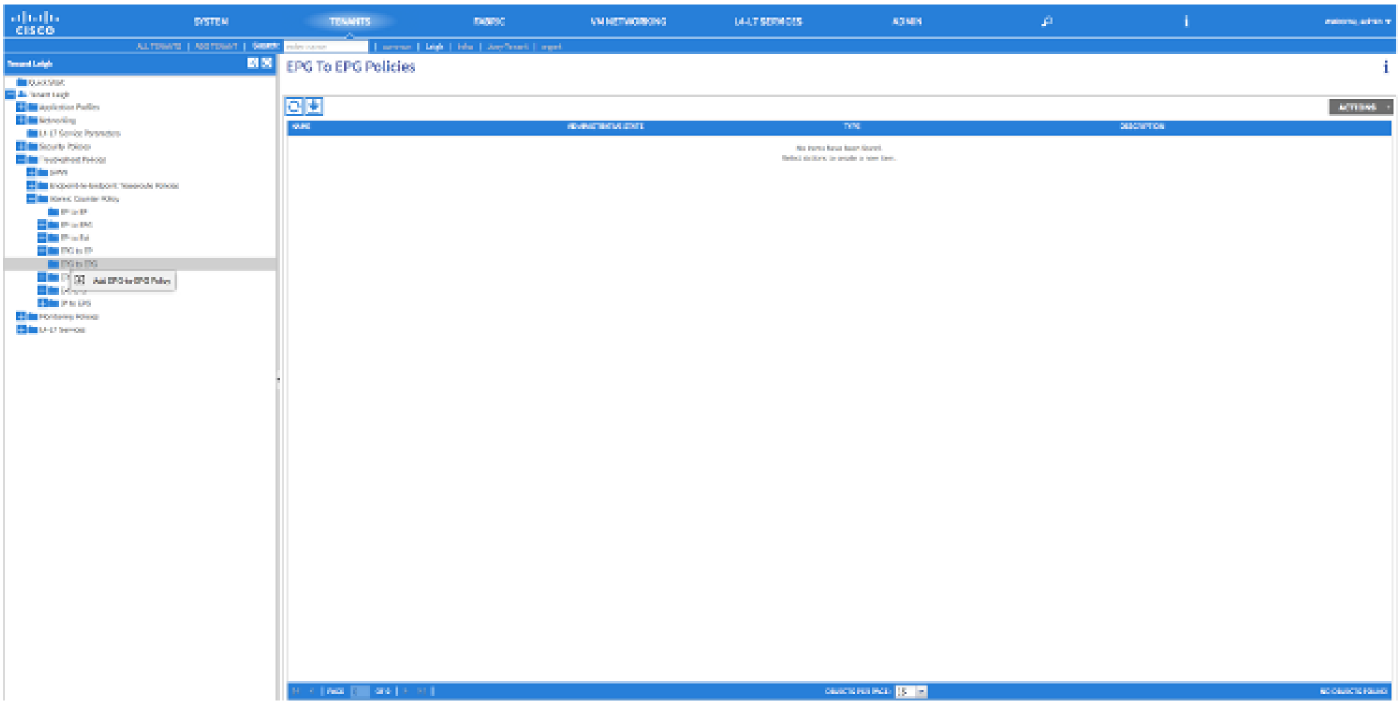

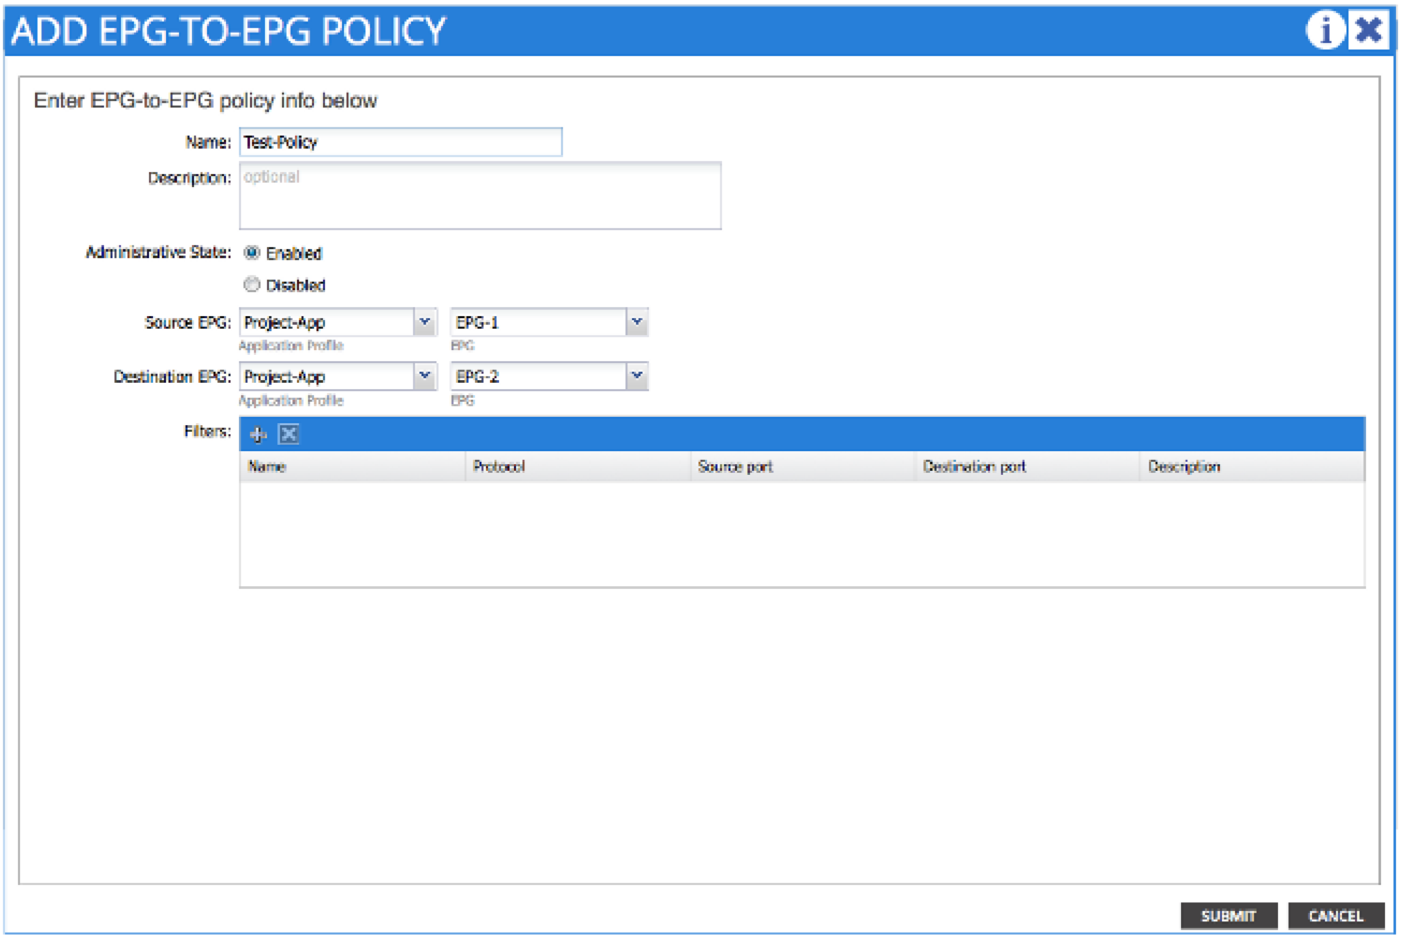
ADD EPG-TO-EPG POLICY

Enter EPG-to-EPG policy info below

Name: Test-Policy

Description: optional

Administrative State: ● Enabled ○ Disabled

Source EPG: Project-App (Application Profile) EPG-1 (EPG)

Destination EPG: Project-App (Application Profile) EPG-2 (EPG)

Filters:

Name	Protocol	Source port	Destination port	Description

EP to EP

You can choose between two source types: EP and IP. If you choose EP, then you select an endpoint that has been learned on the fabric. If you choose IP, then you select an endpoint that has been learned on the fabric as well as an IP address. This allows you to get more granular in deciding between a specific endpoint and a specific IP host that might sit behind an endpoint.

EPG to EPG

Select a source and a destination EPG for the policy. This measures traffic that goes from all endpoints in the source EPG to any endpoints in the destination EPG.

EP to EPG

The process to choose the source is the same as for the "EP to EP" policy. The process to choose the destination is the same as for the "EPG to EPG" policy".

EP to Ext

The process to choose the source is the same as for the "EP to EP" policy. You must enter an "External IP" in order to specify an IP address outside the fabric that will be used as a destination for the counter. You can either choose one specific IP address or an IP address range by putting a "/" after the address and specifying a subnet size.

EPG to EP

The process to choose the source is the same as for the "EPG to EPG" policy. The process to choose the destination is the same as for the "EP to EP" policy.

EPG to IP

The process to choose the source is the same as for the "EPG to EPG" policy. The process to choose the destination is the same as for the "EP to Ext" policy.

Ext to IP

Select a source IP address for the traffic and enter it in the "Source IP" field. It can be either a specific IP address or an IP subnet. The process to choose a destination is the same as for the "EP to EP" policy.

IP to EPG

The process to choose the source is the same as for the "Ext to IP" policy. The process to choose the destination is the same as for the "EPG to EPG" policy.

Add the Filter You Would Like to Use for the Policy

The screen you see here is consistent regardless of what type of policy you configure. Note that the Atomic Counter Filter is a different type of object than the filter you apply to Contracts in the fabric, although they serve similar functions.

- Name - Enter the name for the Atomic Counter Filter here. Note that this filter is specific to this policy only and will not be reused.
- Protocol - You can either choose a protocol from the drop-down list or enter a number that corresponds to the protocol between 0 and 255. The 0 to 255 range corresponds to the IP protocol number contained in the IP packet header.
- Source port - You can either choose one of the commonly used protocols from the drop-down list or enter a number between 0 and 65535.
- Destination port - You can either choose one of the commonly used protocols from the drop-down list or enter a number between 0 and 65535.
- Description - This is just a description for the filter to aid in identification. It will not affect which traffic is or is not identified by this filter.

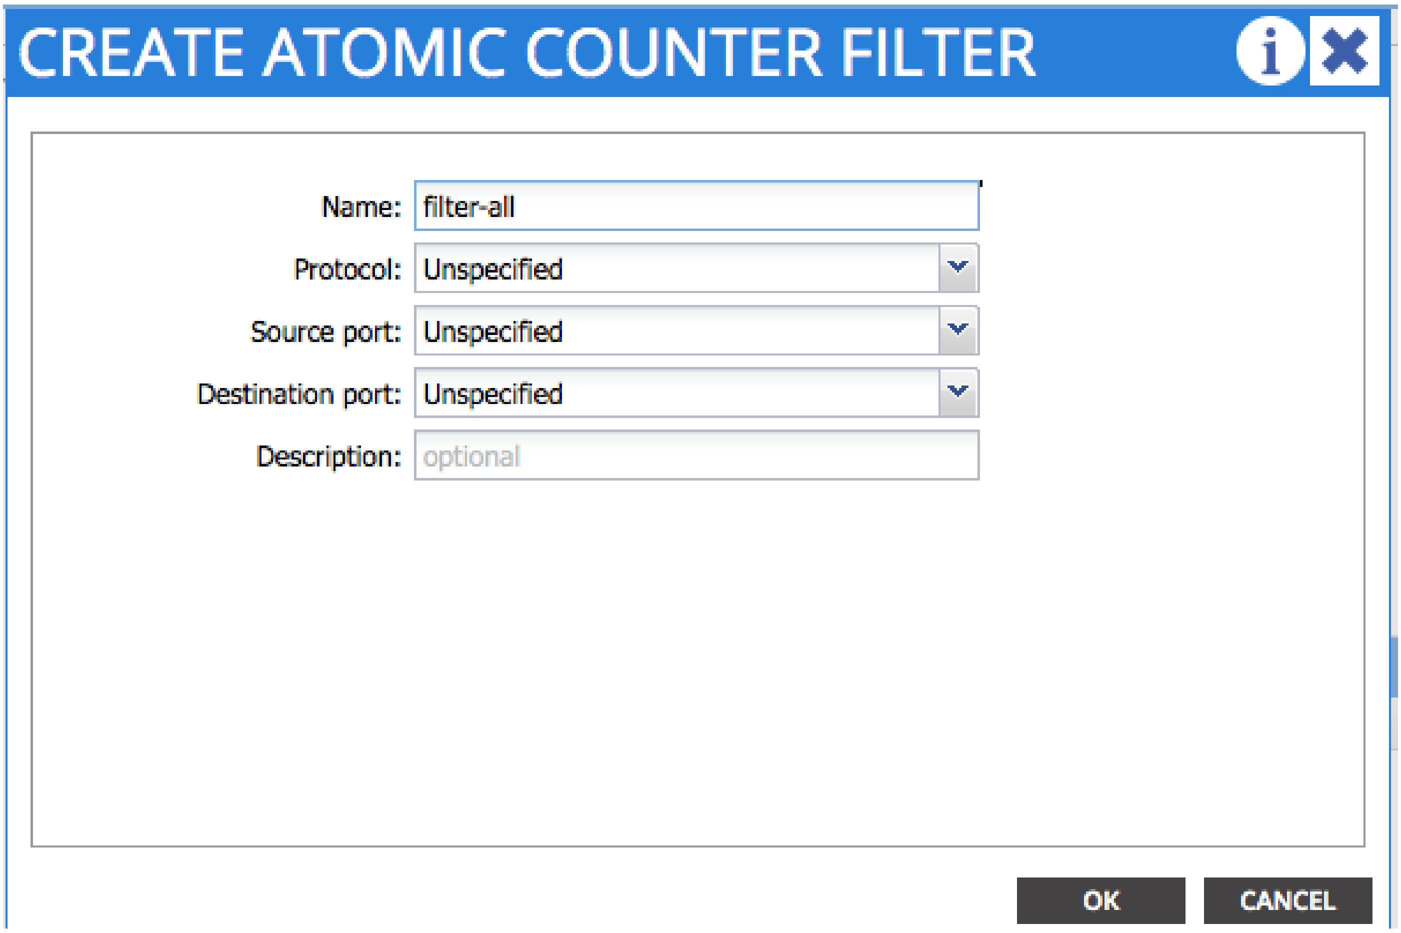

Name: filter-all

Protocol: Unspecified

Source port: Unspecified

Destination port: Unspecified

Description: optional

You can also configure atomic counters with REST API. Here is an example of the POST request used to create an EPG-to-EPG policy:

URL - https://**<apic-ip>**/api/node/mo/uni/tn-Leigh/epgToEpg-Test-Policy.json

JSON

```
{"dbgacEpgToEpg":
   {"attributes":
       {"dn":"uni/tn-Leigh/epgToEpg-Test-Policy",
       "name":"Test-Policy",
       "rn":"epgToEpg-Test-Policy",
       "status":"created"},
       "children":[
            {"dbgacFilter":
                {"attributes":
                    {"dn":"uni/tn-Leigh/epgToEpg-Test-Policy/filt-filter-all",
                    "name":"filter-all",
                    "rn":"filt-filter-all",
                    "status":"created"},
                    "children":[]}},
            {"dbgacRsFromEpg":
                {"attributes":
                    {"tDn":"uni/tn-Leigh/ap-Project-App/epg-EPG-1",
                    "status":"created,modified"},
                    "children":[]}},
            {"dbgacRsToEpgForEpgToEpg":
                {"attributes":
                    {"tDn":"uni/tn-Leigh/ap-Project-App/epg-EPG-2",
                    "status":"created"},
                    "children":[]
                }
            }
        ]
   }
}
```

Verify

Use this section to confirm that your configuration works properly.

The easiest way to verify that the atomic counter policy you configured is operational is to ensure that the "Administrative State" is set to "Enabled" under the "Policy" tab.

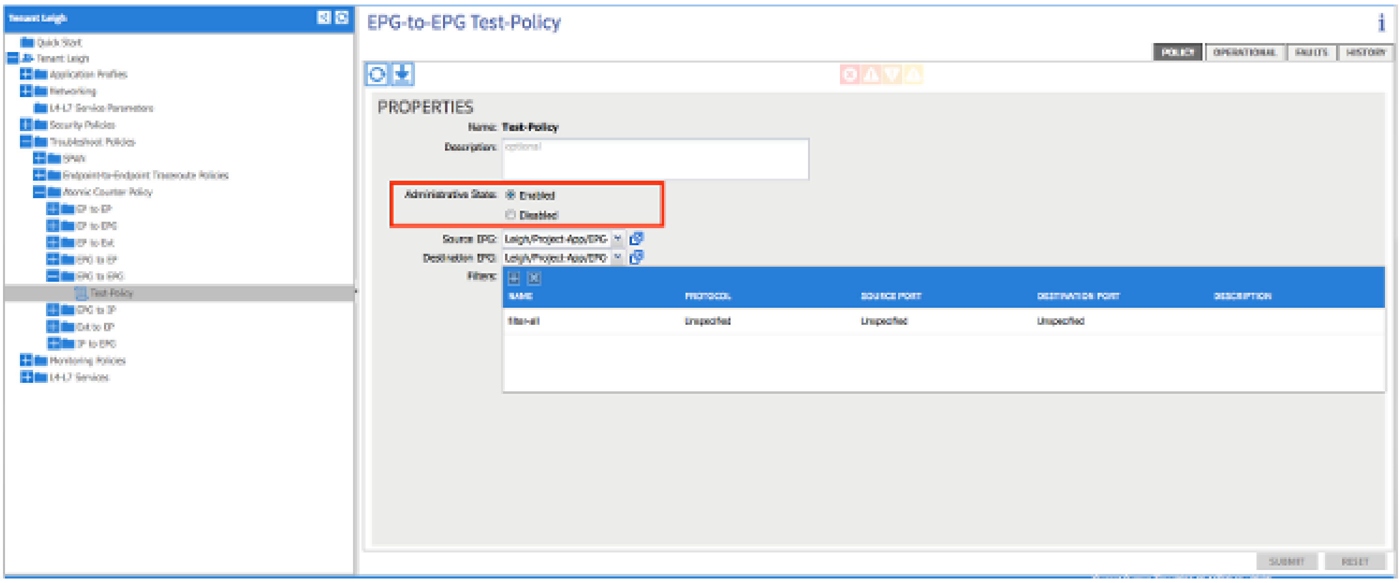

EPG-to-EPG Test-Policy

PROPERTIES

Name: Test-Policy

Description: optional

Administrative State: ● Enabled ○ Disabled

Source EPG: Leigh/Project-App/EPG

Destination EPG: Leigh/Project-App/EPG

Filters:

NAME	PROTOCOL	SOURCE PORT	DESTINATION PORT	DESCRIPTION
filter-all	Unspecified	Unspecified	Unspecified	

In order to see the counters for each statistic on the policy, navigate to the "Operational" tab. Here you should see the number of transmitted and admitted packets increment if traffic flows. A minor fault is triggered if 1% or more of packets are dropped and a major fault is triggered if 5% or more of the packets are dropped.

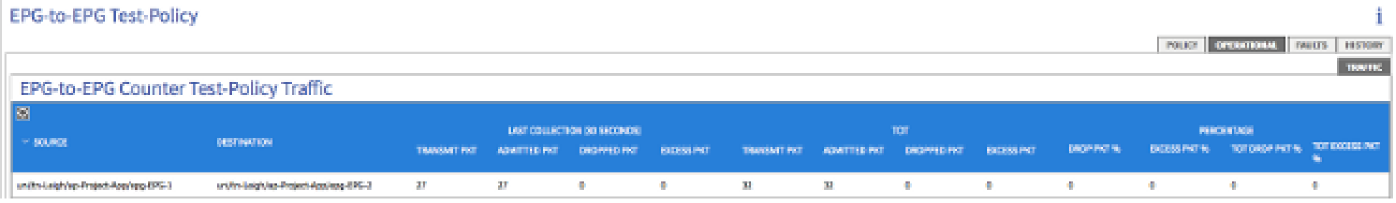

EPG-to-EPG Test-Policy

EPG-to-EPG Counter Test-Policy Traffic

		LAST COLLECTION (30 SECONDS)				TOT				PERCENTAGE			
SOURCE	DESTINATION	TRANSMIT PKT	ADMITTED PKT	DROPPED PKT	EXCESS PKT	TRANSMIT PKT	ADMITTED PKT	DROPPED PKT	EXCESS PKT	DROP PKT %	EXCESS PKT %	TOT DROP PKT %	TOT EXCESS PKT %
uni/tn-Leigh/ap-Project-App/epg-EPG-1	uni/tn-Leigh/ap-Project-App/epg-EPG-2	27	27	0	0	33	33	0	0	0	0	0	0

Troubleshoot

This section provides information you can use to troubleshoot your configuration.

If you do not see any of the counters increment, here are some possible issues that you might face:

- Is the policy enabled?
- Is the filter for the policy configured correctly?
- Are there contracts in place between the two endpoints or devices that you measure traffic between?

If you a certain the policy is configured correctly, enabled, and the endpoints that are tested successfully pass traffic, then the issue is likely that the two endpoints are connected to the same leaf. Due to the design of the hardware architecture, the traffic must go through the Northstar ASICs on the leafs for the counters to increment. If the traffic only passes through one leaf, then you will only see the transmit counters increment.

If you see a high number of dropped or excess packets, then one possibility is that you have an oversubscription between two devices.