Configure and Verify Layer 2 Service Graph Configuration with ASAv

Contents

Introduction

This document describes how to configure and verify Layer 2 Service Graph configuration in Cisco Application Centric Infrastructure (ACI).

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Understanding of how Layer 3 Service Graph works in ACI
- Understanding on how to configure Endpoint policy group, bridge domains and contract in ACI
- Understanding on how to configure (Adaptive Security Applicance Virtual) ASAv as transparent firewall

Components Used

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

- APIC Version: 6.0 (3g)
- Leaf H/W: N9K-C93180YC-FX
- Leaf S/W: n9000-16.0 (3g)
- Leaf Node 101, 102, 103
- ASAv Deployed on ESXi server

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, ensure that you understand the potential impact of any command.

Topology

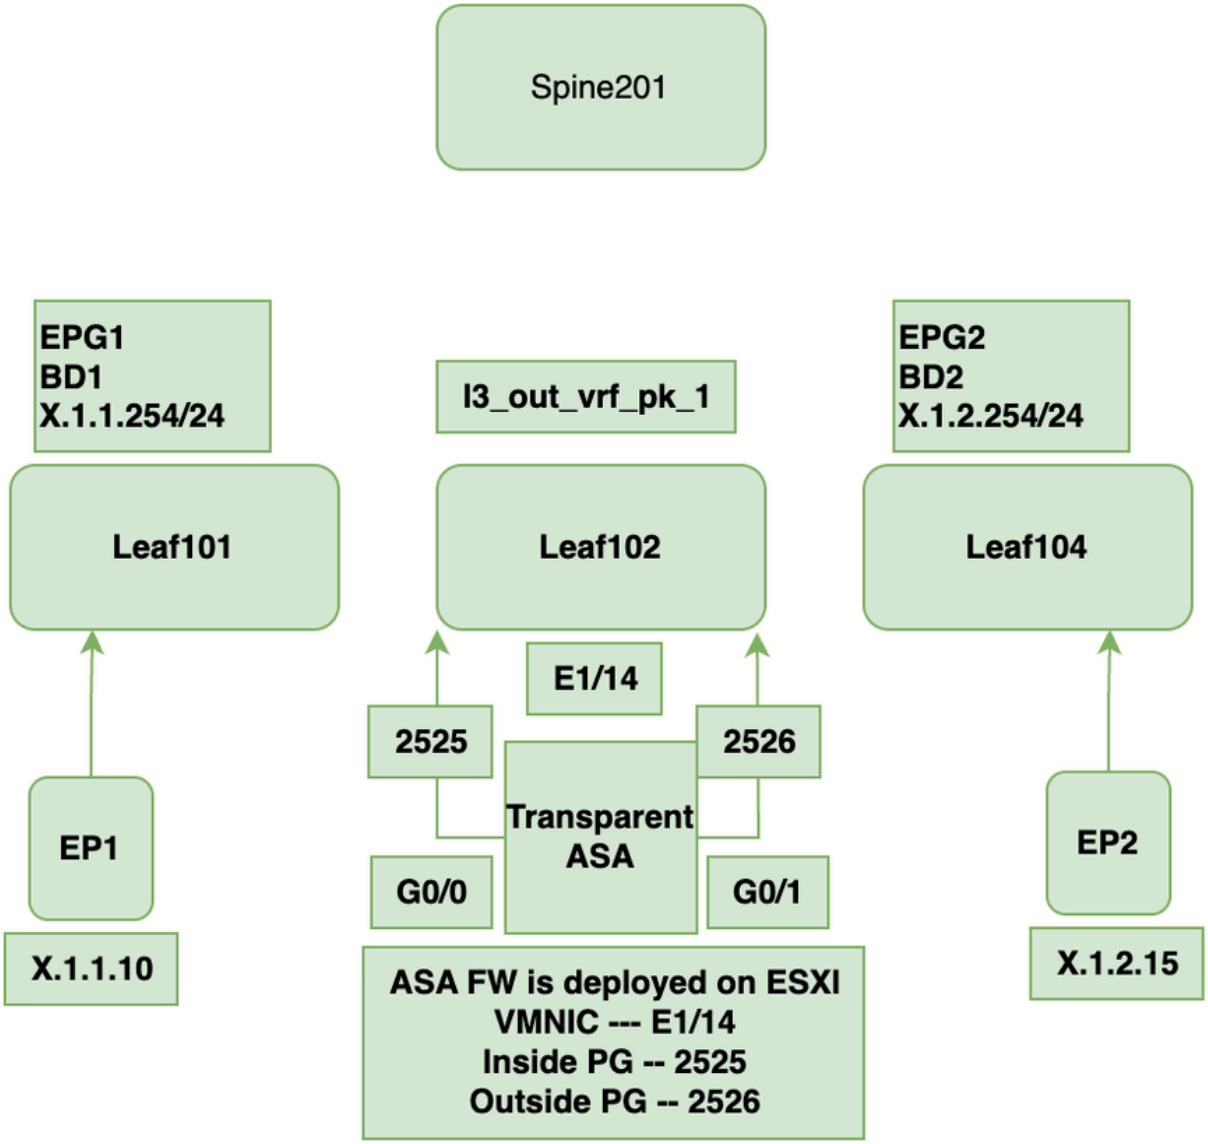

Topology

EPG1 and EPG2 configuration is not shown in this document, it must be connfigured before hand and endpoint must be learnt.

1. Validate EPG1 haș endpoint X.1.1.10 learned (Node 101).

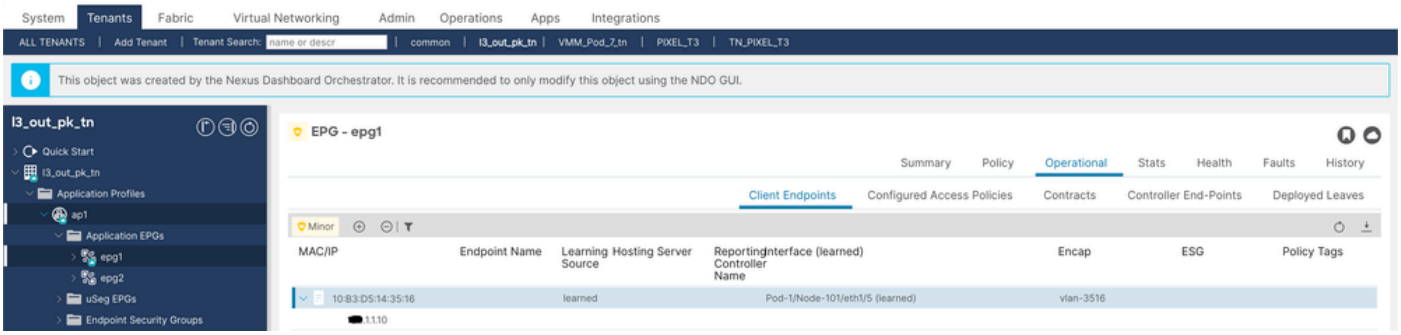

Client Endpoints

2. Contract abc is consumed by EPG1.

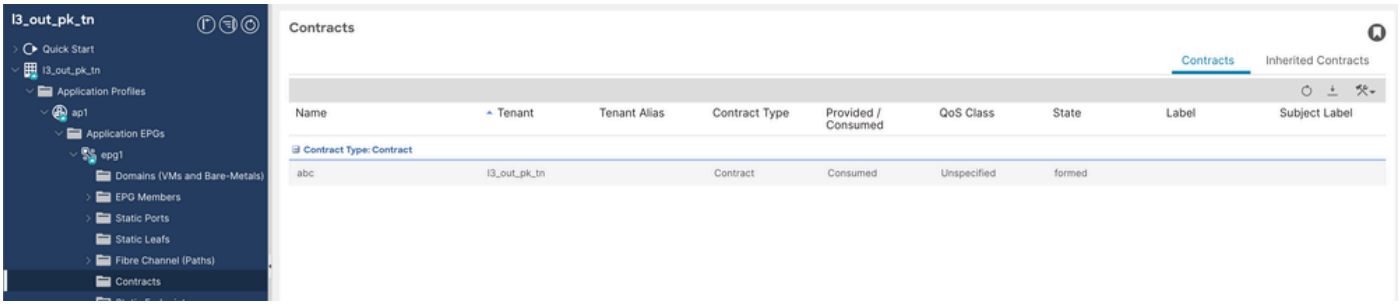

Consumed Contract

3. Validate EPG2 has endpoint X.1.2.15learned (Node 104).

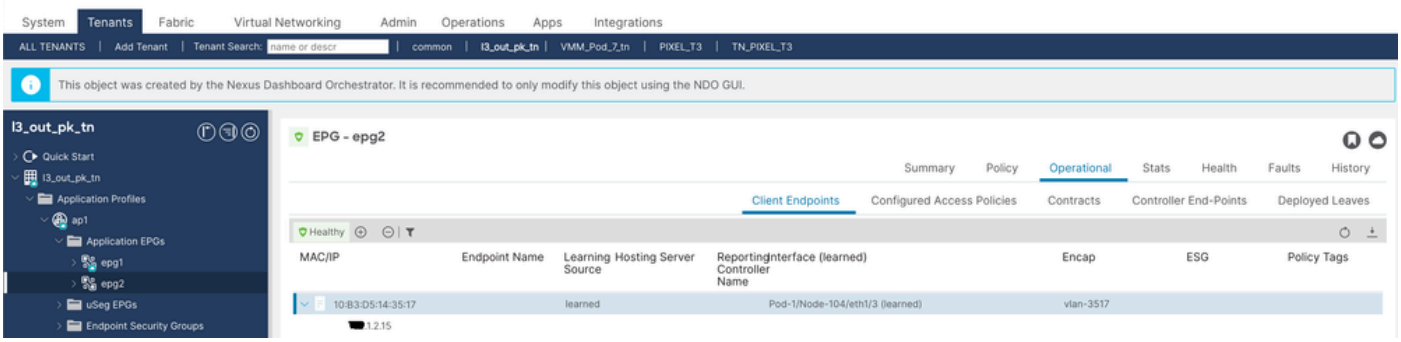

Client Endpoint

4. Contract abc is provided by EPG2.

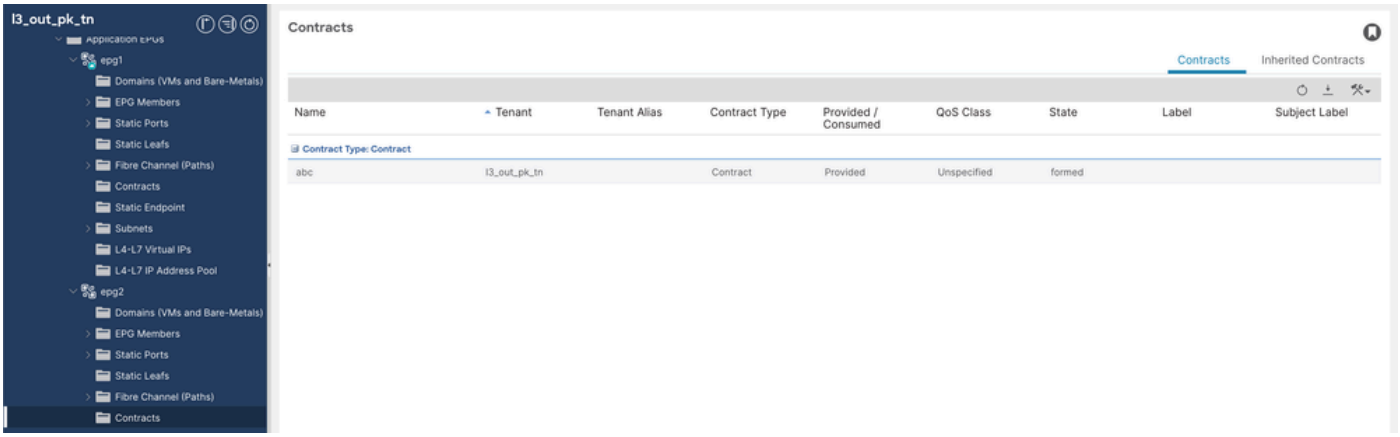

Provided Contract

Why L2 Service Graph is Needed in ACI?

- In Cisco ACI, L4-L7 service devices can be inserted at Layer 3 (L3), Layer 2 (L2), or Layer 1 (L1).
- Layer 3 Service Insertion: The external device (for example, firewall, Intrusion Prevention System (IPS)) makes routing decisions and forwards traffic based on IP addresses.
- Layer 2 Service Insertion: Traffic is forwarded based on MAC addresses without routing involvement. This is useful for transparent firewalls or IPS devices.
- L2 Policy-Based Routing (PBR) is used when inserting an L2 service device, such as an IPS or transparent firewall in ACI.
- The traffic forwarding mechanism remains the same for both L3 and L2 PBR.
- The key difference:
    - L3 PBR: Traffic is redirected to an IP address (device participates in routing).
    - L2 PBR: Traffic is redirected to a MAC address (device operates at Layer 2).
- In L2 PBR, MAC addresses are statically bound to leaf interfaces in order to ensure proper traffic forwarding.

For more information on Active/Stanby or Active/Active L1/L2 PBR use cases, refer to PBR White Paper.

Configuration for L2 Service Graph

Step 1. Configure consumer bd named as con-bd1.

Unicast routing must be enabled, L2 unknown unicast must be set to Hardware proxy and no subnet is required for con and prov Bridge Domains (BDs).

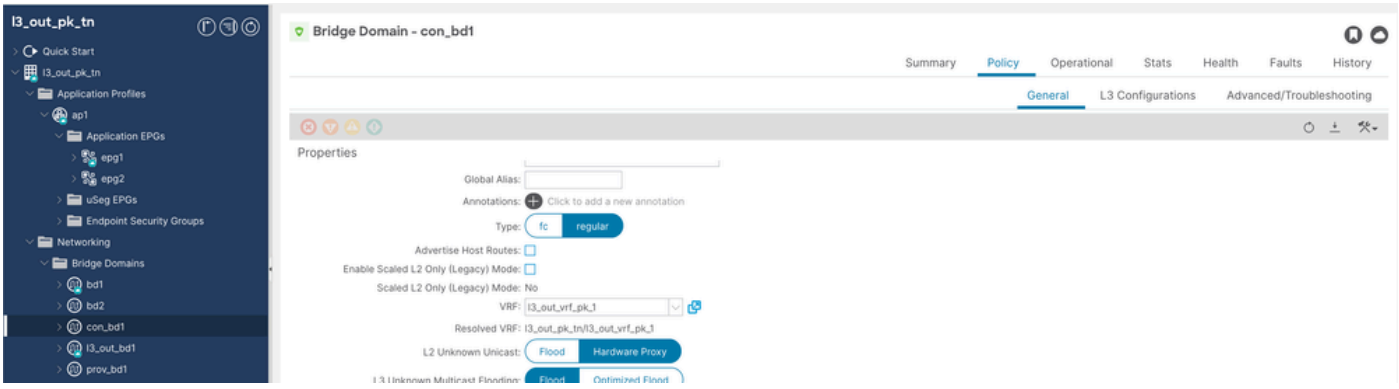

Cons BD Config

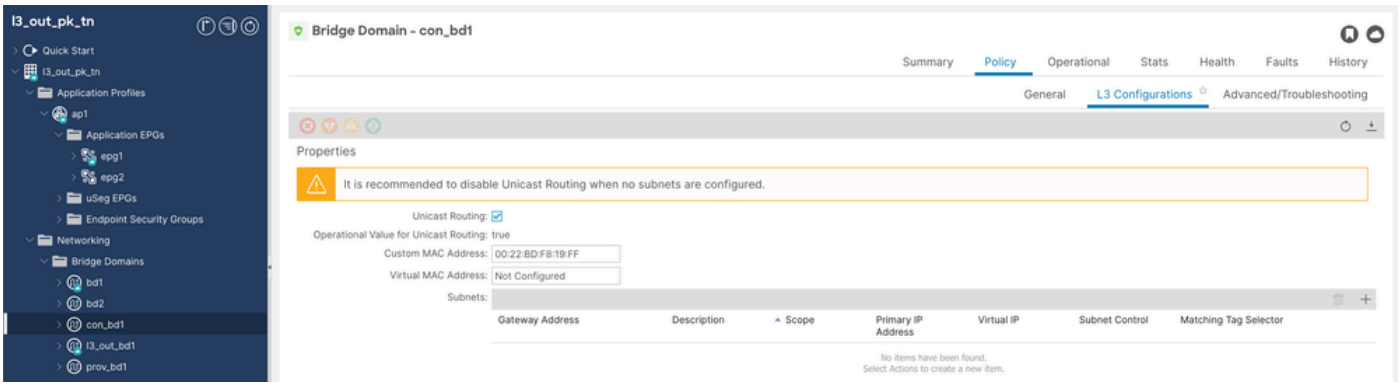

Cons BD Config 2

Step 2. Configure provider bd named as prov-bd1.

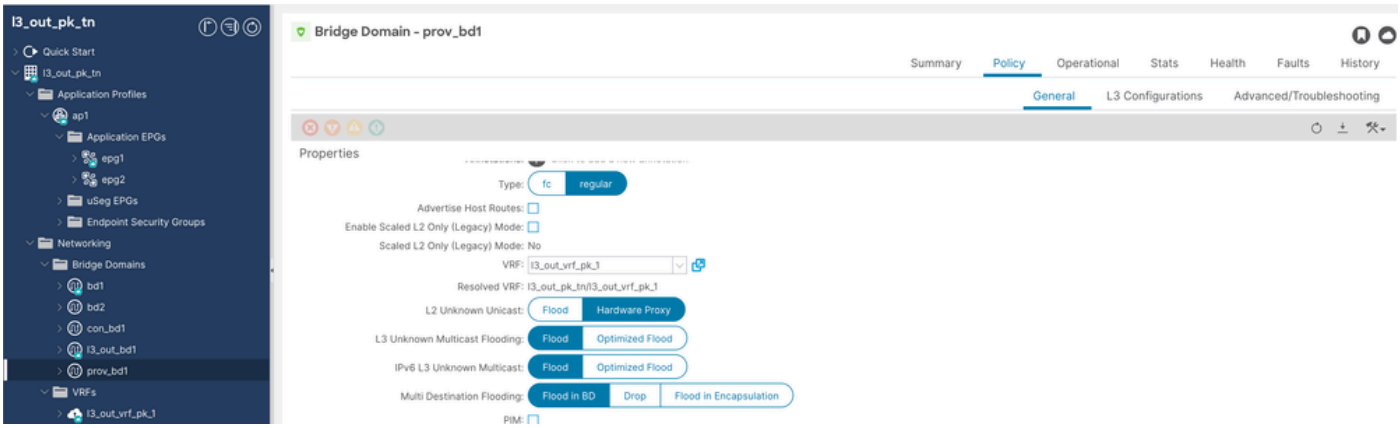

Prov BD Config

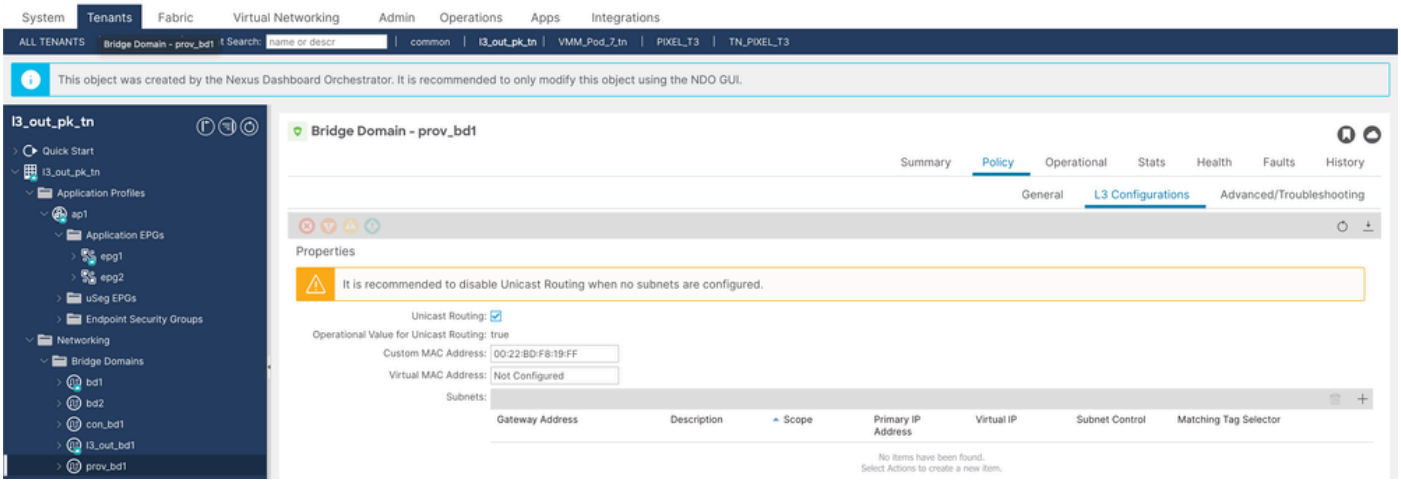

Prov BD Config 2

Step 3. Configure IP Service Level Agreement (SLA) Policy with SLA type L2Ping.

Navigate to**Tenant > Policies > Protocol > IP SLA > IP SLA Monitoring Policies**, then right click and**create policy**.

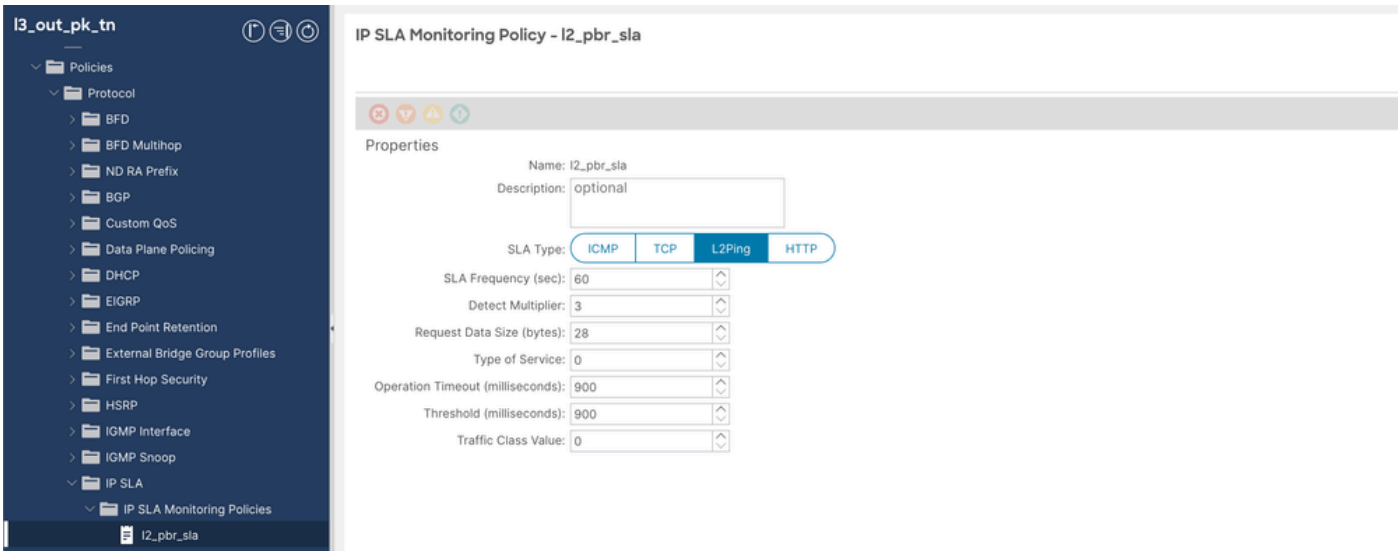

IP SLA Policy

Step 4. Configure L4/L7 device.

Navigate to **Tenant > Services > Devices**, then right click and create L4-L7 device.

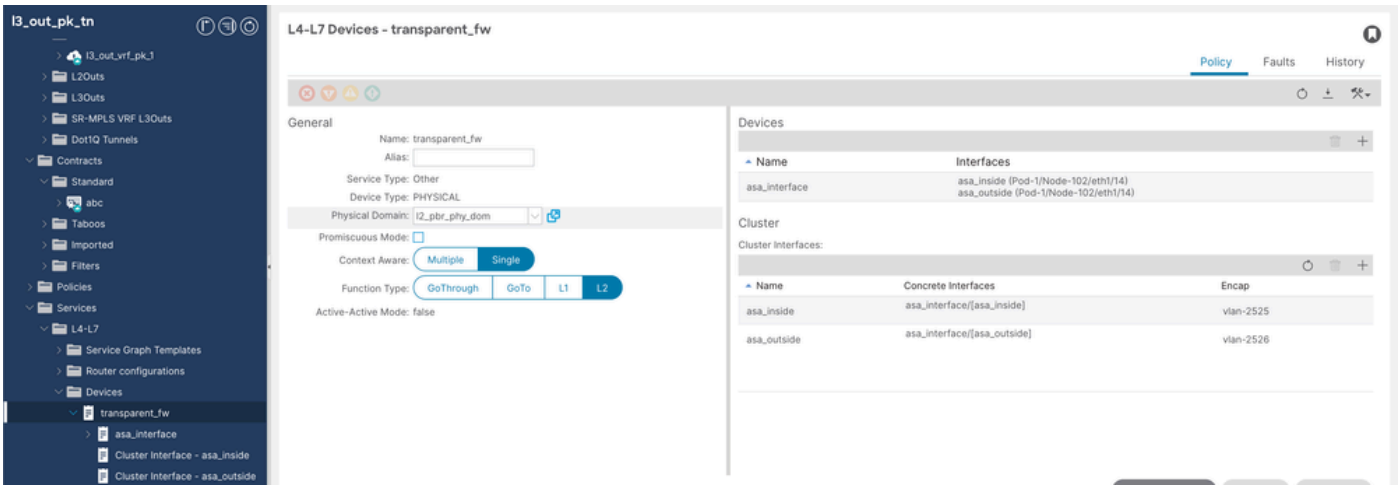

L4-L7 Device

Step 5. Validate Policy Based redirect overview (you can check this after configuring 5a and 5b).

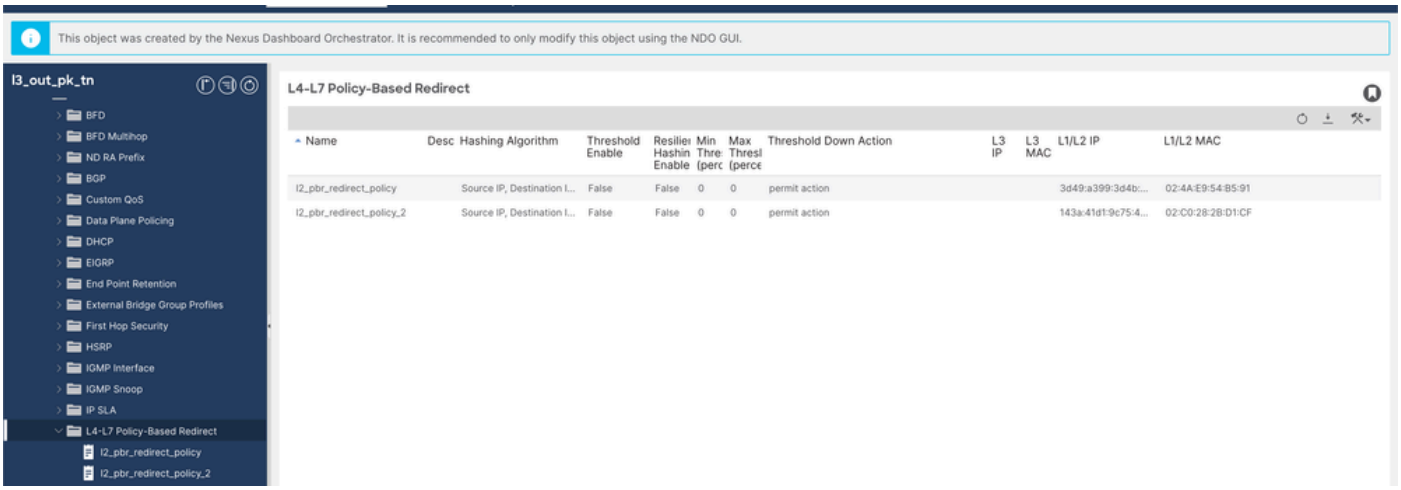

L4-L7 redirect Policy

Step 5.1. Configure L4-L7 policy based redirect policy for Adaptive Security Appliance (ASA) inside interface (no need to specify MAC or IP, it is populated by APIC itself).

Navigate to **Tenant > Policies > Protocol > L4-L7 Policy based redirect**, then right click and **create policy**.

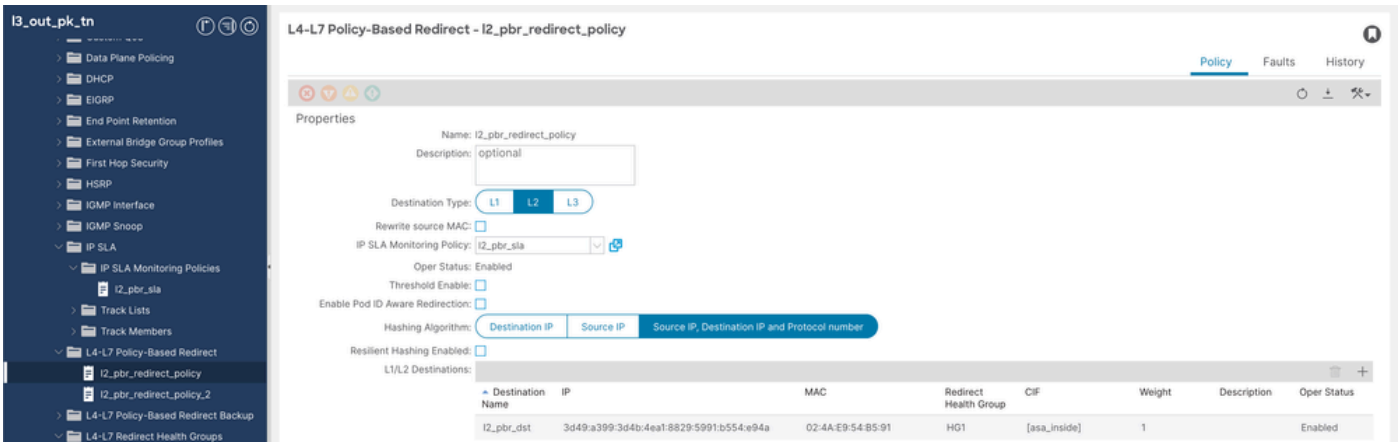

L4-L7 redirect policy config

Step 5.2. Configure L4-L7 policy based redirect policy for ASA outside interface (no need to specify MAC or IP, it is populated by APIC itself).

Navigate to **Tenant > Policies > Protocol > L4-L7 Policy based redirect**, then right click and **create policy**.

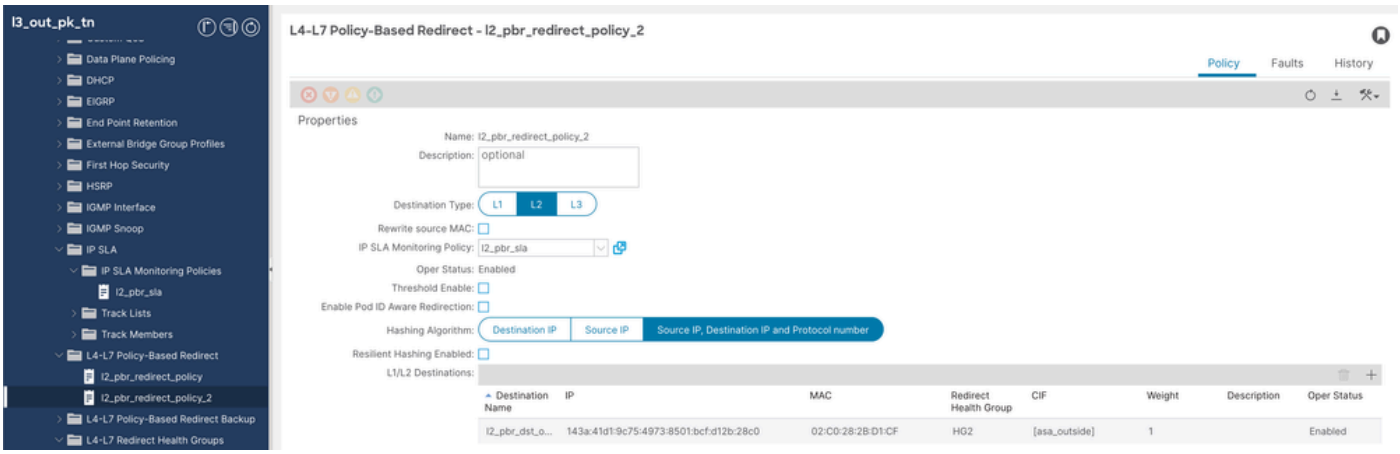

L4-L7 redirect policy config 2

Step 6. Configure Service graph template.

Navigate to **Tenant > Services > Service Graph Template**, then right click and create L4-L7 Service graph template.

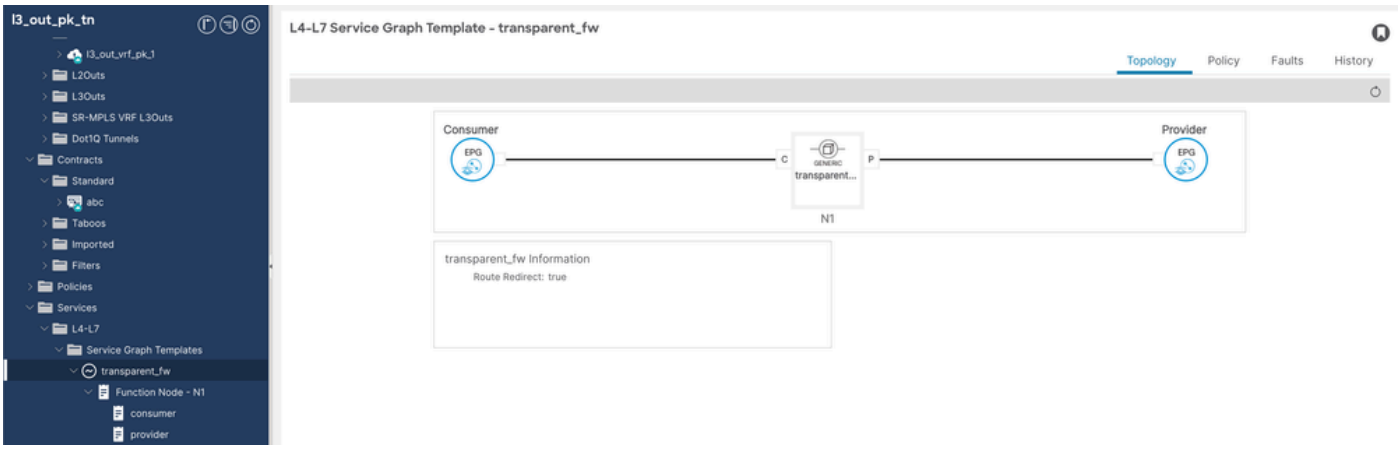

Service Graph Config

Step 7. Configure device selection policy.

Navigate to **Tenant > Services > Device Selection Policy**, then right click and create **Device Selection Policy**.

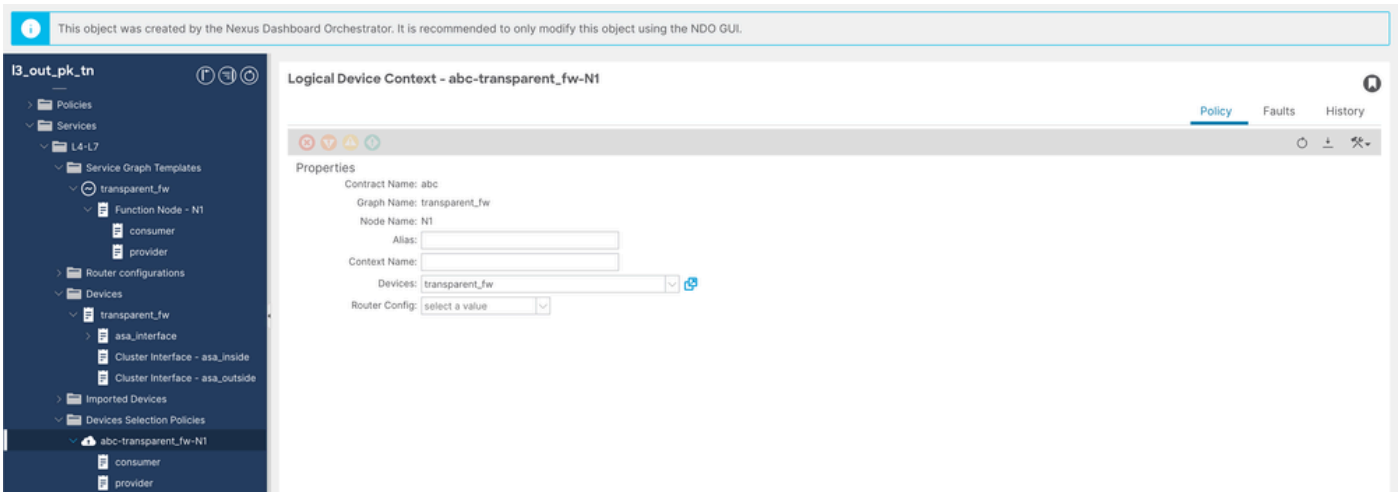

Service Graph Config 2

++ Consumer Logical Interface Context

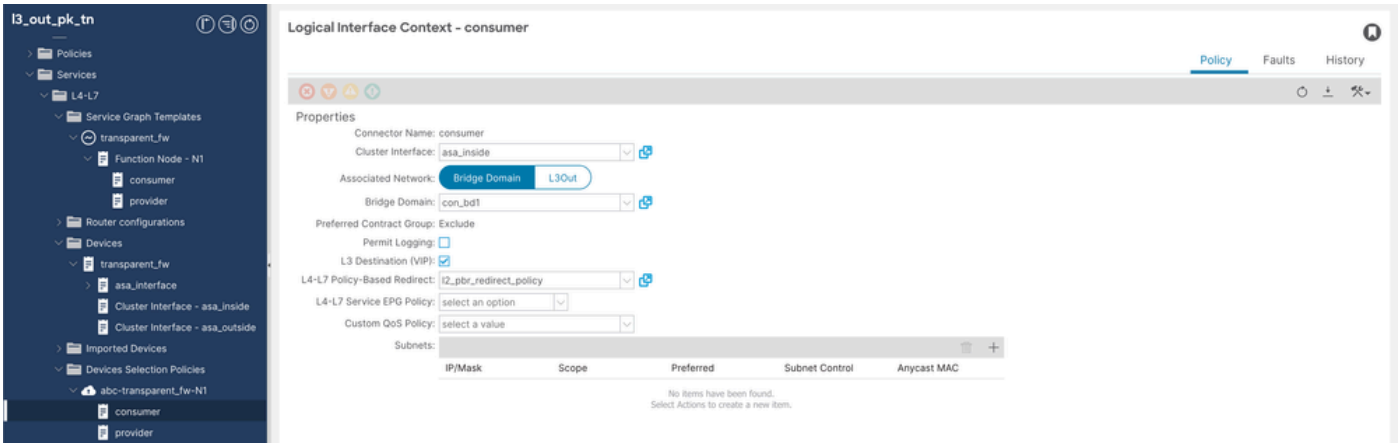

Device Selection Policy Consumer Config

++ Provider Logical Interface Context

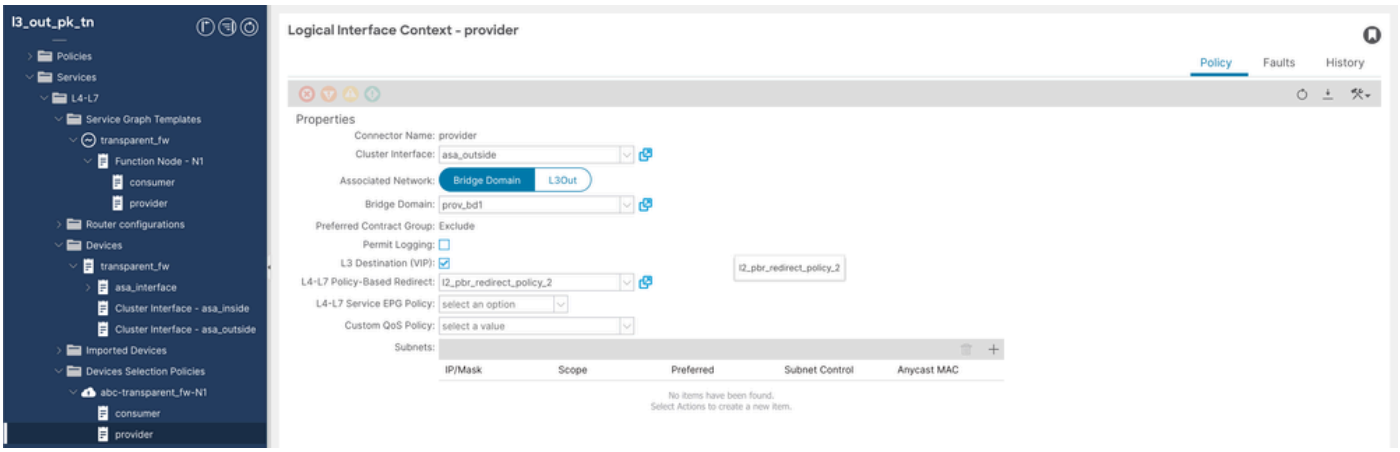

Device Selection Policy Provider config

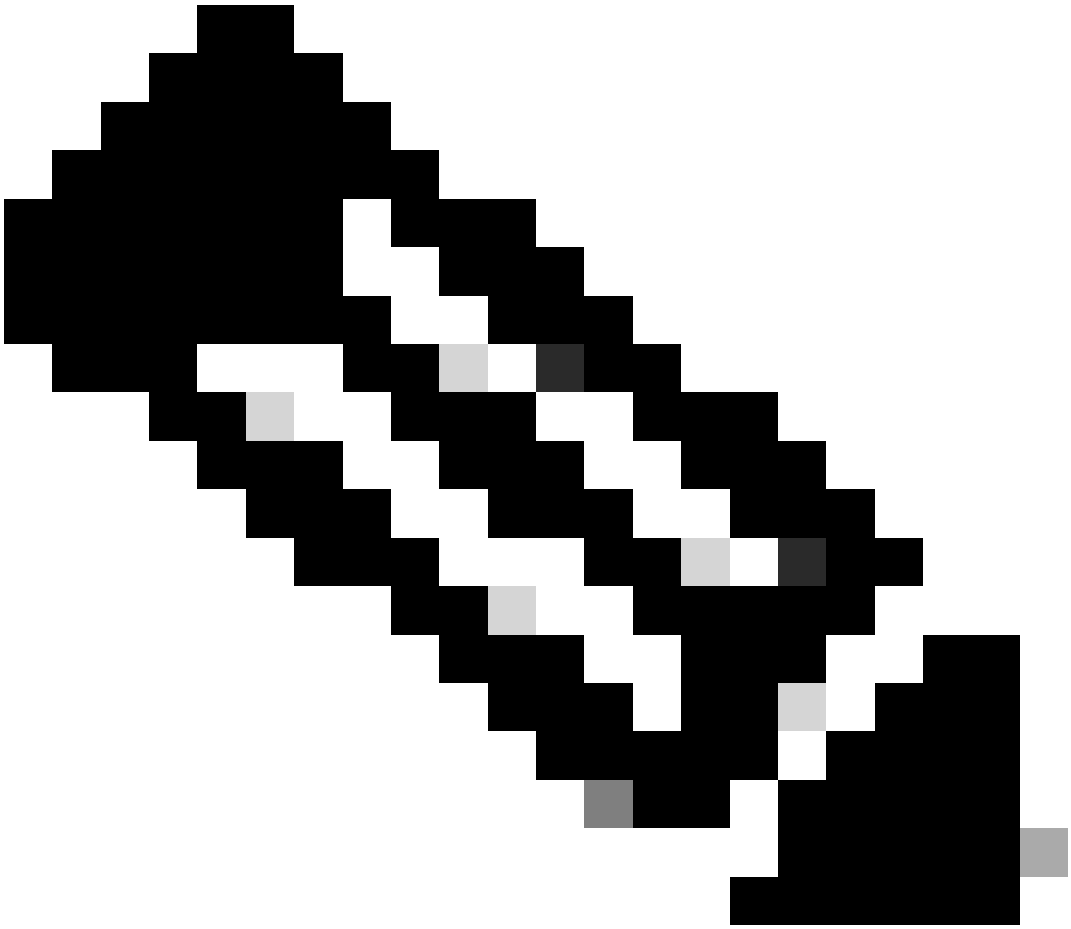

Note: Device selection policy if going to be automatically created incase you are going to use Apply Service Graph Option.

Step 8. Apply PBR in order to contract abc subject.

Navigate to **Tenant > Contract > Contract Subject > L4-L7 Service Graph > transparent_fw**.

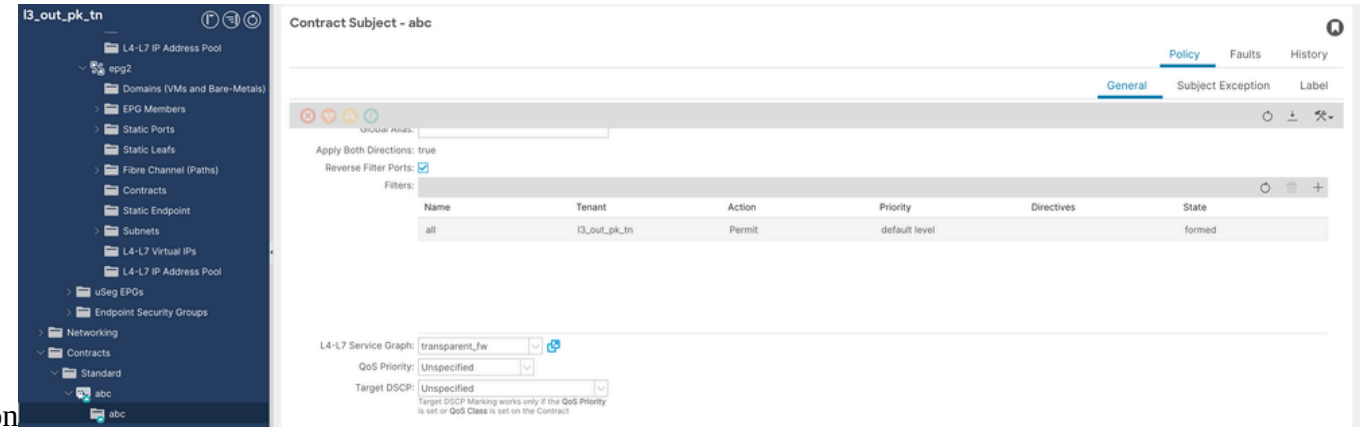

Contract Config

Step 9. If successfully deployed, validate under Deployed instance graph (look for state).

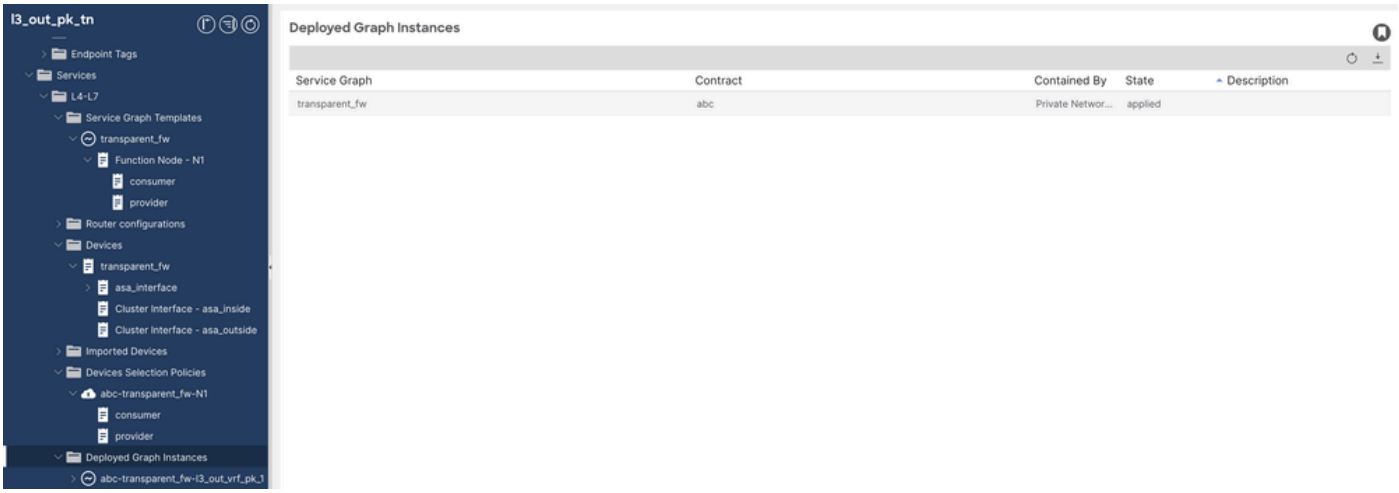

Service Graph Validation

++ Validate cluster interfaces, encap VLANs and function connector class IDs.

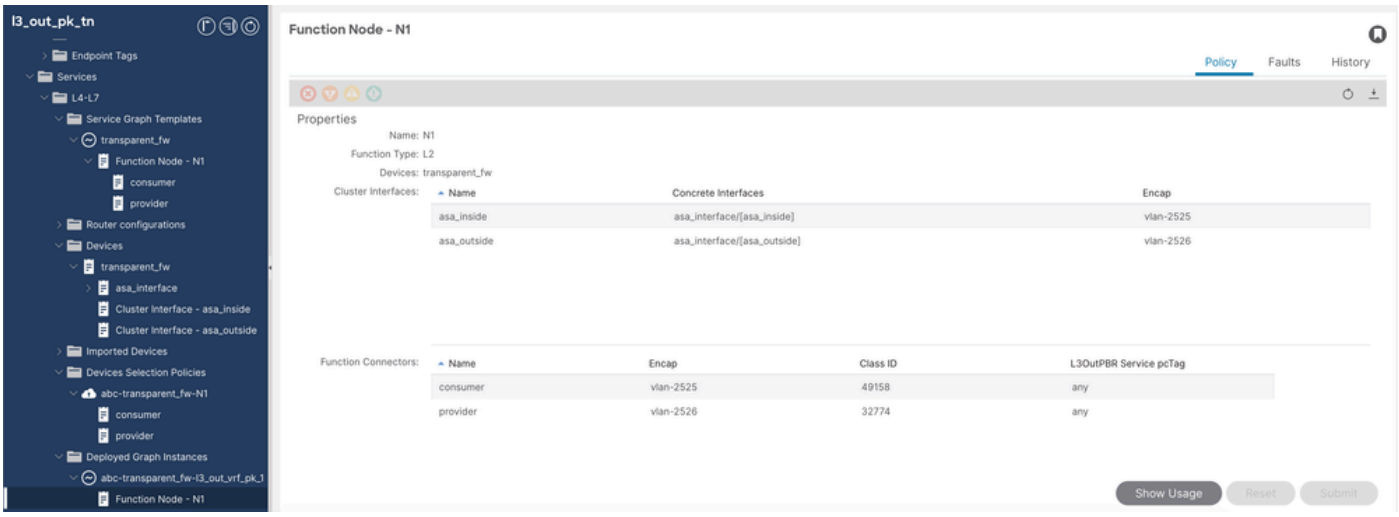

Service Graph Validation 2

Validate L2 PBR Traffic on ASA

Secure Shell (SSH) from Src endpoint to dst endpoint you can see in Conn table entry on ASA.

ASA Validation

Verify L2 PBR on Leaf

1. VLAN programming on Leaf Node 102.

<#root>

```
PBR vlan 2525 and 2526 will get programmed on leaf node 102 and mac addresses will be statically tied t

bgl-aci07-apic100#
fabric 102 show endpoint


--------------------------------------------------------------------
Node 102 (bgl-aci07-leaf2)
--------------------------------------------------------------------
Legend:
S - static                  s - arp              L - local            O - peer-attached
V - vpc-attached            a - local-aged       p - peer-aged        M - span
B - bounce                  H - vtep             R - peer-attached-rl D - bounce-to-proxy
E - shared-service          m - svc-mgr
+----------------------------------+---------------+-----------------+--------------+-------------+
      VLAN/                           Encap           MAC Address       MAC Info/       Interface
      Domain                          VLAN            IP Address        IP Info
+----------------------------------+---------------+-----------------+--------------+-------------+
28/l3_out_pk_tn:l3_out_vrf_pk_1           vlan-2525    024a.e954.b591 LS                    eth1/14
1/l3_out_pk_tn:l3_out_vrf_pk_1            vlan-2526    02c0.282b.d1cf LS                    eth1/14
```

2. Redirect policy and zoning-rule on consumer (101) and provider (104) node.

<#root>

```
++ Redirect policy on consumer node

bgl-aci07-apic100#
fabric 101 show service redir info


--------------------------------------------------------------------
Node 101 (bgl-aci07-leaf1)
--------------------------------------------------------------------
================================================================================================================
LEGEND
TL: Threshold(Low)  |  TH: Threshold(High) |  HP: HashProfile  |  HG: HealthGrp  | BAC: Backup-Dest |
================================================================================================================
List of Dest Groups
GrpID Name            destination                                                HG-name
===== ====            ===========                                                ==============
```

```
7     destgrp-7       dest-[3d49:a399:3d4b:4ea1:8829:5991:b554:e94a]-[vxlan-2228224] l3_out_pk_tn::HG1
8     destgrp-8       dest-[143a:41d1:9c75:4973:8501:bcf:d12b:28c0]-[vxlan-2228224]  l3_out_pk_tn::HG2

List of destinations
Name                                                                bdVnid          vMac                vrf
====                                                                ======          ====                ===
dest-[3d49:a399:3d4b:4ea1:8829:5991:b554:e94a]-[vxlan-2228224] vxlan-16744328  02:4A:E9:54:B5:91    l3_
dest-[143a:41d1:9c75:4973:8501:bcf:d12b:28c0]-[vxlan-2228224]  vxlan-16056296  02:C0:28:2B:D1:CF    l3_

List of Health Groups
HG-Name                                     HG-OperSt  HG-Dest
=======                                     =========  =======
l3_out_pk_tn::HG1                           enabled    dest-[3d49:a399:3d4b:4ea1:8829:5991:b554:e94a]-[v
l3_out_pk_tn::HG2                           enabled    dest-[143a:41d1:9c75:4973:8501:bcf:d12b:28c0]-[vx

List of Backup Destinations
Name                                                                primaryDestName
====                                                                ==============
List of AclRules
AclRuleVnid     DestGroup       OperSt   OperStQual
===========     =========       ======   ==========

++ Zoning rule on consumer Node

bgl-aci07-apic100#
```

`fabric 101 show zoning-rule  | grep redir`

```
|   4228  | 32771  | 49157  | default  |     bi-dir     | enabled | 2228224  |
|   4231  | 49157  | 32771  | default  | uni-dir-ignore | enabled | 2228224  |
|   4230  | 32771  |   15   | default  |    uni-dir     | enabled | 2228224  |
|   4229  | 16386  | 32771  | default  |    uni-dir     | enabled | 2228224  |
```

<#root>

```
++ Redirect Policy on Provider Node
bgl-aci07-apic100#
```

`fabric 104 show service redir info`

```
------------------------------------------------------------------
Node 104 (bgl-aci07-leaf4)
------------------------------------------------------------------
================================================================================================================
LEGEND
TL: Threshold(Low)  |  TH: Threshold(High) |  HP: HashProfile  |  HG: HealthGrp  | BAC: Backup-Dest |  T
================================================================================================================
List of Dest Groups
GrpID Name            destination                                                   HG-name
===== ====            ===========                                                   ==============
3     destgrp-3       dest-[3d49:a399:3d4b:4ea1:8829:5991:b554:e94a]-[vxlan-2228224] l3_out_pk_tn::HG1
4     destgrp-4       dest-[143a:41d1:9c75:4973:8501:bcf:d12b:28c0]-[vxlan-2228224]  l3_out_pk_tn::HG2

List of destinations
Name                                                                bdVnid          vMac                vrf
====                                                                ======          ====                ===
dest-[3d49:a399:3d4b:4ea1:8829:5991:b554:e94a]-[vxlan-2228224] vxlan-16744328  02:4A:E9:54:B5:91    l3_
dest-[143a:41d1:9c75:4973:8501:bcf:d12b:28c0]-[vxlan-2228224]  vxlan-16056296  02:C0:28:2B:D1:CF    l3_
```

```
List of Health Groups
HG-Name                                    HG-OperSt  HG-Dest
=======                                    =========  =======
l3_out_pk_tn::HG1                          enabled    dest-[3d49:a399:3d4b:4ea1:8829:5991:b554:e94a]-[v
l3_out_pk_tn::HG2                          enabled    dest-[143a:41d1:9c75:4973:8501:bcf:d12b:28c0]-[vx
List of Backup Destinations
Name                                                          primaryDestName
====                                                          ===============

++ Zoning rule on provider node
bgl-aci07-apic100#
```

fabric 104 show zoning-rule | grep redir

```
|   4220  | 32771  | 49157  | default  |     bi-dir     | enabled | 2228224  |
|   4221  | 49157  | 32771  | default  | uni-dir-ignore | enabled | 2228224  |
```

Faults Seen in Case L2Ping Failed

In case L2pings are failing across PBR device, you will observe PBR is still in deployed state and faults F4203, F2833, and F2911 raised stating track/health group are down.

Capturing L2 Pings

You can capture L2Pings using **tcpdump** on tahoe interface in order to know if they are correctly sent and recieved. If you only see CPU transmit sent and not recieved, then the earlier mentioned faults are expected and you must further troubleshoot on ASA why they are dropped (refer to ASA configuration section).

<#root>

```
Capturing L2Pings using tcpdump on PBR Node 102

bgl-aci07-leaf2#
```

tcpdump -i tahoe0 -w /data/techsupport/l2_pbr1.pcap

```
tcpdump: listening on tahoe0, link-type EN10MB (Ethernet), capture size 262144 bytes
^C4858 packets captured
4875 packets received by filter
0 packets dropped by kernel
```

In order to deocde the tcpdump

cat /data/techsupport/l2_pbr1.pcap | knet_parser.py --decode tahoe --pcap | less

```
** Search for mac 00ab.8752.3100

++ CPU transmit packets
Frame 505
Time: 2024-10-29T05:55:28.707136+00:00
Header: ieth
```

CPU Transmit

sup_tx:1, ttl_bypass:0, opcode:0x0, bd:0x207, outer_bd:0x0, dl:0, span:0, traceroute:0, tclass:5
src_idx:0x0, src_chip:0x0, src_port:0x0, src_is_tunnel:0, src_is_peer:0
dst_idx:0x0, dst_chip:0x0, dst_port:0x0, dst_is_tunnel:0
Len: 72
Eth:

00ab.8752.3100 > 024a.e954.b591

, len/

ethertype:0x721

Frame 506
Time: 2024-10-29T05:55:28.707297+00:00
Header: ieth CPU Transmit
sup_tx:1, ttl_bypass:0, opcode:0x0, bd:0x208, outer_bd:0x0, dl:0, span:0, traceroute:0, tclass:5
src_idx:0x0, src_chip:0x0, src_port:0x0, src_is_tunnel:0, src_is_peer:0
dst_idx:0x0, dst_chip:0x0, dst_port:0x0, dst_is_tunnel:0
Len: 72
Eth:

00ab.8752.3100 > 02c0.282b.d1cf

, len/

ethertype:0x721

++CPU recived packets

Frame 509
Time: 2024-10-10T20:16:37.580855+00:00
Header: ieth_extn

CPU Receive

sup_qnum:0x33, sup_code:0x4d, istack:

ISTACK_SUP_CODE_PBR_TRACK_REFRESH

(0x4d)
Header: ieth
sup_tx:0, ttl_bypass:0, opcode:0x0, bd:0x209, outer_bd:0x2, dl:0, span:0, traceroute:0, tclass:0
src_idx:0x32, src_chip:0x0, src_port:0x6, src_is_tunnel:0, src_is_peer:0
dst_idx:0x1, dst_chip:0x0, dst_port:0x3d, dst_is_tunnel:0
Len: 76
Eth:

00ab.8752.3100 > 024a.e954.b591

, len/ethertype:0x8100(802.1q)
802.1q:

vlan:2526

, cos:0, len/

ethertype:0x721

Frame 510
Time: 2024-10-10T20:16:37.580891+00:00

Header: ieth_extn

CPU Receive

sup_qnum:0x33, sup_code:0x4d, istack:

ISTACK_SUP_CODE_PBR_TRACK_REFRESH(0x4d)

Header: ieth
sup_tx:0, ttl_bypass:0, opcode:0x0, bd:0x20a, outer_bd:0x2, dl:0, span:0, traceroute:0, tclass:0
src_idx:0x32, src_chip:0x0, src_port:0x6, src_is_tunnel:0, src_is_peer:0
dst_idx:0x1, dst_chip:0x0, dst_port:0x3d, dst_is_tunnel:0
Len: 76
Eth:

00ab.8752.3100 > 02c0.282b.d1cf

, len/ethertype:0x8100(802.1q)
802.1q:

vlan:2525

, cos:0, len/

ethertype:0x721

Traffic Flow From Src to Dst Endpoint

<#root>

++ Endpoint X.1.1.10 want to send traffic to X.1.2.15
++ If destination is not learned on consumer/source leaf, PBR will be performed on destination leaf
++ For this case we are assuming endpoint X.1.2.15 is learned on Leaf 101 so PBR/Redirection will be pe

bgl-aci07-apic100#

fabric 101 show endpoint

```
--------------------------------------------------------------------
Node 101 (bgl-aci07-leaf1)
--------------------------------------------------------------------
Legend:
S - static           s - arp              L - local            O - peer-attached
V - vpc-attached     a - local-aged       p - peer-aged        M - span
B - bounce           H - vtep             R - peer-attached-rl D - bounce-to-proxy
E - shared-service   m - svc-mgr
+----------------------------------+--------------+-----------------+-------------+-------------+
      VLAN/                          Encap          MAC Address       MAC Info/       Interface
      Domain                         VLAN           IP Address        IP Info
+----------------------------------+--------------+-----------------+-------------+-------------+
l3_out_pk_tn:l3_out_vrf_pk_1                            X.1.2.15                      tunnel6 ===> |
17                                   vlan-3516    10b3.d514.3516 L                    eth1/5 ===>
l3_out_pk_tn:l3_out_vrf_pk_1         vlan-3516        X.1.1.10 L                     eth1/5
```

++ EPM entry to get the PC TAG
bgl-aci07-apic100#

fabric 101 show system internal epm endpoint ip X.1.1.10

```
------------------------------------------------------------------
Node 101 (bgl-aci07-leaf1)
------------------------------------------------------------------
MAC : 10b3.d514.3516 ::: Num IPs : 1
IP# 0 : X.1.1.10 ::: IP# 0 flags :  ::: l3-sw-hit: No
Vlan id : 17 ::: Vlan vnid : 11792 ::: VRF name : l3_out_pk_tn:l3_out_vrf_pk_1
BD vnid : 16744307 ::: VRF vnid : 2228224
Phy If : 0x1a004000 ::: Tunnel If : 0
Interface : Ethernet1/5
Flags : 0x80005c04 ::: sclass :
```

32771

```
 ::: Ref count : 5 ==> sclass
EP Create Timestamp : 10/11/2024 09:15:44.430334
EP Update Timestamp : 10/29/2024 10:45:35.458416
EP Flags : local|IP|MAC|host-tracked|sclass|timer|

bgl-aci07-apic100#
```

fabric 101 show system internal epm endpoint ip X.1.2.15

```
------------------------------------------------------------------
Node 101 (bgl-aci07-leaf1)
------------------------------------------------------------------
MAC : 0000.0000.0000 ::: Num IPs : 1
IP# 0 : X.1.2.15 ::: IP# 0 flags :  ::: l3-sw-hit: No
Vlan id : 0 ::: Vlan vnid : 0 ::: VRF name : l3_out_pk_tn:l3_out_vrf_pk_1
BD vnid : 0 ::: VRF vnid : 2228224
Phy If : 0 ::: Tunnel If : 0x18010006
Interface : Tunnel6
Flags : 0x80004400 ::: sclass :
```

49157

```
 ::: Ref count : 3 ==> sclass
EP Create Timestamp : 10/29/2024 10:38:34.949150
EP Update Timestamp : 10/29/2024 10:45:55.571786
EP Flags : IP|sclass|timer|

++ Traffic will be redirected based on redir(destgrp-7)
bgl-aci07-apic100#
```

fabric 101 show zoning-rule src-epg 32771 dst-epg 49157

```
------------------------------------------------------------------
Node 101 (bgl-aci07-leaf1)
------------------------------------------------------------------
+---------+--------+--------+----------+--------+---------+---------+------+------------------+--------
| Rule ID | SrcEPG | DstEPG | FilterID |  Dir   |  operSt |  Scope  | Name |      Action      |    Prio
+---------+--------+--------+----------+--------+---------+---------+------+------------------+--------
|   4228  | 32771  | 49157  | default  | bi-dir | enabled | 2228224 |      | redir(destgrp-7) | src_dst
+---------+--------+--------+----------+--------+---------+---------+------+------------------+--------

++ Based on redirect policy traffic will be redirected to mac
```

02:4A:E9:54:B5:91

```
bgl-aci07-apic100#
```

fabric 101 show service redir info

```
------------------------------------------------------------
Node 101 (bgl-aci07-leaf1)
------------------------------------------------------------
================================================================================================================
LEGEND
TL: Threshold(Low)  |  TH: Threshold(High) |  HP: HashProfile  |  HG: HealthGrp  | BAC: Backup-Dest |  T
================================================================================================================
List of Dest Groups
GrpID Name            destination                                                    HG-name
===== ====            ===========                                                    ==============
7     destgrp-7       dest-[3d49:a399:3d4b:4ea1:8829:5991:b554:e94a]-[vxlan-2228224] l3_out_pk_tn::HG1
List of destinations
Name                                                          bdVnid          vMac                  vrf
====                                                          ======          ====                  ====
dest-[3d49:a399:3d4b:4ea1:8829:5991:b554:e94a]-[vxlan-2228224] vxlan-16744328
```

02:4A:E9:54:B5:91

```
    l3_out_pk_tn:l3_out_vrf_pk_1 enabled    no-oper-dest    l3_out_pk_tn::HG
1

++ PBR mac addresses are never learnt remotely as IP/MAC learning is disabled for PBR BD
++ PBR mac addresses are statically binded to interfaces where L4/L7 device is connected and reported t
++ Traffic will be forwarded to SPINE PROXY
++ Spine has an COOP entry for 02:4A:E9:54:B5:91

bgl-aci07-apic100#
```

fabric 201 show coop internal info repo ep key 16744328 02:4A:E9:54:B5:91

```
------------------------------------------------------------
Node 201 (bgl-aci07-spine1)
------------------------------------------------------------
Repo Hdr Checksum : 49503
Repo Hdr record timestamp : 10 29 2024 10:15:07 658496921
Repo Hdr last pub timestamp : 10 29 2024 10:15:07 661679296
Repo Hdr last dampen timestamp : 01 01 1970 00:00:00 0
Repo Hdr dampen penalty : 0
Repo Hdr flags : IN_OBJ ACTIVE
EP bd vnid : 16744328
EP mac :
```

02:4A:E9:54:B5:91

```
 <<<<========== ASA MAC
flags : 0x480
repo flags : 0x102
Vrf vnid : 2228224
PcTag : 0x100c006
EVPN Seq no : 0
Remote publish timestamp: 01 01 1970 00:00:00 0
Snapshot timestamp: 10 29 2024 10:15:07 658496921
Tunnel nh : 10.0.144.66
MAC Tunnel  : 10.0.144.66
IPv4 Tunnel : 10.0.144.66
IPv6 Tunnel : 10.0.144.66
ETEP Tunnel : 0.0.0.0
num of active ipv4 addresses : 0
num of anycast ipv4 addresses : 0
num of ipv4 addresses : 0
num of active ipv6 addresses : 0
num of anycast ipv6 addresses : 0
```

```
num of ipv6 addresses : 0
Primary Path:
Current published TEP :
```

10.0.144.66

```
Backup Path:
BackupTunnel nh : 0.0.0.0
Current Backup (publisher_id): 0.0.0.0
Anycast_flags : 0
Current citizen (publisher_id): 10.0.144.66
Previous citizen : 10.0.144.66
Prev to Previous citizen : 10.0.144.66
Synthetic Flags : 0x5
Synthetic Vrf : 411
Synthetic IP : X.X.83.223
Tunnel EP entry: 0x7f20900167a8
Backup Tunnel EP entry: (nil)
TX Status: COOP_TX_DONE\
Damp penalty: 0
Damp status: NORMAL
Exp status: 0
Exp timestamp: 01 01 1970 00:00:00 0
Hash: 3209430840 owner: 10.0.144.65

++ Spine will forward this to PBR Leaf Node 102 based on COOP entry
++ PBR Leaf Node will forward this to ASA FW on interface E1/14
++ ASA FW will forward the traffic based on mac address table and send it back to PBR Leaf Node 102
++ PBR Leaf Node will look for Dst IP in the traffic and route it to Leaf 104 if remote endpoint entry
++ Leaf 104 will get this traffic forwarded to actual EP X.1.2.15 (Leaf4 does not learn the client IP a
```

ASA Configuration

Step 1. Interface Conifguration.

<#root>

```
ASA(config)#
```

show running-config interface

```
!
interface GigabitEthernet0/0
 bridge-group 1
 nameif inside
 security-level 100
!
interface GigabitEthernet0/1
 bridge-group 1
 nameif outside
 security-level 0
!
interface BVI1
ip address 192.168.100.1 255.255.255.0 ==> In case BVI IP is not defined ASA will not switch the packet
!
```

Step 2. MAC learning must be disabled.

<#root>

```
ASA(config)#
```

show run mac-learn

```
mac-learn inside disable
mac-learn outside disable
```

PBR:

Step 3. Static mac-address-table for PBR Mac.

<#root>

```
The mac statically binded to inside interface is the PBR mac generated by provider and vice versa
ASA(config)#
```

show run mac-address-table

```
mac-address-table static outside 024a.e954.b591
mac-address-table static inside 02c0.282b.d1cf
```

Step 4. Configure Access Control List (ACL) in order to pass L2pings.

<#root>

```
ASA(config)#
```

show access-list

```
access-list L2_PBR ethertype permit 721

ASA(config)# show run access-group
access-group L2_PBR in interface inside
access-group L2_PBR in interface outside
```