

Case Study: L3 Multicast in the ACI Fabric

Contents

[Introduction](#)

[Design Requirements](#)

[Solution](#)

[Configuration](#)

[Verify](#)

Introduction

Layer 3 Multicast routing is supported in the ACI fabric starting with release 2.0 onwards and requires EX switches (i.e N9K-C93180YC-EX). Prior to release 2.0, ACI only supported L2 multicast within the bridge domain. This is still a valid option in 2.0 and can be used for non-EX switches.

In ACI release 2.0, supported multicast routing features include: PIM ASM, PIM SSM, Static RP, PIM Auto-RP, and PIM BSR.

In this document, we describe a validated solution for a real customer deployment scenario for L3 multicast routing on ACI fabric. Selected ACI release is 2.1(1h). This release does not support RP on the fabric, thus an external RP is needed for PIM ASM.

Design Requirements

Customer requires an End-to-End solution for L3 multicast routing within and outside the fabric. The deployment scenario has the following requirements:


- One VRF deployed for all tenants.

Note: Multicast requires dedicated L3out for each VRF. If there are multiple VRFs in the fabric, shared L3out is not supported for multicast routing.

- Receivers in fabric with external sources
- Sources in fabric with external receivers
- Sources and receivers in fabric
- Static RP or Auto-RP

Solution

Topology Review



In the topology, there are two major components: the ACI module and the Tiersed LAN Network module. Both modules are connected to core devices via point-to-point L3 links running OSPF and PIM. In ACI module, the external routed network is called L3out-to-Core associated with the VRF common:default. It includes the four links from the two border leaves to the core devices. The tiersed LAN network module, referred as external to the fabric, consists of traditional access layer and aggregation layer with vPC.

L3 multicast flows run across the ACI fabric and the legacy LAN network through the Core layer. For Static RP scenario, RP is deployed on the WAN edge device.

We use Spirent traffic generators (STC) to simulate the internal and external sources and receivers. The Spirent ports are connected to different locations in the ACI module and the Tiersed LAN network module. Receivers send IGMP v2 membership join messages.

Internal Sources attached to Sleaf2: source IPs are 10.101.101.115~124, sending to group addresses: 228.0.0.1~10

Internal Receivers attached to Bleaf1, Cleaf1 and Cleaf2: multicast enabled BDs are BD90, BD91, BD31, BD32, interested groups: 228.0.0.1~10 and 229.0.0.1~10.

External Sources attached to the access layer in LAN network: source IPs are 10.103.103.40~49, sending to group addresses: 229.0.0.1~10.

External Receivers attached to the access layer in LAN network: vlan131, vlan132, interested groups: 228.0.0.1~10.

Configuration

Step 0: Set up RP on the simulated WAN device attached to core, enable PIM sparse mode on

the tiered LAN network devices.

```
!!!! RP configuration

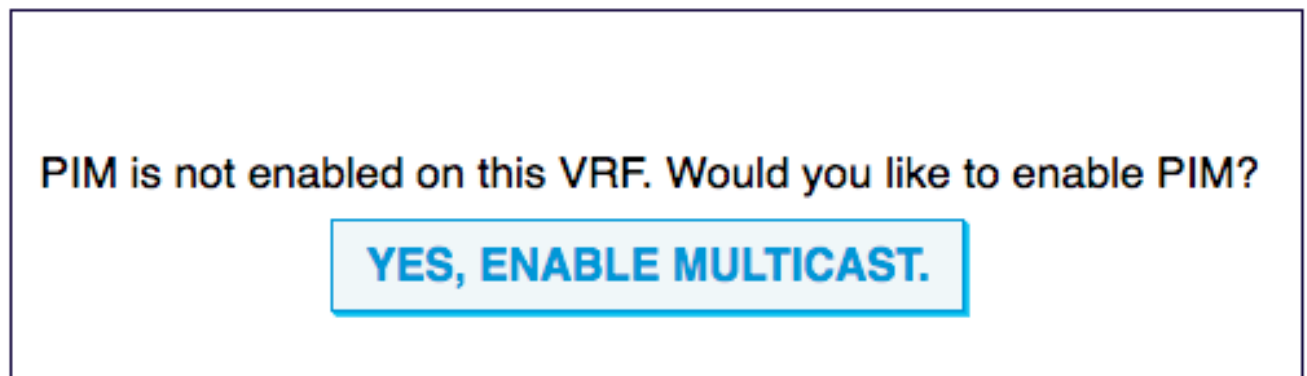
ip pim rp-address 99.99.99.99 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8

interface loopback99
 ip address 99.99.99.99/32
 ip router ospf 65017 area 0.0.0.0
 ip pim sparse-mode

interface Ethernet2/1
 ip pim sparse-mode

interface Ethernet2/2
 ip pim sparse-mode
```

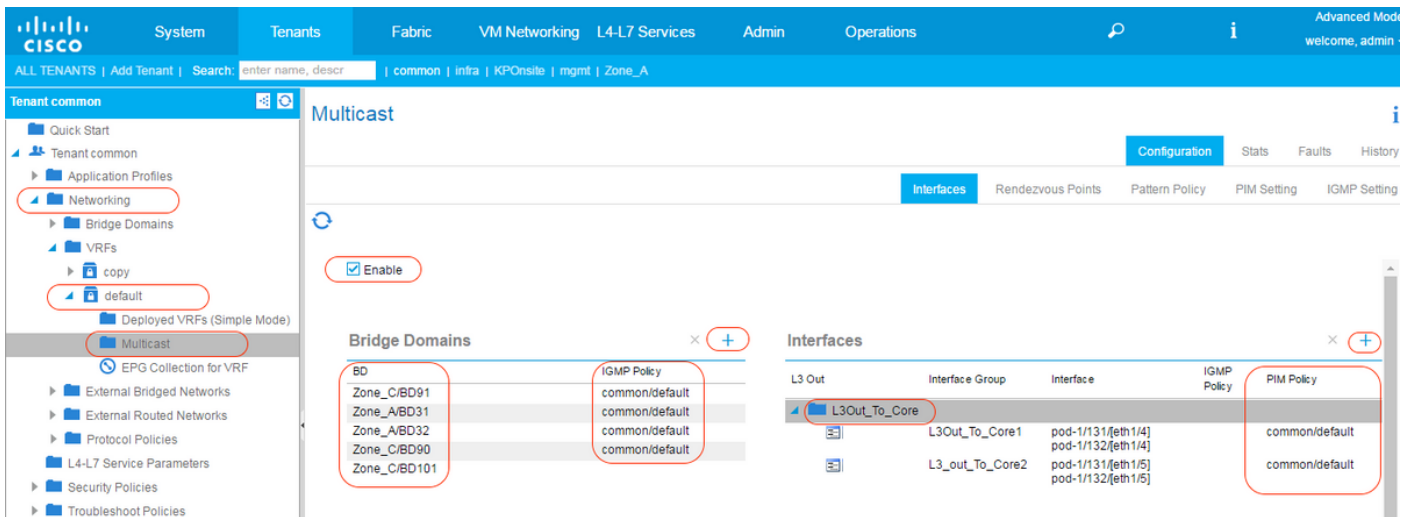
Step 1: Enable Multicast at the VRF. In Tenant space, navigate to Networking > VRFs > Multicast, on the working panel, click the button to enable Multicast.



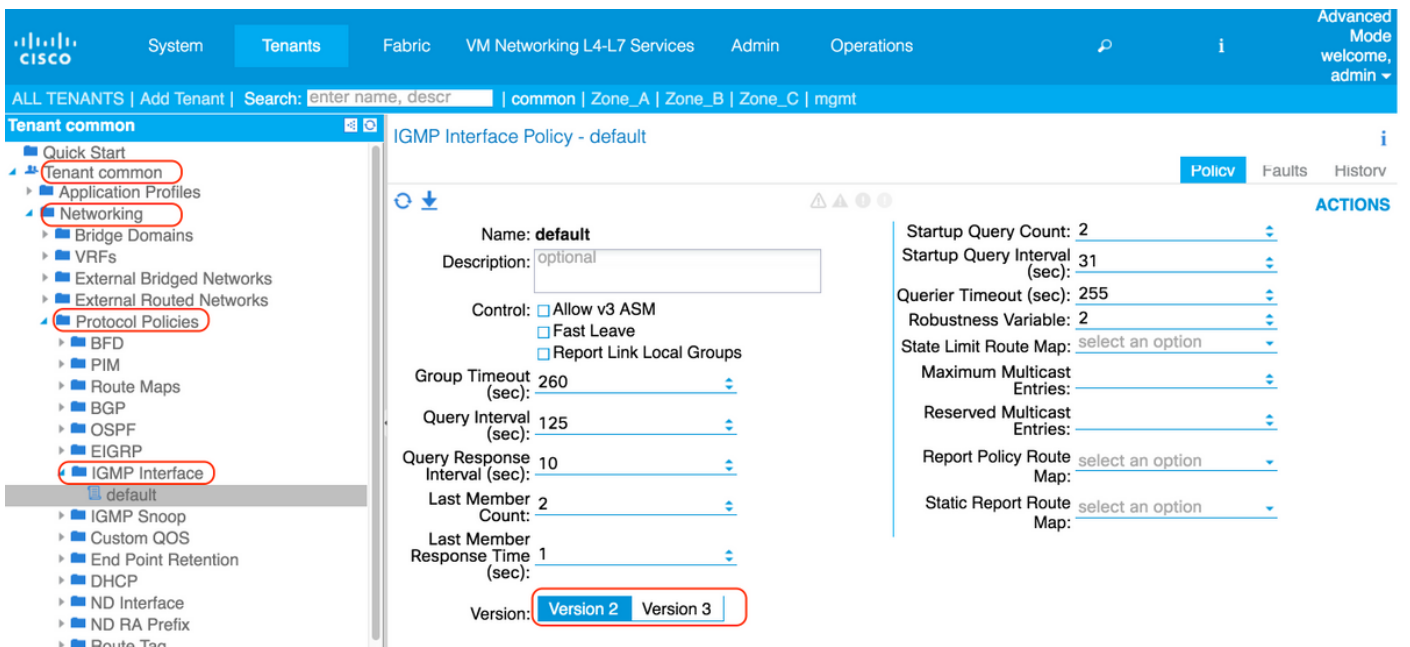
Step 2: Enable Multicast at BD and L3out levels, enable IGMP for receiver BDs. Navigate to Networking > VRFs > VRF name > Multicast, on the working panel, select Configuration > Interface tab, click "+" to add the Bridge Domains where multicast traffic is expected. Enable IMGP policy for multicast enabled BD.

Then click "+" to add L3out for this VRF. When multicast is enabled for an L3out it will enable PIM on all interfaces under the L3out, and all the border leaves for that L3out is enabled with multicast routing. Select the PIM policy for L3out interface groups.

Here assume the BDs and L3out are provisioned already.



Attaching IGMP policy to the BDs makes it an IGMP querier as well. IGMP policy is configured under Tenant > Networking > Protocol Policies > IGMP Interface. The Default IGMP policy has the following parameters where you can define query intervals. If no policy is specified, the interface will use the default policy.



PIM policy is also configured under Tenant > Networking > Protocol Policies > PIM.

The Default PIM policy has the following parameters where you can define hello intervals.

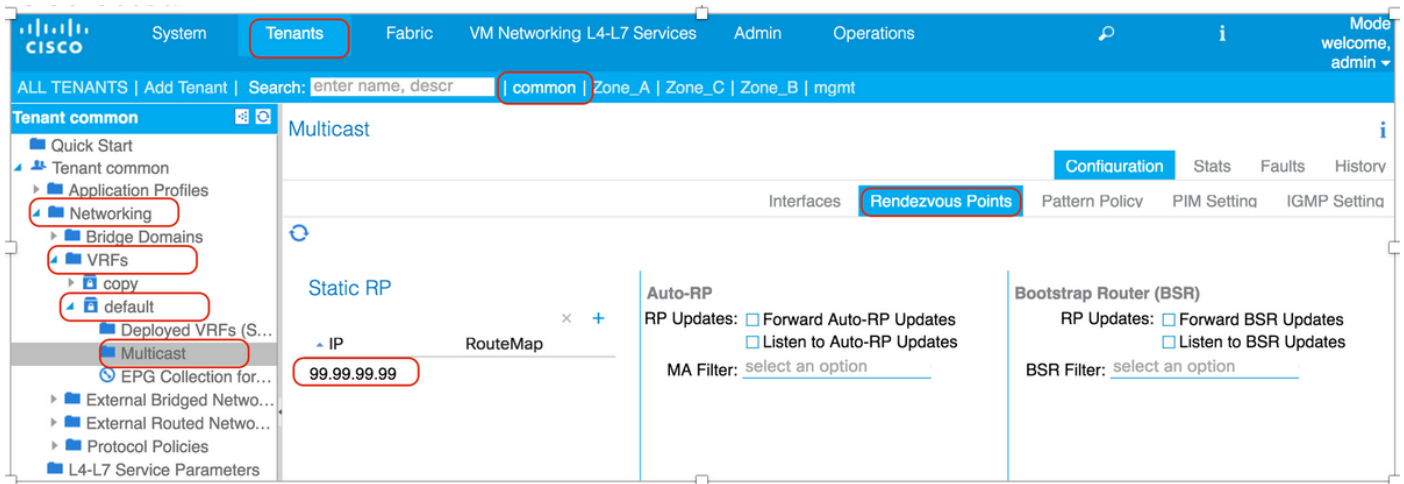
The screenshot shows the 'Edit Interface Policy' configuration page in the Cisco SD-WAN GUI. The left sidebar shows a tree view with 'Networking' and 'PIM Policies' expanded. The main panel shows the configuration for a policy named 'default'. The 'Auth Type' is set to 'MD5 HMAC authentication'. The 'Hello Interval (milliseconds)' is set to 30000. Other settings include 'Designated Router Delay (seconds): 3', 'Designated Router Priority: 1', and 'Join-Prune Interval Policy (seconds): 60'.

L3Outs on border leaf must be configured with loopback addresses enabled in the Node profile.

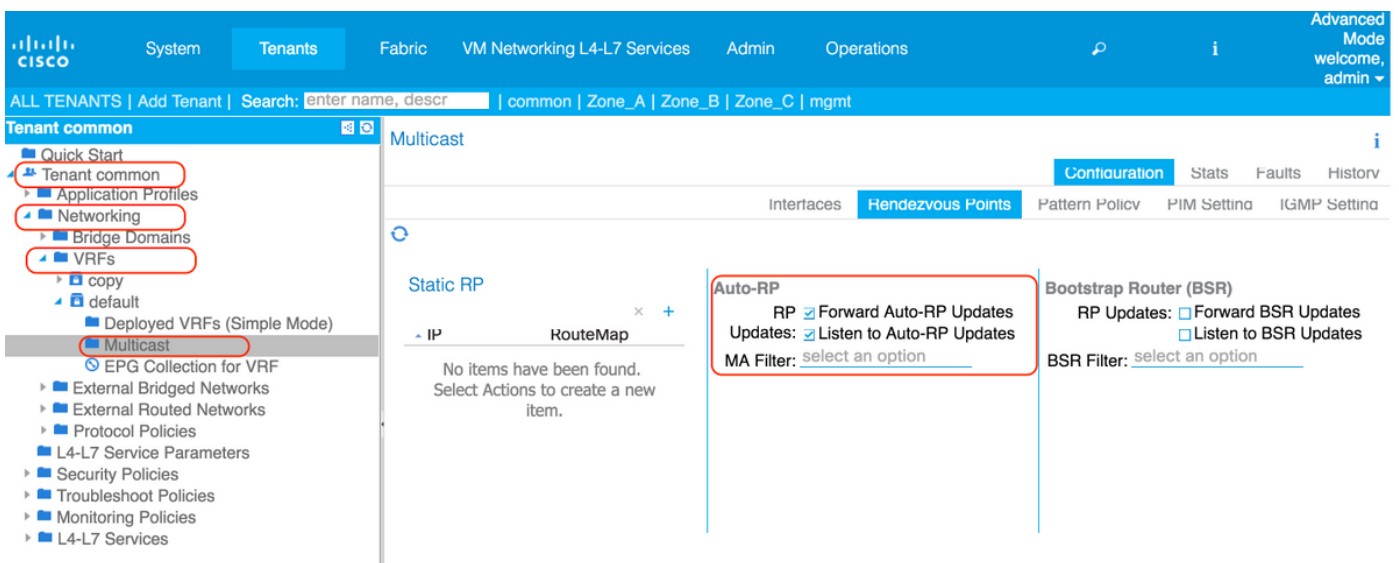
The screenshot shows the 'Logical Node Profile - L3Out_To_Core_NP' configuration page. The left sidebar shows a tree view with 'L3Out_To_Core' and 'Logical Node Profiles' expanded. The main panel shows the configuration for the profile. The 'Nodes' table lists two nodes with their Router IDs and Loopback Addresses.

Node ID	Router ID	Static Routes	Loopback Address
topology/pod-1/node-131	131.131.131.1		131.131.131.1
topology/pod-1/node-132	132.132.132.1		132.132.132.1

Step 3: Configure RP for PIM ASM. Navigate to Tenant > VRF > Multicast, in the working panel select Configuration > Rendezvous Points. In this example, Static RP is selected. Click "+" to add the RP.



For Auto-RP configuration, check the boxes for “Forward Auto-RP updates” and “Listen to Auto-RP Updates” on the “Rendezvous Points” page.



Outside the ACI fabric, the AUTO-RP configurations on NX-OS platforms remain the same.

!!! On RP candidate

```
ip pim send-rp-announce loopback99 group-list 224.0.0.0/4
ip pim send-rp-discovery loopback99 scope 32
```

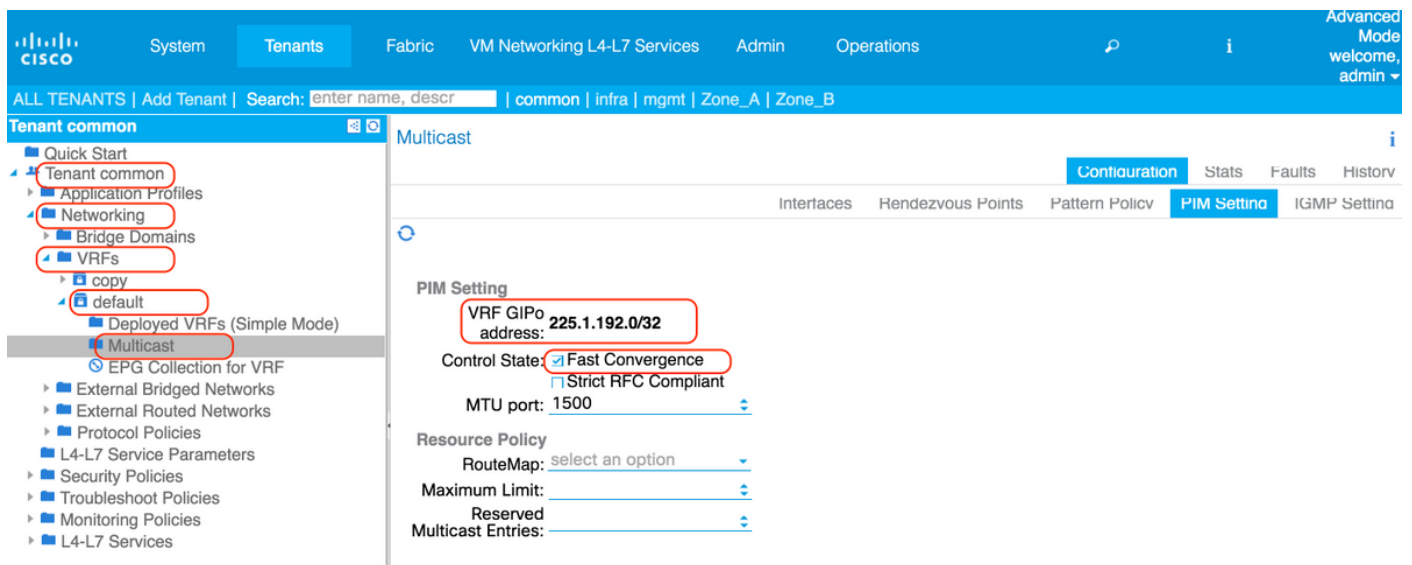
!!! On RP listeners:

```
ip pim auto-rp listen forward
```

Step 4: Configure necessary PIM settings. Navigate to Tenant > VRF -> Multicast, in the working panel, select Configuration > PIM settings, note the VRF GIPO address 225.1.192.0/32 which is assigned by APIC from the multicast group address pool. The VRF GIPO will be used as the outside group IP address for multicast traffic for BDs enabled with PIM.

When **Fast Convergence** mode is enabled (default is disable), all border leaves enabled with PIM will send joins towards the external network, but only one border leaf will forward traffic onto the fabric to prevent duplicates. The border leaf that forwards traffic for the group is the **designated forwarder** for the group. Enabling Fast Convergence helps reduce the packet drop duration for multicast flows with outside sources and inside receivers, when there is a stripe winner change due to the border leaf is down. There is no latency incurred by joining the PIM tree from the new stripe winner. This comes at the cost of the additional bandwidth usage on the non-stripe winners' external links since all of the border leaves pull traffic from the external source.

About the Stripe Winners - Currently ACI uses a BSR (Bootstrap Router) hash to calculate the BL stripe winner. The hash is computed using the S,G and Loopback IP of the Leaf. As of ACI 3.0(1) there is no way to influence stripe winner election for the user.

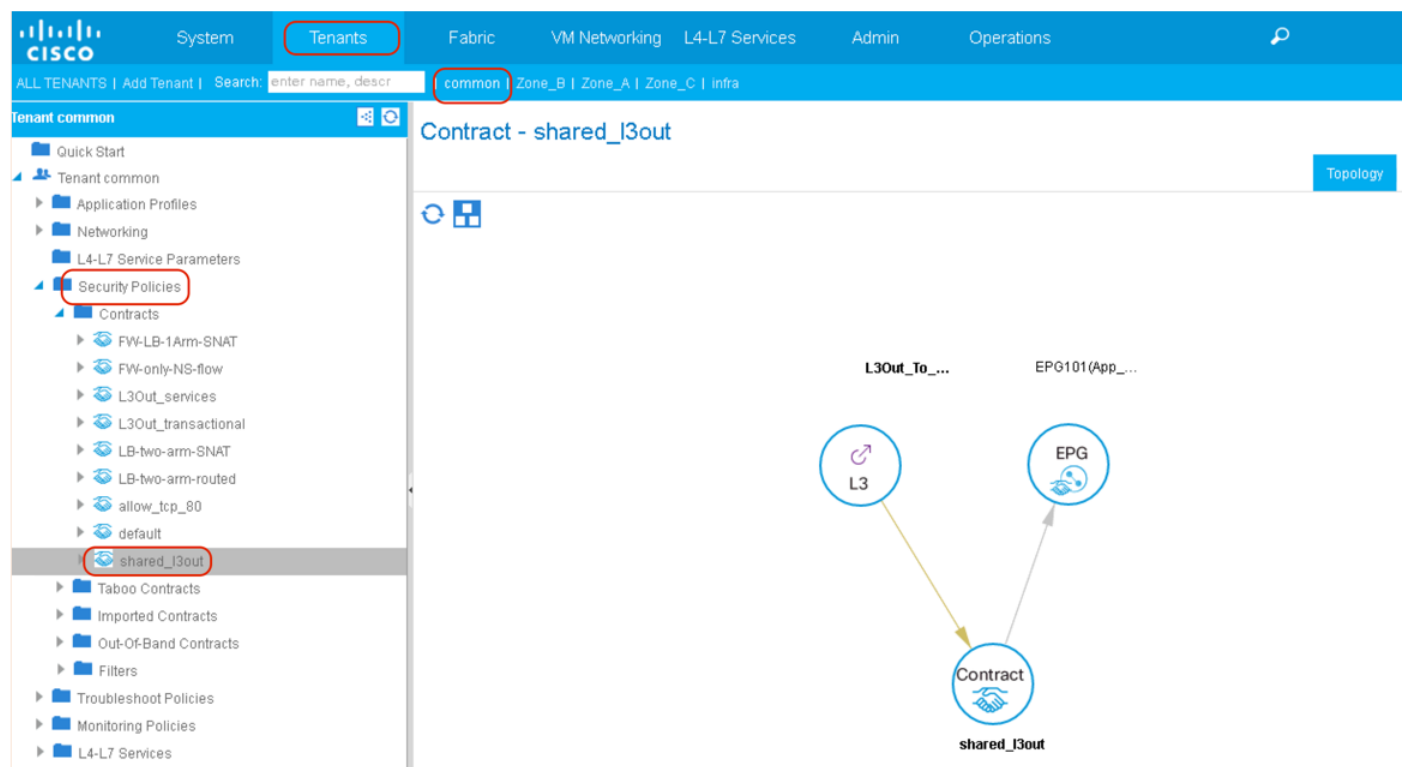


Step 5: Create needed contract to allow multicast traffic:

- Source and Receiver inside fabric (contract not needed)
- Receiver inside fabric, External source (contract not needed)
- Source inside fabric, External receiver (contract required)*

*The contract is not required if the BD is deployed on the border leaf

In our case, we have receivers outside the fabric, apply contract between L3out_to_Core and mcast source in EPG101.



Verify

PIM Verification

When the VRF is enabled for multicast routing, a fabric interface (tunnel) is created for multicast routing within the fabric. The PIM control plane packets are sent via the fabric interfaces within the fabric. The tunnel destination will be the VRF GiPo multicast address. On border leaf switches the tunnel source will be the loopback interface on the border leaf. On non-border leaf switches the tunnel source will be a loopback address (127.0.0.100).

Border leaves send PIM hellos on the fabric interface. L3Out interfaces run PIM in normal mode including sending and receiving hellos, electing DR, etc. Non-border leaves run in passive mode on the fabric interface; they listen for PIM hellos from the border leaves but do not send PIM hellos. Non-border leaves will not show up in output to "show ip pim neighbor".

```
!!!!!! Border Leaf Node bleaf1 !!!!!
```

```
bleaf1# show ip pim neighbor
```

```
PIM Neighbor information for Dom:common:default
```

Neighbor	Interface	Uptime	Expires	DRPriority
Bidir	BFDState			
132.132.132.1/32	tunnell16	06:20:40	00:01:21	1
no	n/a			
10.1.20.25/32	eth1/5	06:23:12	00:01:35	1
yes	n/a			
10.1.20.1/32	eth1/4	06:23:12	00:01:24	1
yes	n/a			

```
bleaf1# show interface tunnel 16
```

```
Tunnell16 is up
  MTU 9000 bytes, BW 0 Kbit
  Transport protocol is in VRF "common:default"
  Tunnel protocol/transport is ipvlan
  Tunnel source 131.131.131.1
  Tunnel destination 225.1.192.0/32
  Last clearing of "show interface" counters never
  Tx
  0 packets output, 1 minute output rate 0 packets/sec
  Rx
  0 packets input, 1 minute input rate 0 packets/sec
```

```
bleaf1#
```

```
!!!!!! Border Leaf Node bleaf2 !!!!!
```

```
bleaf2# show ip pim neighbor
```

```
PIM Neighbor information for Dom:common:default
```

Neighbor	Interface	Uptime	Expires	DRPriority
Bidir	BFDState			
131.131.131.1/32	tunnell16	06:23:26	00:01:30	1
no	n/a			
10.1.20.29/32	eth1/5	06:38:26	00:01:43	1
yes	n/a			
10.1.20.5/32	eth1/4	06:38:27	00:01:20	1
yes	n/a			

```
bleaf2# show interface tunnel 16
```

```
Tunnell16 is up
  MTU 9000 bytes, BW 0 Kbit
```



```
Transport protocol is in VRF "common:default"
Tunnel protocol/transport is ivxlan
Tunnel source 132.132.132.1
Tunnel destination 225.1.192.0/32
Last clearing of "show interface" counters never
Tx
0 packets output, 1 minute output rate 0 packets/sec
Rx
0 packets input, 1 minute input rate 0 packets/sec
```

bleaf2#

!!!! RP !!!!!

bleaf1# show ip pim rp vrf all

```
PIM RP Status Information for VRF:"common:default"
BSR: Not Operational
Auto-RP RPA: 192.168.1.2/32
RP: 99.99.99.99, uptime: 26d21h, expires: 00:02:38,
  priority: 0, RP-source: 192.168.1.2 (A), group-map: None, group ranges:
  224.0.0.0/4
```

bleaf1#

bleaf2# show ip pim rp vrf all

```
PIM RP Status Information for VRF:"common:default"
BSR: Not Operational
Auto-RP RPA: 192.168.1.2/32
RP: 99.99.99.99, uptime: 26d21h, expires: 00:02:38,
  priority: 0, RP-source: 192.168.1.2 (A), group-map: None, group ranges:
  224.0.0.0/4
```

bleaf2#

!!!! Non border leaf Node !!!!!

cleaf1# show ip pim neighbor

PIM Neighbor information for Dom:common:default		Interface	Uptime	Expires	DRPriority
Neighbor	BFDState				
132.132.132.1/32		tunnel16	06:32:43	00:01:37	1
no	n/a				
131.131.131.1/32		tunnel16	06:32:43	00:01:17	1
no	n/a				

cleaf1# show interface tunnel 16

```
Tunnel16 is up
  MTU 9000 bytes, BW 0 Kbit
  Transport protocol is in VRF "common:default"
  Tunnel protocol/transport is ivxlan
  Tunnel source 127.0.0.100/32
  Tunnel destination 225.1.192.0/32
Last clearing of "show interface" counters never
Tx
0 packets output, 1 minute output rate 0 packets/sec
Rx
0 packets input, 1 minute input rate 0 packets/sec
```

cleaf1#

cleaf2# show ip pim neighbor vrf all

PIM Neighbor information for Dom:common:default		Interface	Uptime	Expires	DRPriority
Neighbor					

```

Bidir      BFDState
132.132.132.1/32  tunnel16      06:33:17      00:01:33      1
no         n/a
131.131.131.1/32  tunnel16      06:33:17      00:01:41      1
no         n/a

```

```

cleaf2# show interface tunnel 16 Tunnel16 is up MTU 9000 bytes, BW 0 Kbit Transport protocol is
in VRF "common:default" Tunnel protocol/transport is ipvlan Tunnel source 127.0.0.100/32

```

```
Tunnel destination 225.1.192.0/32
```

```
Last clearing of "show interface" counters never
```

```
Tx
```

```
0 packets output, 1 minute output rate 0 packets/sec
```

```
Rx
```

```
0 packets input, 1 minute input rate 0 packets/sec
```

```
cleaf2#
```

```
!!!!!! Core Router !!!!!
```

```
N7K-core-1# show ip pim neighbor
```

```
PIM Neighbor Status for VRF "default"
```

Neighbor	Interface	Uptime	Expires	DR	Priority	Bidir- Capable	BFD State
10.1.20.2	Ethernet1/1	3d22h	00:01:43	1	1	no	n/a
10.1.20.6	Ethernet1/2	3d22h	00:01:36	1	1	no	n/a
10.1.20.10	Ethernet1/3	2w6d	00:01:30	1	1	yes	n/a
10.1.20.14	Ethernet1/4	2w6d	00:01:18	1	1	yes	n/a
10.1.20.42	Ethernet1/5	2w6d	00:01:28	1	1	yes	n/a

```
N7K-core-1#
```

```
N7K-core-2# sh ip pim neighbor
```

```
PIM Neighbor Status for VRF "default"
```

Neighbor	Interface	Uptime	Expires	DR	Priority	Bidir- Capable	BFD State
10.1.20.26	Ethernet1/1	3d22h	00:01:23	1	1	no	n/a
10.1.20.30	Ethernet1/2	3d22h	00:01:17	1	1	no	n/a
10.1.20.18	Ethernet1/3	2w6d	00:01:38	1	1	yes	n/a
10.1.20.22	Ethernet1/4	2w6d	00:01:41	1	1	yes	n/a
10.1.20.46	Ethernet1/5	2w6d	00:01:17	1	1	yes	n/a

```
N7K-core-2#
```

Active Border Leaf Verification

If there are more than one border leaf is enabled with multicast routing, APIC chooses one stripe winner for each group address across all active border leaves. The border leaf that is the stripe winner for a group is responsible for sending PIM joins on behalf of the fabric and forwarding multicast traffic into the fabric.

The stripe winner for the group decides on the Designated Forwarder. If the stripe winner has reachability to the root, then the stripe winner is the DF also. If the stripe winner does not have external connectivity to the root, then that BL chooses a DF by sending a PIM join over the fabric interface.

```
!!!!!! Enter into vsh mode to execute the command !!!!!
```

```
bleaf2# vsh
```

```
Cisco iNX-OS Debug Shell
```

```
This shell should only be used for internal commands and exists
for legacy reasons. User should use ibash infrastructure as this
will be deprecated.
```

```
bleaf2# show ip pim internal stripe-winner 228.0.0.1 vrf common:default
```

```
PIM Stripe Winner info for VRF "common:default" (BL count: 2)
```

```
(*, 228.0.0.1)
```

```
BLs: 132.132.132.1 hash: 2081913316 (local)
```

```

131.131.131.1 hash: 1024236260
Winner: 132.132.132.1 best_hash: 2081913316
bleaf2#
bleaf2#
bleaf2# show ip pim internal stripe-winner 229.0.0.1 vrf common:default
PIM Stripe Winner info for VRF "common:default" (BL count: 2)
(*, 229.0.0.1)
BLs: 132.132.132.1 hash: 1595374052 (local)
131.131.131.1 hash: 2047646436
Winner: 131.131.131.1 best_hash: 2047646436
bleaf2#

```

Fast Convergence Verification

```

!!! Verify if fast convergence is enabled
bleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name          VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
                  ID       If          ID          Role     Mode     IP
common:default    4        Tunnel16   2162688    BL       Fast     131.131.131.1
bleaf1#

```

!!! None-border leaf

```

cleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name          VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
                  ID       If          ID          Role     Mode     IP
common:default    4        Tunnel16   2162688    Leaf    Fast     127.0.0.100
cleaf1#

```

IGMP Verification

```

!!! Verify if fast convergence is enabled
bleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name          VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
                  ID       If          ID          Role     Mode     IP
common:default    4        Tunnel16   2162688    BL       Fast     131.131.131.1
bleaf1#

```

!!! None-border leaf

```

cleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name          VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
                  ID       If          ID          Role     Mode     IP
common:default    4        Tunnel16   2162688    Leaf    Fast     127.0.0.100
cleaf1#

```

MRIB Verification

The leaf node sleaf2 which is the FHR has directly attached multicast sources. Its RPF neighbor is 10.0.176.64 on spine1. The incoming interface is the fabric interface (tunnel16) which is peering with the border leaf via PIM.

For simplicity, the output shown is for one multicast IP address from each group range: 228.0.0.1 for inside sources, and 229.0.0.1 for the outside sources.

```

!!! Verify if fast convergence is enabled
bleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name          VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
                  ID       If          ID          Role     Mode     IP

```

```
common:default      4      Tunnel16      2162688      BL      Fast      131.131.131.1
bleaf1#
```

!!! None-border leaf

```
cleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name      VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
              ID      If          ID          Role     Mode     IP
common:default  4      Tunnel16    2162688    Leaf     Fast     127.0.0.100
cleaf1#
```

Receivers for 228.0.0.1 are attached to bleaf2 (node 132), cleaf1 (Node 101) and cleaf2 (Node 102). Bleaf2 forwards the mcast to group 228.0.0.1 to internal receivers via tunnel 16, and external receivers via L3out to the core devices.

```
!!! Verify if fast convergence is enabled
bleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name      VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
              ID      If          ID          Role     Mode     IP
common:default  4      Tunnel16    2162688    BL      Fast     131.131.131.1
bleaf1#
```

!!! None-border leaf

```
cleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name      VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
              ID      If          ID          Role     Mode     IP
common:default  4      Tunnel16    2162688    Leaf     Fast     127.0.0.100
cleaf1#
```

Bleaf1 forwards group of 228.0.0.1 to outside via L3out interface, but it does not forward onto the fabric via the fabric interfaces since it is not the stripe winner for 228.0.0.1

```
!!! Verify if fast convergence is enabled
bleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name      VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
              ID      If          ID          Role     Mode     IP
common:default  4      Tunnel16    2162688    BL      Fast     131.131.131.1
bleaf1#
```

!!! None-border leaf

```
cleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name      VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
              ID      If          ID          Role     Mode     IP
common:default  4      Tunnel16    2162688    Leaf     Fast     127.0.0.100
cleaf1#
```

Bleaf1 is the active border leaf/stripe winner for group 229.0.0.1. bleaf1 receives the multicast to group 229.0.0.1 via the external core device, and then forwards to internal receivers in BD90, BD91, BD31, BD32. (note the vlan ID is only internal to leaf node as the pervasive GW).

```
!!! Verify if fast convergence is enabled
bleaf1# show fabric multicast vrf common:default
Fabric Multicast Enabled VRFs
VRF Name      VRF      Vprime      VN-Seg      VRF      Conv      Tunnel
              ID      If          ID          Role     Mode     IP
common:default  4      Tunnel16    2162688    BL      Fast     131.131.131.1
bleaf1#
```

!!! None-border leaf

cleaf1# show fabric multicast vrf common:default

Fabric Multicast Enabled VRFs

VRF Name	VRF ID	Vprime If	VN-Seg ID	VRF Role	Conv Mode	Tunnel IP
common:default	4	Tunnell16	2162688	Leaf	Fast	127.0.0.100

cleaf1#

Non-border leaf Cleaf1 and Cleaf2 has attached receivers in BD31, BD32, BD91. Non-border leaf node install (*, G) only, (S,G) is not supported.

cleaf1# show ip mroute 228.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"

(*, 228.0.0.1/32), uptime: 3w5d, igmp ip pim
Incoming interface: Tunnell14, RPF nbr: 10.0.80.91
Outgoing interface list: (count: 2)
Vlan4, uptime: 1w5d, igmp
Vlan7, uptime: 3w5d, igmp

cleaf1# show ip mroute 229.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"

(*, 229.0.0.1/32), uptime: 3w5d, igmp ip pim
Incoming interface: Tunnell14, RPF nbr: 10.0.80.91
Outgoing interface list: (count: 2)
Vlan4, uptime: 1w5d, igmp
Vlan7, uptime: 3w5d, igmp

cleaf1#

cleaf1# show interface vlan 4

Vlan4 is up, line protocol is up
Hardware EtherSVI, address is 0000.0c07.ac1f
Internet Address is 10.31.31.1/24
MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
reliability 255/255, txload 1/255, rxload 1/255
Carrier delay is 10 sec
Encapsulation ARPA, loopback not set
Keepalive not supported
ARP type: ARPA
Last clearing of "show interface" counters never
30 seconds input rate 0 bits/sec, 0 packets/sec
30 seconds output rate 0 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
L3 Switched:
input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
L3 in Switched:
ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
L3 out Switched:
ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes

cleaf1# show interface vlan 7

Vlan7 is up, line protocol is up
Hardware EtherSVI, address is 0000.0c07.ac20
Internet Address is 10.32.32.1/24
MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
reliability 255/255, txload 1/255, rxload 1/255
Carrier delay is 10 sec

```
Encapsulation ARPA, loopback not set
Keepalive not supported
ARP type: ARPA
Last clearing of "show interface" counters never
30 seconds input rate 0 bits/sec, 0 packets/sec
30 seconds output rate 0 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
L3 Switched:
  input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
L3 in Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
L3 out Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

cleaf1#

!!!! Non-border leaf node has (*, G) only, (S,G) is not supported.

```
cleaf2# show ip mroute 228.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(*, 228.0.0.1/32), uptime: 3w5d, igmp ip pim
  Incoming interface: Tunnell16, RPF nbr: 10.0.80.91
  Outgoing interface list: (count: 3)
    Vlan3, uptime: 1w5d, igmp
    Vlan30, uptime: 3w5d, igmp
    Vlan9, uptime: 3w5d, igmp
```

```
cleaf2# show ip mroute 229.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(*, 229.0.0.1/32), uptime: 3w5d, igmp ip pim
  Incoming interface: Tunnell16, RPF nbr: 10.0.80.91
  Outgoing interface list: (count: 3)
    Vlan3, uptime: 1w5d, igmp
    Vlan30, uptime: 3w5d, igmp
    Vlan9, uptime: 3w5d, igmp
```

cleaf2#

```
cleaf2# show interface vlan 3
```

```
Vlan3 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.ac1f
  Internet Address is 10.31.31.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
  reliability 255/255, txload 1/255, rxload 1/255
  Carrier delay is 10 sec
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA
  Last clearing of "show interface" counters never
  30 seconds input rate 0 bits/sec, 0 packets/sec
  30 seconds output rate 0 bits/sec, 0 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  L3 Switched:
    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
  L3 in Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
  L3 out Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2# show interface vlan 30
Vlan30 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.ac5b
  Internet Address is 10.91.91.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
    reliability 255/255, txload 1/255, rxload 1/255
  Carrier delay is 10 sec
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA
  Last clearing of "show interface" counters never
  30 seconds input rate 0 bits/sec, 0 packets/sec
  30 seconds output rate 0 bits/sec, 0 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  L3 Switched:
    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
  L3 in Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
  L3 out Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2# show interface vlan 9
Vlan9 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.ac20
  Internet Address is 10.32.32.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
    reliability 255/255, txload 1/255, rxload 1/255
  Carrier delay is 10 sec
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA
  Last clearing of "show interface" counters never
  30 seconds input rate 0 bits/sec, 0 packets/sec
  30 seconds output rate 0 bits/sec, 0 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  L3 Switched:
    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
  L3 in Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
  L3 out Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2#
```

On the core routers, N7K-core-1 and N7K-core-2 are load sharing for multicast flows sourced in LAN network, if fast convergence is not enabled, only one border leaf (bleaf1) is sending join towards the source.

```
cleaf1# show ip mroute 228.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(* , 228.0.0.1/32), uptime: 3w5d, igmp ip pim
  Incoming interface: Tunnell4, RPF nbr: 10.0.80.91
  Outgoing interface list: (count: 2)
    Vlan4, uptime: 1w5d, igmp
    Vlan7, uptime: 3w5d, igmp
```

```
cleaf1# show ip mroute 229.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(* , 229.0.0.1/32), uptime: 3w5d, igmp ip pim
  Incoming interface: Tunnell4, RPF nbr: 10.0.80.91
```



```
Outgoing interface list: (count: 2)
  Vlan4, uptime: 1w5d, igmp
  Vlan7, uptime: 3w5d, igmp
```

```
cleaf1#
```

```
cleaf1# show interface vlan 4
Vlan4 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.ac1f
  Internet Address is 10.31.31.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
    reliability 255/255, txload 1/255, rxload 1/255
  Carrier delay is 10 sec
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA
  Last clearing of "show interface" counters never
  30 seconds input rate 0 bits/sec, 0 packets/sec
  30 seconds output rate 0 bits/sec, 0 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  L3 Switched:
    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
  L3 in Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
  L3 out Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf1# show interface vlan 7
Vlan7 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.ac20
  Internet Address is 10.32.32.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
    reliability 255/255, txload 1/255, rxload 1/255
  Carrier delay is 10 sec
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA
  Last clearing of "show interface" counters never
  30 seconds input rate 0 bits/sec, 0 packets/sec
  30 seconds output rate 0 bits/sec, 0 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  L3 Switched:
    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
  L3 in Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
  L3 out Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf1#
```

```
!!!! Non-border leaf node has (*, G) only, (S,G) is not supported.
cleaf2# show ip mroute 228.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(*, 228.0.0.1/32), uptime: 3w5d, igmp ip pim
  Incoming interface: Tunnell6, RPF nbr: 10.0.80.91
  Outgoing interface list: (count: 3)
    Vlan3, uptime: 1w5d, igmp
    Vlan30, uptime: 3w5d, igmp
```

Vlan9, uptime: 3w5d, igmp

```
cleaf2# show ip mroute 229.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(*, 229.0.0.1/32), uptime: 3w5d, igmp ip pim
  Incoming interface: Tunnell16, RPF nbr: 10.0.80.91
  Outgoing interface list: (count: 3)
    Vlan3, uptime: 1w5d, igmp
    Vlan30, uptime: 3w5d, igmp
    Vlan9, uptime: 3w5d, igmp
```

```
cleaf2#
```

```
cleaf2# show interface vlan 3
```

```
Vlan3 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.ac1f
  Internet Address is 10.31.31.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
    reliability 255/255, txload 1/255, rxload 1/255
  Carrier delay is 10 sec
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA
  Last clearing of "show interface" counters never
  30 seconds input rate 0 bits/sec, 0 packets/sec
  30 seconds output rate 0 bits/sec, 0 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  L3 Switched:
    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
  L3 in Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
  L3 out Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2# show interface vlan 30
```

```
Vlan30 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.ac5b
  Internet Address is 10.91.91.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
    reliability 255/255, txload 1/255, rxload 1/255
  Carrier delay is 10 sec
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA
  Last clearing of "show interface" counters never
  30 seconds input rate 0 bits/sec, 0 packets/sec
  30 seconds output rate 0 bits/sec, 0 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  L3 Switched:
    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
  L3 in Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
  L3 out Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2# show interface vlan 9
```

```
Vlan9 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.ac20
  Internet Address is 10.32.32.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
```

```
reliability 255/255, txload 1/255, rxload 1/255
Carrier delay is 10 sec
Encapsulation ARPA, loopback not set
Keepalive not supported
ARP type: ARPA
Last clearing of "show interface" counters never
30 seconds input rate 0 bits/sec, 0 packets/sec
30 seconds output rate 0 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
L3 Switched:
  input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
L3 in Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
L3 out Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2#
```

```
cleaf1# show ip mroute 228.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(* , 228.0.0.1/32), uptime: 3w5d, igmp ip pim
  Incoming interface: Tunnell14, RPF nbr: 10.0.80.91
  Outgoing interface list: (count: 2)
    Vlan4, uptime: 1w5d, igmp
    Vlan7, uptime: 3w5d, igmp
```

```
cleaf1# show ip mroute 229.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(* , 229.0.0.1/32), uptime: 3w5d, igmp ip pim
  Incoming interface: Tunnell14, RPF nbr: 10.0.80.91
  Outgoing interface list: (count: 2)
    Vlan4, uptime: 1w5d, igmp
    Vlan7, uptime: 3w5d, igmp
```

```
cleaf1#
```

```
cleaf1# show interface vlan 4
```

```
Vlan4 is up, line protocol is up
  Hardware EtherSVI, address is 0000.0c07.aclf
  Internet Address is 10.31.31.1/24
  MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
  reliability 255/255, txload 1/255, rxload 1/255
  Carrier delay is 10 sec
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA
  Last clearing of "show interface" counters never
  30 seconds input rate 0 bits/sec, 0 packets/sec
  30 seconds output rate 0 bits/sec, 0 packets/sec
  Load-Interval #2: 5 minute (300 seconds)
    input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
  L3 Switched:
    input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
  L3 in Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
  L3 out Switched:
    ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf1# show interface vlan 7
```

```
Vlan7 is up, line protocol is up
```

```
Hardware EtherSVI, address is 0000.0c07.ac20
Internet Address is 10.32.32.1/24
MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
  reliability 255/255, txload 1/255, rxload 1/255
Carrier delay is 10 sec
Encapsulation ARPA, loopback not set
Keepalive not supported
ARP type: ARPA
Last clearing of "show interface" counters never
30 seconds input rate 0 bits/sec, 0 packets/sec
30 seconds output rate 0 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
L3 Switched:
  input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
L3 in Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
L3 out Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

cleaf1#

!!!! Non-border leaf node has (*, G) only, (S,G) is not supported.

```
cleaf2# show ip mroute 228.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(*, 228.0.0.1/32), uptime: 3w5d, igmp ip pim
Incoming interface: Tunnell6, RPF nbr: 10.0.80.91
Outgoing interface list: (count: 3)
  Vlan3, uptime: 1w5d, igmp
  Vlan30, uptime: 3w5d, igmp
  Vlan9, uptime: 3w5d, igmp
```

```
cleaf2# show ip mroute 229.0.0.1 vrf common:default
IP Multicast Routing Table for VRF "common:default"
```

```
(*, 229.0.0.1/32), uptime: 3w5d, igmp ip pim
Incoming interface: Tunnell6, RPF nbr: 10.0.80.91
Outgoing interface list: (count: 3)
  Vlan3, uptime: 1w5d, igmp
  Vlan30, uptime: 3w5d, igmp
  Vlan9, uptime: 3w5d, igmp
```

cleaf2#

```
cleaf2# show interface vlan 3
```

```
Vlan3 is up, line protocol is up
Hardware EtherSVI, address is 0000.0c07.ac1f
Internet Address is 10.31.31.1/24
MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
  reliability 255/255, txload 1/255, rxload 1/255
Carrier delay is 10 sec
Encapsulation ARPA, loopback not set
Keepalive not supported
ARP type: ARPA
Last clearing of "show interface" counters never
30 seconds input rate 0 bits/sec, 0 packets/sec
30 seconds output rate 0 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
L3 Switched:
  input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
```

```
L3 in Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
L3 out Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2# show interface vlan 30
Vlan30 is up, line protocol is up
Hardware EtherSVI, address is 0000.0c07.ac5b
Internet Address is 10.91.91.1/24
MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
  reliability 255/255, txload 1/255, rxload 1/255
Carrier delay is 10 sec
Encapsulation ARPA, loopback not set
Keepalive not supported
ARP type: ARPA
Last clearing of "show interface" counters never
30 seconds input rate 0 bits/sec, 0 packets/sec
30 seconds output rate 0 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
L3 Switched:
  input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
L3 in Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
L3 out Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2# show interface vlan 9
Vlan9 is up, line protocol is up
Hardware EtherSVI, address is 0000.0c07.ac20
Internet Address is 10.32.32.1/24
MTU 9000 bytes, BW 10000000 Kbit, DLY 1 usec
  reliability 255/255, txload 1/255, rxload 1/255
Carrier delay is 10 sec
Encapsulation ARPA, loopback not set
Keepalive not supported
ARP type: ARPA
Last clearing of "show interface" counters never
30 seconds input rate 0 bits/sec, 0 packets/sec
30 seconds output rate 0 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
L3 Switched:
  input: 0 pkts, 0 bytes - output: 0 pkts, 0 bytes
L3 in Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
L3 out Switched:
  ucast: 0 pkts, 0 bytes - mcast: 0 pkts, 0 bytes
```

```
cleaf2#
```

References

[ACI 2.0 Multicast Routing](#)