

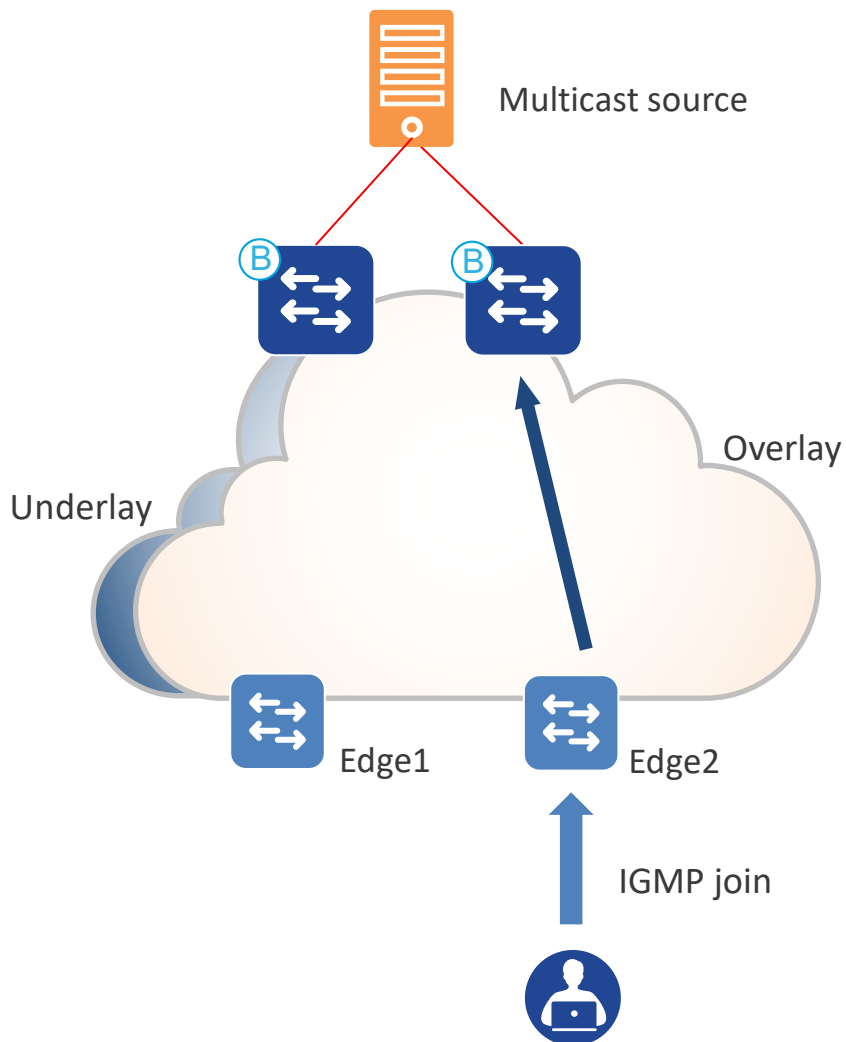
Multicast in SDA

The document describes how to configure and validate multicast configuration for SDA. A number of scenarios have been tested:

1. PIM ASM with multicast source outside the fabric and receivers within the fabric
2. PIM ASM with multicast source and receivers within the fabric
3. PIM ASM with dual RP configuration (MSDP)
4. PIM SSM with multicast source outside the fabric and receivers within the fabric

Please note that currently there is a support for multicast implementation over Locator/ID Separation Protocol (LISP) using ingress replication. *It means that unicast Routing Locator (RLOC) core is used to transport identity information (EID) multicast.*

The topology being used:



1. PIM ASM with multicast source outside the fabric and receivers within the fabric

In order to configure multicast in SDA, it is necessary to designate a separate network pool that will be used for multicast overlay control plane connectivity (**Design -> Network Settings -> IP Address Pools -> Add IP Pool**).

Edit IP Pool ✕

IP Pool Name *
Mcast_pool

IP Subnet *
100.100.100.0

CIDR Prefix
/24 (255.255.255.0)

Gateway IP Address *
100.100.100.1

DHCP Server(s) ▼

DNS Server(s) ▼

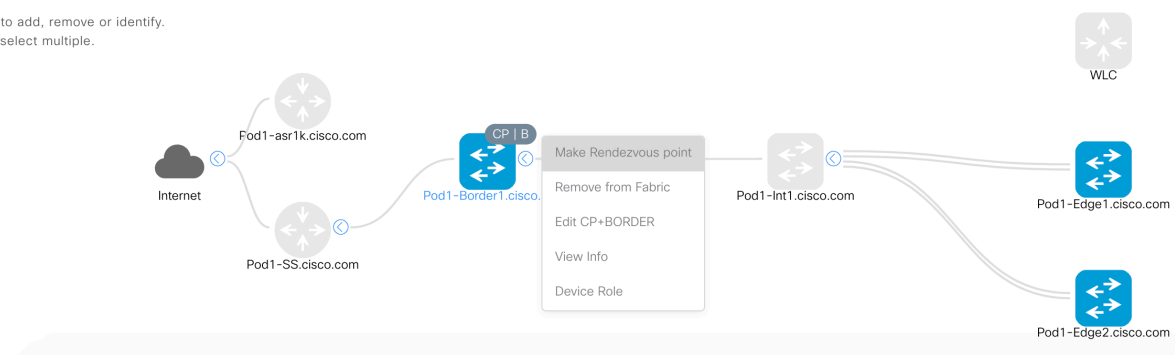
Overlapping

Cancel
Save

Once configured, the next step would be to enable Rendezvous point in the Fabric (**Provision -> Fabric -> Select Devices -> Make Rendezvous point**). The optimal choice from the traffic flow perspective is to designate Border router for this function. (this function can be also handled by the control plane node which will be demonstrated in the subsequent section).

Search Topology

Select Devices to add, remove or identify.
Shift + Click to select multiple.





Once selected, DNAC will prompt the list of Virtual Networks to enable multicast.

Rendezvous Point: Pod1 - ×
Border1.cisco.com

Select Virtual Network

Campus Guest

DEFAULT_VN

Cancel Enable

If all the steps has been followed, the “M” character should appear for all corresponding Virtual Networks. The last step in the configuration is to enable the newly created network pool in the Host Onboarding section.

University -->

Select Devices Host Onboarding

Select Authentication template

Validation Type ▾

Save

Closed Authentication Easy Connect No Authentication Open Authentication

Virtual Networks

Campus M Guest DEFAULT_VN INFRA_VN

Choose Pool Type ×

Host Pools Multicast Pools

Cancel

Next



Edit Virtual Network: Campus

<input type="checkbox"/>	IP Pool Name	Address Pool
<input type="checkbox"/>	BorderWAN	192.168.19.0/24
<input type="checkbox"/>	BorderWan2	192.168.20.0/24
<input type="checkbox"/>	GuestNet	172.16.19.0/24
<input type="checkbox"/>	Infrastructure	172.16.12.0/24
<input checked="" type="checkbox"/>	Mcast_pool	100.100.100.0/24
<input type="checkbox"/>	Production	172.16.11.0/24
<input type="checkbox"/>	Underlay_Global	192.168.80.0/24
<input type="checkbox"/>	WirelessAP	192.168.18.0/24
<input type="checkbox"/>	WirelessUser	172.16.18.0/24

Show entries Showing 1 - 9 of 9 Previous 1 Next

Cancel Update

Once saved, DNAC pushes all multicast relevant configuration to the Edge and Border/RP devices. As of now, the SDA fabric is multicast enabled and allows to forward multicast streams southbound (the border downstream) and east-west (between the hosts connected to the Fabric edge switches).

What happens behind the scenes

Configuration pushed from DNAC

When DNAC provisions multicast in the Fabric it performs several configuration changes on the Fabric Border and Fabric Edge switches.

On Fabric Edge switches:

1. Enables Multicast Routing across all the switches in the Fabric

```
ip multicast-routing
ip multicast-routing vrf Campus
```

2. Enables IGMP and PIM on the SVI overlay interfaces on Fabric Edge switches

```
interface Vlan1022
description Configured from apic-em
mac-address 0000.0c9f.f45d
vrf forwarding Campus
ip address 172.16.12.1 255.255.255.0
ip helper-address 192.168.11.201
```



```
no ip redirects
ip local-proxy-arp
ip pim sparse-mode
ip route-cache same-interface
ip igmp version 3
no lisp mobility liveness test
lisp mobility 172_16_12_0-Campus
end
```

3. Registers Edge address with LISP for overlay connectivity

```
router lisp
<snip>
instance-id 4099
  service ipv4
    eid-table vrf Campus
    database-mapping 100.100.100.3/32 locator-set rloc_3374259c-bfb8-481f-9a60-6d2f31bd286f
<snip>
```

4. Configures static RP across all the switches belonging to the fabric (Fabric Edge and Fabric Border). The RP ip address is chosen from the multicast designated pool (Mcast_pool).

```
ip pim ssm default
ip pim vrf Campus rp-address 100.100.100.1
ip pim vrf Campus register-source Loopback4099
ip pim vrf Campus ssm default
```

5. Creates a dedicated Loopback interface for overlay multicast connectivity

```
interface Loopback4099
vrf forwarding Campus
ip address 100.100.100.3 255.255.255.255
```

6. Adds PIM under LISP interface

```
interface LISP0.4099
ip pim sparse-mode
```

On Border Edge/RP switches:

1. Enables Multicast Routing across all the switches in the Fabric

```
ip multicast-routing
ip multicast-routing vrf Campus
```

2. Registers RP address with LISP for overlay connectivity

```
router lisp
<snip>
instance-id 4099
  service ipv4
    eid-table vrf Campus
    database-mapping 100.100.100.1/32 locator-set rloc_18a7f2ea-6638-4d06-8d70-0f5adcc8e458
<snip>
```

3. Configures static RP across all the switches belonging to the fabric (Fabric Edge and Fabric Border). The RP ip address is chosen from the multicast designated pool (Mcast_pool).

```
ip pim ssm default
ip pim vrf Campus rp-address 100.100.100.1
ip pim vrf Campus register-source Loopback4099
ip pim vrf Campus ssm default
```



4. Creates a dedicated Loopback interface for overlay multicast connectivity

```
interface Loopback4099
vrf forwarding Campus
ip address 100.100.100.1 255.255.255.255
ip pim sparse-mode
```

5. Adds PIM under LISP interface

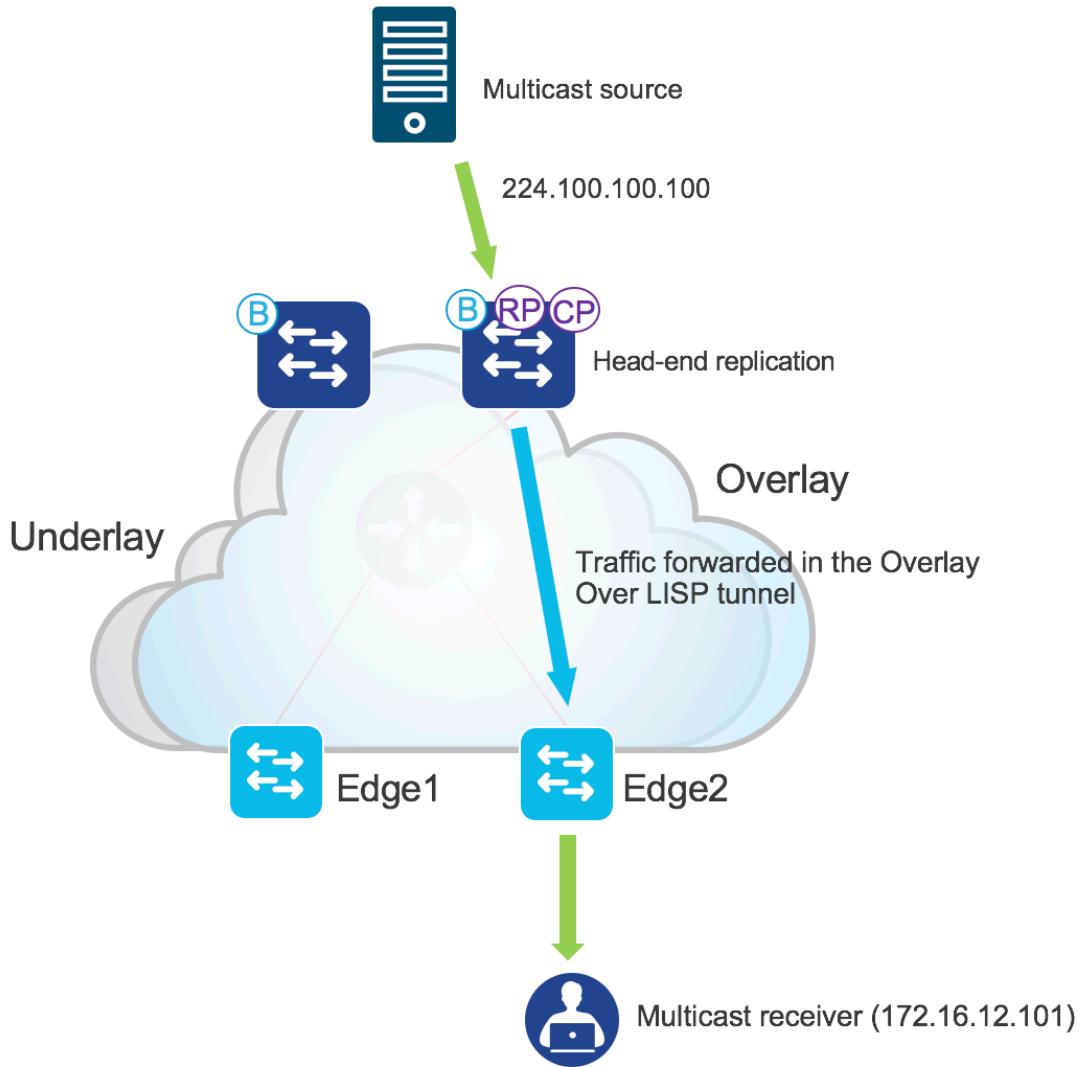
```
interface LISP0.4099
ip pim sparse-mode
```

6. Adds PIM on the L3 overlay hand-off interface towards the Fusion router

```
interface Vlan3003
description vrf interface to External router
vrf forwarding Campus
ip address 192.168.19.9 255.255.255.252
no ip redirects
ip pim sparse-mode
ip route-cache same-interface
end
```

Please note that there is no Multicast underlay configuration being pushed from DNAC. Currently there is a support for Head-end replication only – meaning the ITR in the fabric is responsible for multicast replication according to the **Location-set**

PIM ASM with multicast source outside the fabric and receivers within the fabric (single receiver)





Multicast traffic flow verification

Border/RP:

```
Pod1-Border1#sh ip pim vrf Campus interface
```

Address	Interface	Ver/ Mode	Nbr Count	Query Intvl	DR Prior	DR
172.16.12.1	LISP0.4099	v2/S	0	30	1	172.16.12.1
192.168.19.9	Vlan3003	v2/S	0	30	1	192.168.19.9
100.100.100.1	Loopback4099	v2/S	0	30	1	100.100.100.1

```
Pod1-Border1#sh ip pim vrf Campus rp mapping
```

```
PIM Group-to-RP Mappings
```

```
Group(s): 224.0.0.0/4, Static  
RP: 100.100.100.1 (?)
```

PIM Hello messages are not exchanged between the routers within the fabric hence we do not see the PIM neighborship in the output of the following command. PIM join/prune are the only messages exchanged.

```
Pod1-Border1#sh ip pim vrf Campus neighbor
```

```
PIM Neighbor Table
```

```
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,  
P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,  
L - DR Load-balancing Capable
```

Neighbor Address	Interface	Uptime/Expires	Ver	DR Prio/Mode
---------------------	-----------	----------------	-----	-----------------

The traffic not yet initiated from the source:

```
Pod1-Border1#sh ip mroute vrf Campus
```

```
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,  
L - Local, P - Pruned, R - RP-bit set, F - Register flag,  
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,  
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,  
U - URD, I - Received Source Specific Host Report,  
Z - Multicast Tunnel, z - MDT-data group sender,  
Y - Joined MDT-data group, y - Sending to MDT-data group,  
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,  
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,  
Q - Received BGP S-A Route, q - Sent BGP S-A Route,  
V - RD & Vector, v - Vector, p - PIM Joins on route,  
x - VxLAN group, c - PFP-SA cache created entry
```

```
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
```

```
Timers: Uptime/Expires
```

```
Interface state: Interface, Next-Hop or VCD, State/Mode
```

```
(*, 239.255.255.250), 00:00:31/00:03:05, RP 100.100.100.1, flags: S
```

```
Incoming interface: Null, RPF nbr 0.0.0.0
```

```
Outgoing interface list:
```

```
LISP0.4099, 192.168.10.142, Forward/Sparse, 00:00:24/00:03:05
```

```
LISP0.4099, 192.168.10.150, Forward/Sparse, 00:00:31/00:02:58
```

```
(*, 224.0.1.40), 00:00:53/00:02:57, RP 100.100.100.1, flags: SJCL
```

```
Incoming interface: Null, RPF nbr 0.0.0.0
```

```
Outgoing interface list:
```

```
LISP0.4099, 192.168.10.150, Forward/Sparse, 00:00:32/00:02:57
```

```
LISP0.4099, 192.168.10.142, Forward/Sparse, 00:00:45/00:02:44
```

```
Vlan3003, Forward/Sparse, 00:00:53/00:02:32
```




Initiating multicast stream from the Fusion router (directly connected to the Fabric Border) in the Vlan belonging to the Overlay space (vrf Campus on the Fabric Border):

```
Pod1-SS#ping 224.100.100.100 source Vlan3003 repeat 500
Type escape sequence to abort.
Sending 500, 100-byte ICMP Echos to 224.100.100.100, timeout is 2 seconds:
Packet sent with a source address of 192.168.19.10
...
```

The source being registered at the RP – since Border is acting as RP, the traffic is getting pruned (no active receivers).

```
Pod1-Border1#sh ip mroute vrf Campus
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.255.255.250), 00:03:21/00:03:11, RP 100.100.100.1, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
  LISP0.4099, 192.168.10.142, Forward/Sparse, 00:03:15/00:03:11
  LISP0.4099, 192.168.10.150, Forward/Sparse, 00:03:21/00:03:06

(*, 224.100.100.100), 00:00:20/stopped, RP 100.100.100.1, flags: SPF
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list: Null

(192.168.19.10, 224.100.100.100), 00:00:20/00:02:41, flags: PFT
Incoming interface: Vlan3003, RPF nbr 0.0.0.0
Outgoing interface list: Null -----> No outgoing interface

(*, 224.0.1.40), 00:03:43/00:03:04, RP 100.100.100.1, flags: SJCL
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
  LISP0.4099, 192.168.10.150, Forward/Sparse, 00:03:22/00:03:04
  LISP0.4099, 192.168.10.142, Forward/Sparse, 00:03:35/00:02:51
  Vlan3003, Forward/Sparse, 00:03:43/00:02:47
```

Multicast traffic forwarding at the Fabric Border (C9500):

```
Pod1-Border1#sh ip mfib vrf Campus 224.100.100.100 192.168.19.10 verbose
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
ET - Data Rate Exceeds Threshold, K - Keepalive
DDE - Data Driven Event, HW - Hardware Installed
ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
MS - MoFRR Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
NS - Negate Signalling, SP - Signal Present,
A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
MA - MFIB Accept, A2 - Accept backup,
RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts: Total/RPF failed/Other drops
I/O Item Counts: FS Pkt Count/PS Pkt Count
VRF Campus
(192.168.19.10,224.100.100.100) Flags: K HW DDE
0xD0 OIF-IC count: 0, OIF-A count: 1
```



```
SW Forwarding: 1/0/100/0, Other: 0/0/0
HW Forwarding: 344/0/100/0, Other: 0/0/0 -----> Hardware Forwarding counters
Vlan3003 Flags: RA A MA -----> A = Accept = Incoming interface
-----> No outgoing interface
```

Pod1-Edge2#sh ip mroute vrf Campus

IP Multicast Routing Table

Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,

L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry

Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join

Timers: Uptime/Expires

Interface state: Interface, Next-Hop or VCD, State/Mode

(* , 239.255.255.250), 01:14:18/00:02:36, RP 100.100.100.1, flags: SJC

Incoming interface: LISPO.4099, RPF nbr 192.168.10.1

Outgoing interface list:

Vlan1022, Forward/Sparse, 01:14:18/00:02:36

(* , 224.0.1.40), 01:14:21/00:02:40, RP 100.100.100.1, flags: SJCL

Incoming interface: LISPO.4099, RPF nbr 192.168.10.1

Outgoing interface list:

Vlan1022, Forward/Sparse, 01:14:21/00:02:40

User sitting behind Fabric Edge2 subscribes to the multicast stream (224.100.100.100):

Pod1-Edge2#sh ip mroute vrf Campus 224.100.100.100

IP Multicast Routing Table

Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,

L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry

Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join

Timers: Uptime/Expires

Interface state: Interface, Next-Hop or VCD, State/Mode

(* , 224.100.100.100), 00:00:14/stopped, RP 100.100.100.1, flags: SJC

Incoming interface: LISPO.4099, RPF nbr 192.168.10.1

Outgoing interface list:

Vlan1022, Forward/Sparse, 00:00:14/00:02:45

(192.168.19.10, 224.100.100.100), 00:00:13/00:02:46, flags: JT

Incoming interface: LISPO.4099, RPF nbr 192.168.10.200

Outgoing interface list:

Vlan1022, Forward/Sparse, 00:00:13/00:02:46

Pod1-Edge2#

Jan 30 17:00:19.425: PIM(17): Check RP 100.100.100.1 into the (, 224.100.100.100) entry

Jan 30 17:00:19.425: PIM(17): Building Triggered (,G) Join / (S,G,RP-bit) Prune message for 224.100.100.100

Jan 30 17:00:19.425: PIM(17): Upstream mode for (, 224.100.100.100) changed from 0 to 1

Jan 30 17:00:19.425: PIM(17): Insert (,224.100.100.100) join in nbr 192.168.10.1's queue

*Jan 30 17:00:19.425: PIM(17): Building Join/Prune packet for nbr 192.168.10.1



```
*Jan 30 17:00:19.425: PIM(17): Adding v2 (100.100.100.1/32, 224.100.100.100), WC-bit, RPT-bit, S-bit Join
*Jan 30 17:00:19.425: PIM(17): Adding LISP Unicast transport attribute in join/prune to 192.168.10.1 (LISPO.4099)
```

```
Pod1-Border1#
*Jan 30 17:00:19.425: PIM(6): Received v2 Join/Prune on LISPO.4099 from 100.100.100.4
*Jan 30 17:00:19.425: PIM(6): J/P Transport Attribute, Transport Type: Unicast, to us
*Jan 30 17:00:19.425: PIM(6): Join-list: (*, 224.100.100.100), RPT-bit set, WC-bit set, S-bit set
*Jan 30 17:00:19.425: PIM(6): Add LISPO.4099/192.168.10.150 to (*, 224.100.100.100), Forward state, by PIM *G Join
*Jan 30 17:00:19.425: PIM(6): Add LISPO.4099/192.168.10.150 to (192.168.19.10, 224.100.100.100), Forward state, by PIM *G Join
```

OIL list on the Border is Null and the flag indicates that the traffic switched to SPT.

```
Pod1-Border1#sh ip mroute vrf Campus 224.100.100.100
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.100.100.100), 00:08:13/00:02:49, RP 100.100.100.1, flags: SF
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    LISPO.4099, 192.168.10.150, Forward/Sparse, 00:05:40/00:02:49

(192.168.19.10, 224.100.100.100), 00:08:13/00:02:48, flags: PFT -----> T-flag indicating traffic
switched to SPT
  Incoming interface: Vlan3003, RPF nbr 0.0.0.0
  Outgoing interface list: Null
```

Traffic forwarding (Fabric Border):

```
Pod1-Border1#sh ip mfib vrf Campus 224.100.100.100 192.168.19.10 verbose
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
            ET - Data Rate Exceeds Threshold, K - Keepalive
            DDE - Data Driven Event, HW - Hardware Installed
            ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
            MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
            MS - MoFRR Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
              NS - Negate Signalling, SP - Signal Present,
              A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
              MA - MFIB Accept, A2 - Accept backup,
              RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:  FS Pkt Count/PS Pkt Count
VRF Campus
(192.168.19.10,224.100.100.100) Flags: K HW DDE
0x98 OIF-IC count: 0, OIF-A count: 1
SW Forwarding: 1/0/100/0, Other: 0/0/0
HW Forwarding: 271/0/100/0, Other: 0/0/0 -----> Hardware Forwarding counters
Vlan3003 Flags: RA A MA NS -----> A = Accept = Incoming interface
Tunnell Flags: RF F -----> F indicating traffic forwarded over Tunnell
interface
CEF: Adjacency with MAC: 4500000000000000FF672ACC64646401646464012100DEFF00000000
Pkts: 0/0
```



```
Pod1-Border1#sh int tunnelli
Tunnelli is up, line protocol is up
Hardware is Tunnel
Description: Pim Register Tunnel (Encap) for RP 100.100.100.1 on VRF Campus
Interface is unnumbered. Using address of Loopback4099 (100.100.100.1)
MTU 17864 bytes, BW 100 Kbit/sec, DLY 50000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation TUNNEL, loopback not set
Keepalive not set
Tunnel linestate evaluation up
Tunnel source 100.100.100.1 (Loopback4099), destination 100.100.100.1
Tunnel Subblocks:
  src-track:
    Tunnelli source tracking subblock associated with Loopback4099
    Set of tunnels with source Loopback4099, 1 member (includes iterators), on interface <OK>
Tunnel protocol/transport PIM/IPv4
Tunnel TTL 255
Tunnel transport MTU 1486 bytes
Tunnel is transmit only
Tunnel transmit bandwidth 8000 (kbps)
Tunnel receive bandwidth 8000 (kbps)
Last input never, output 00:15:40, output hang never
Last clearing of "show interface" counters 04:04:25
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/0 (size/max)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts (0 IP multicasts)
  0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
  3 packets output, 300 bytes, 0 underruns
  0 output errors, 0 collisions, 0 interface resets
  0 unknown protocol drops
  0 output buffer failures, 0 output buffers swapped out
```

Traffic forwarding (Fabric Edge):

```
Pod1-Edge2#sh ip mfib vrf Campus 224.100.100.100 192.168.19.10 verbose
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS - MoFRR Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:  FS Pkt Count/PS Pkt Count
VRF Campus
(192.168.19.10,224.100.100.100) Flags: K HW DDE
0xE8 OIF-IC count: 0, OIF-A count: 1
SW Forwarding: 0/0/0/0, Other: 0/0/0
HW Forwarding: 256/0/100/0, Other: 0/0/0 -----> Hardware Forwarding counters
LISPO.4099 Flags: RA A MA NS -----> A = Accept = Incoming interface
Vlan1022 Flags: RF F NS -----> F indicating traffic forwarded over Vlan1022
CEF: Adjacency with MAC: 01005E6464640000C9FF45D0800
Pkts: 0/0
```

```
Pod1-Edge2#sh ip igmp vrf Campus groups 224.100.100.100
IGMP Connected Group Membership
Group Address  Interface          Uptime    Expires    Last Reporter  Group Accounted
224.100.100.100  Vlan1022          00:13:57  00:02:01  172.16.12.101
```

```
Pod1-Edge2#sh ip igmp snooping groups vlan 1022
Vlan    Group          Type          Version    Port List
```



```
-----  
1022      224.100.100.100      igmp      v2      Gi1/0/14  
1022      239.255.255.250      igmp      v2      Gi1/0/14
```

Traffic capture (Fabric Edge):

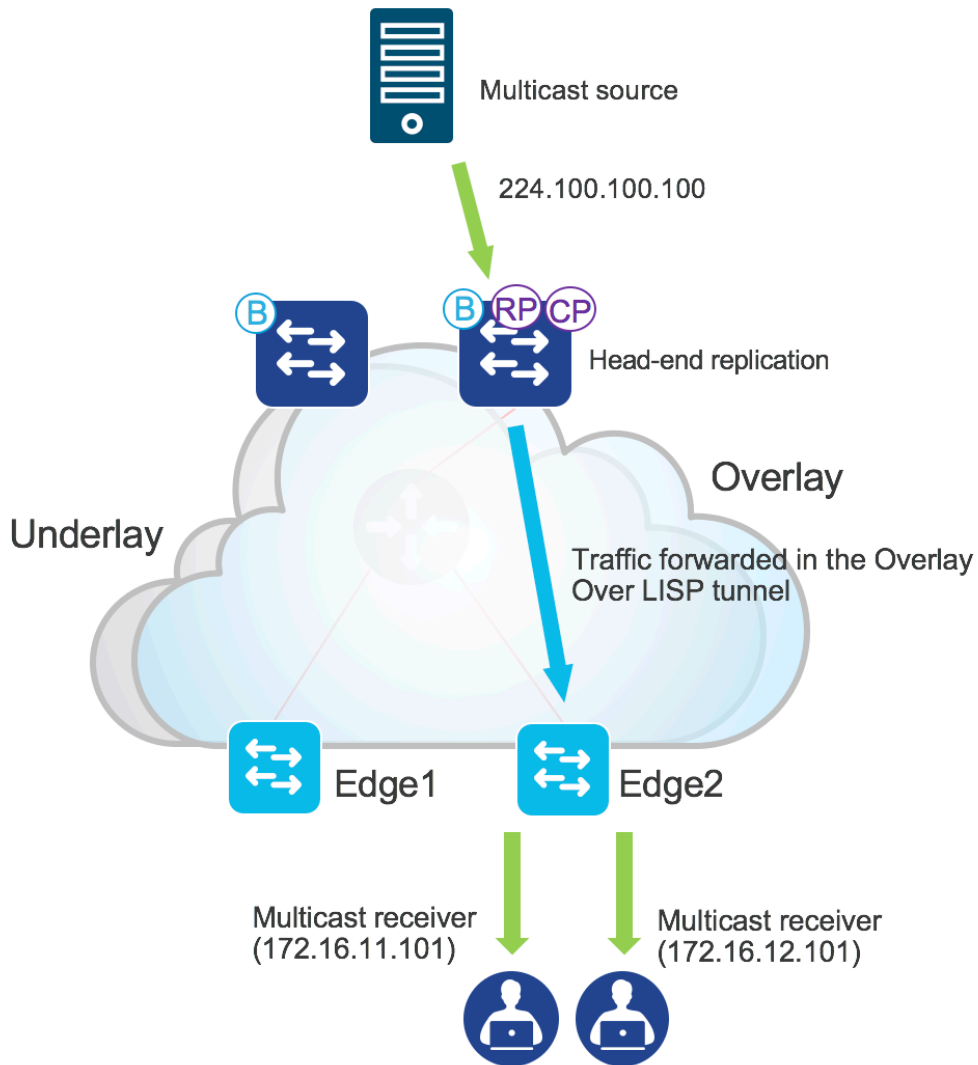
103	42.905319	192.168.19.10	224.100.100.100	ICMP	164 Echo (ping) request	id=0x0020, seq=21476/58451, ttl=254 (no...
104	42.905488	192.168.19.10	224.100.100.100	ICMP	164 Echo (ping) request	id=0x0020, seq=21477/58707, ttl=254 (no...
105	42.905781	192.168.19.10	224.100.100.100	ICMP	164 Echo (ping) request	id=0x0020, seq=21478/58963, ttl=254 (no...
106	42.905978	192.168.19.10	224.100.100.100	ICMP	164 Echo (ping) request	id=0x0020, seq=21479/59219, ttl=254 (no...
107	42.906887	192.168.19.10	224.100.100.100	ICMP	164 Echo (ping) request	id=0x0020, seq=21480/59475, ttl=254 (no...
108	42.907137	192.168.19.10	224.100.100.100	ICMP	164 Echo (ping) request	id=0x0020, seq=21481/59731, ttl=254 (no...
109	42.907415	192.168.19.10	224.100.100.100	ICMP	164 Echo (ping) request	id=0x0020, seq=21482/59987, ttl=254 (no...

- ▶ Frame 104: 164 bytes on wire (1312 bits), 164 bytes captured (1312 bits)
- ▶ Ethernet II, Src: 00:42:5a:38:b3:56 (00:42:5a:38:b3:56), Dst: 70:6b:b9:29:ee:d6 (70:6b:b9:29:ee:d6)
- ▶ Internet Protocol Version 4, Src: 192.168.10.130, Dst: 192.168.10.150
- ▶ User Datagram Protocol, Src Port: 65409 (65409), Dst Port: 4789 (4789)
- ▶ Virtual eXtensible Local Area Network
 - ▶ Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
 - Group Policy ID: 0
 - VXLAN Network Identifier (VNI): 4099
 - Reserved: 0
- ▶ Ethernet II, Src: 00:a3:d1:45:aa:72 (00:a3:d1:45:aa:72), Dst: IPv4mcast_64:64:64 (01:00:5e:64:64:64)
- ▶ Internet Protocol Version 4, Src: 192.168.19.10, Dst: 224.100.100.100
 - 0100 = Version: 4
 - ... 0101 = Header Length: 20 bytes
 - ▶ Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
 - Total Length: 100
 - Identification: 0x66d1 (26321)
 - ▶ Flags: 0x00
 - Fragment offset: 0
 - Time to live: 254
 - Protocol: ICMP (1)
 - ▶ Header checksum: 0x3d4c [validation disabled]
 - Source: 192.168.19.10
 - Destination: 224.100.100.100
 - [Source GeoIP: Unknown]
 - [Destination GeoIP: Unknown]
- ▶ Internet Control Message Protocol

The traffic is replicated at the Fabric Border for each Fabric Edge switch having receivers being subscribed to the multicast stream. Traffic is being sent using unicast VxLAN encapsulation through LISP tunnel.

PIM ASM with multicast source outside the fabric and receivers within the fabric (multiple receiver)

This case represents multicast stream being sent from outside of the fabric towards two receivers sitting in a separate VLANs (the same Virtual Network).





Multicast traffic flow verification

```
Pod1-Edge1# sh ip pim vrf Campus tunnel verbose
```

```
Tunnell
Type       : PIM Encap
RP         : 100.100.100.1
Source     : 100.100.100.3
State      : UP
Last event : Created (1d03h)
Refcnt     : 0
```

```
Pod1-Edge2#sh ip pim vrf Campus tunnel verbose
```

```
Tunnell
Type       : PIM Encap
RP         : 100.100.100.1
Source     : 100.100.100.4
State      : UP
Last event : Created (1d03h)
Refcnt     : 0
```

```
Pod1-Edge2#sh ip rpf vrf Campus 100.100.100.1
```

```
RPF information for ? (100.100.100.1)
RPF interface: LISPO.4099
RPF neighbor: ? (192.168.10.1)
RPF route/mask: 100.100.100.1/32
RPF type: unicast ()
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base
```

```
Pod1-Edge2#sh ip mroute vrf Campus 224.100.100.100
```

```
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
```

```
(*, 224.100.100.100), 01:47:45/stopped, RP 100.100.100.1, flags: SJC
```

```
Incoming interface: LISPO.4099, RPF nbr 192.168.10.1
```

```
Outgoing interface list:
```

```
Vlan1021, Forward/Sparse, 00:00:15/00:02:44
```

```
Vlan1022, Forward/Sparse, 01:47:45/00:02:10
```

```
(192.168.19.10, 224.100.100.100), 00:01:50/00:01:09, flags: JT
```

```
Incoming interface: LISPO.4099, RPF nbr 192.168.10.200
```

```
Outgoing interface list:
```

```
Vlan1021, Forward/Sparse, 00:00:15/00:02:44
```

```
Vlan1022, Forward/Sparse, 00:01:50/00:02:10
```

```
Pod1-Border1#sh ip mroute vrf Campus 224.100.100.100
```

```
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry
```



Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(* , 224.100.100.100), 00:03:35/00:02:42, RP 100.100.100.1, flags: SF
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
LISP0.4099, 192.168.10.150, Forward/Sparse, 00:02:46/00:02:42

(192.168.19.10, 224.100.100.100), 00:03:35/00:01:30, flags: PFT
Incoming interface: Vlan3003, RPF nbr 0.0.0.0
Outgoing interface list: Null

Pod1-Edge2# sh ip mfib vrf Campus 224.100.100.100 192.168.19.10 verbose
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
ET - Data Rate Exceeds Threshold, K - Keepalive
DDE - Data Driven Event, HW - Hardware Installed
ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
MS - MoFRR Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
NS - Negate Signalling, SP - Signal Present,
A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
MA - MFIB Accept, A2 - Accept backup,
RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second

Other counts: Total/RPF failed/Other drops

I/O Item Counts: FS Pkt Count/PS Pkt Count

VRF Campus

(192.168.19.10,224.100.100.100) Flags: K HW DDE

0xA7 OIF-IC count: 0, OIF-A count: 1

SW Forwarding: 0/0/0/0, Other: 0/0/0

HW Forwarding: 108/0/100/0, Other: 0/0/0 -----> Hardware Forwarding counters

LISP0.4099 Flags: RA A MA -----> A = Accept = Incoming interface

Vlan1021 Flags: RF F NS -----> F indicating traffic forwarded over Vlan1021

CEF: Adjacency with MAC: 01005E6464640000C9FF45C0800

Pkts: 0/0

Vlan1022 Flags: RF F NS -----> F indicating traffic forwarded over Vlan1022

CEF: Adjacency with MAC: 01005E6464640000C9FF45D0800

Pkts: 0/0

```
Pod1-Edge2# sh ip igmp snooping groups
```

Vlan	Group	Type	Version	Port List
1021	224.100.100.100	igmp	v3	Gi1/0/15
1021	239.255.255.250	igmp	v3	Gi1/0/15
1022	224.100.100.100	igmp	v2	Gi1/0/14
1022	239.255.255.250	igmp	v2	Gi1/0/14



There is a single unicast packet sent via Fabric Border switch towards Fabric Edge 2 switch.

```

> Frame 774: 164 bytes on wire (1312 bits), 164 bytes captured (1312 bits)
> Ethernet II, Src: 00:42:5a:38:b3:56 (00:42:5a:38:b3:56), Dst: 70:6b:b9:29:ee:d6 (70:6b:b9:29:ee:d6)
> Internet Protocol Version 4, Src: 192.168.10.130, Dst: 192.168.10.150
  User Datagram Protocol, Src Port: 65409 (65409), Dst Port: 4789 (4789)
    Source Port: 65409
    Destination Port: 4789
    Length: 130
    > Checksum: 0x0000 (none)
      [Stream index: 1]
  Virtual eXtensible Local Area Network
    > Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
      Group Policy ID: 0
      VXLAN Network Identifier (VNI): 4099
      Reserved: 0
    > Ethernet II, Src: 40:00:00:11:aa:72 (40:00:00:11:aa:72), Dst: IPv4mcast_64:64:64 (01:00:5e:64:64:64)
  Internet Protocol Version 4, Src: 192.168.19.10, Dst: 224.100.100.100
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes
    > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
      Total Length: 100
      Identification: 0x8210 (33296)
    > Flags: 0x00
      Fragment offset: 0
      Time to live: 254
      Protocol: ICMP (1)
    > Header checksum: 0x220d [validation disabled]
      Source: 192.168.19.10
      Destination: 224.100.100.100
      [Source GeoIP: Unknown]
      [Destination GeoIP: Unknown]
  Internet Control Message Protocol

```

The multicast stream is replicated at the Fabric Edge 2 switch and sent towards both receivers.

```

> Frame 38: 114 bytes on wire (912 bits), 114 bytes captured (912 bits)
> Ethernet II, Src: CiscoInc_9f:f4:5c (00:00:0c:9f:f4:5c), Dst: IPv4mcast_64:64:64 (01:00:5e:64:64:64)
  Internet Protocol Version 4, Src: 192.168.19.10, Dst: 224.100.100.100
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes
    > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
      Total Length: 100
      Identification: 0x1b58 (7000)
    > Flags: 0x00
      Fragment offset: 0
      Time to live: 252
      Protocol: ICMP (1)
    > Header checksum: 0x8ac5 [validation disabled]
      Source: 192.168.19.10
      Destination: 224.100.100.100
      [Source GeoIP: Unknown]
      [Destination GeoIP: Unknown]
  Internet Control Message Protocol

> Frame 39: 114 bytes on wire (912 bits), 114 bytes captured (912 bits)
> Ethernet II, Src: CiscoInc_9f:f4:5d (00:00:0c:9f:f4:5d), Dst: IPv4mcast_64:64:64 (01:00:5e:64:64:64)
  Internet Protocol Version 4, Src: 192.168.19.10, Dst: 224.100.100.100
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes
    > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
      Total Length: 100
      Identification: 0x1b59 (7001)
    > Flags: 0x00
      Fragment offset: 0
      Time to live: 252
      Protocol: ICMP (1)
    > Header checksum: 0x8ac4 [validation disabled]
      Source: 192.168.19.10
      Destination: 224.100.100.100
      [Source GeoIP: Unknown]
      [Destination GeoIP: Unknown]
  Internet Control Message Protocol

```

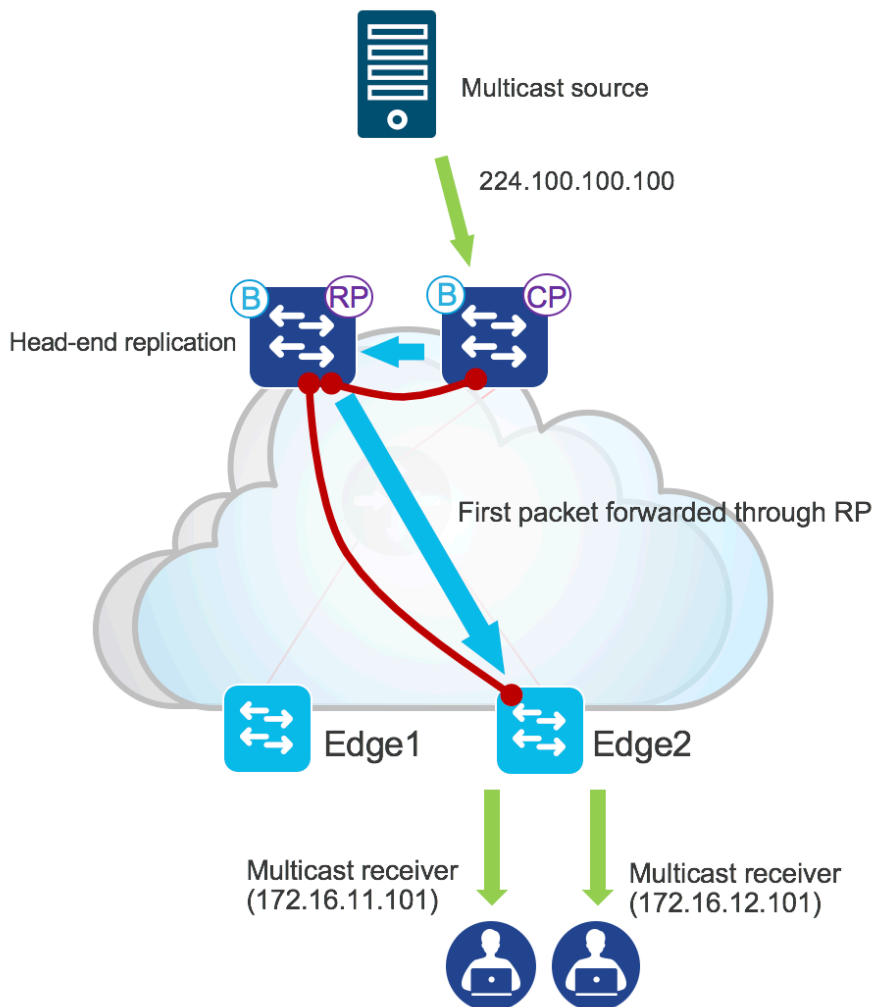
Pod1-Edge2#sh int vlan 1021 | i Hardware

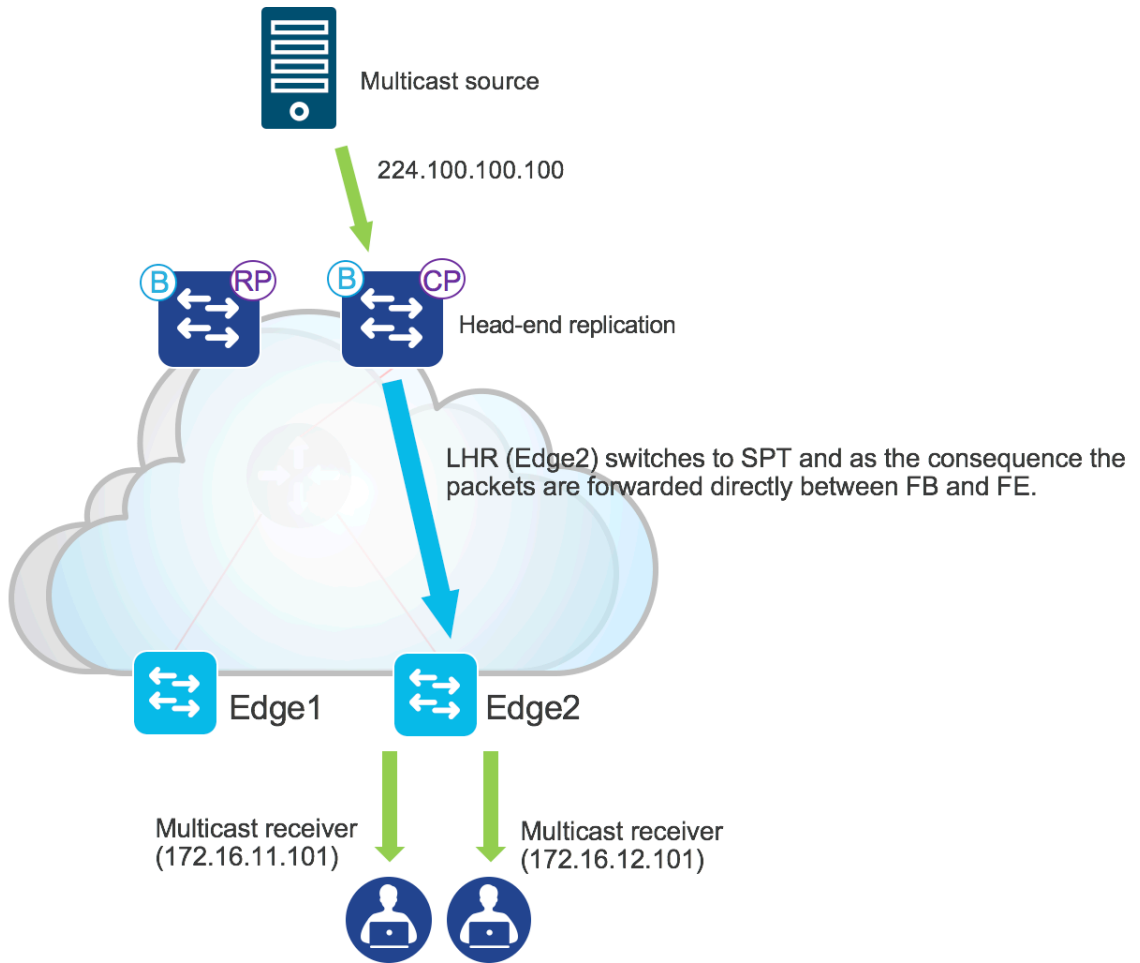


```
Hardware is Ethernet SVI, address is 0000.0c9f.f45c (bia 706b.b929.eef5)
Pod1-Edge2#sh int vlan 1022 | i Hardware
Hardware is Ethernet SVI, address is 0000.0c9f.f45d (bia 706b.b929.eef5)
```

PIM ASM with multicast source outside the fabric and receivers within the fabric (multiple receivers) / RP on a separate Border node

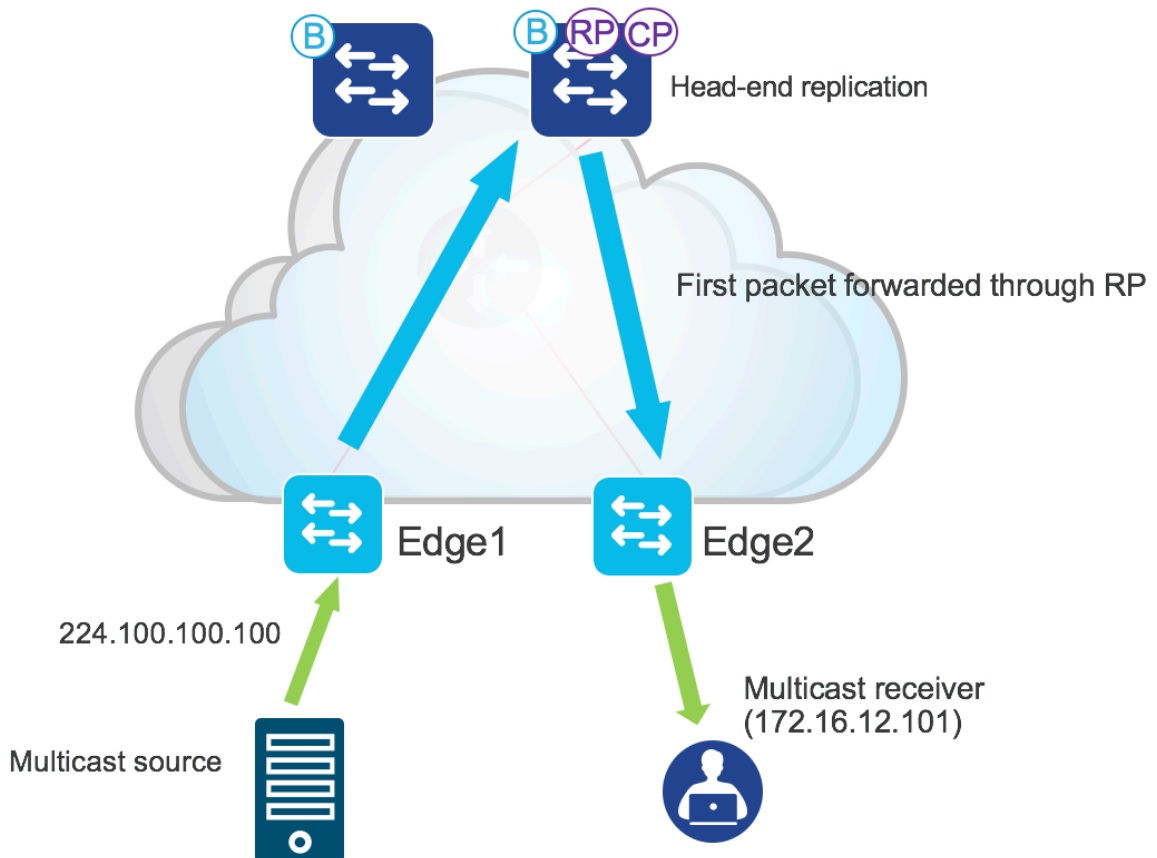
In this scenario, the Rendezvous point is located on the separate Border Fabric node. The very first packet is forwarded through the RP (shared tree) and all subsequent packets are directly routed between the ingressing Fabric Border and Fabric Edge switches.





PIM PIM ASM with multicast source and receivers within the fabric

In this scenario, the source and receiver are located within the Fabric domain on two separate Fabric Edge switches. Both are sitting in the same VLAN.





```
Pod1-Edge2#sh ip pim vrf Campus tunnel verbose
```

```
Tunnell
Type      : PIM Encap
RP        : 100.100.100.1
Source    : 100.100.100.4
State     : UP
Last event : Created (1d03h)
Refcnt    : 0
Pod1-Edge2#sh ip rp
```

```
Pod1-Edge2#sh ip rpf vrf Campus 100.100.100.1
```

```
RPF information for ? (100.100.100.1)
RPF interface: LISP0.4099
RPF neighbor: ? (192.168.10.1)
RPF route/mask: 100.100.100.1/32
RPF type: unicast ()
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base
```

```
Pod1-Edge1#sh ip mroute vrf Campus
```

```
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD and Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.255.255.250), 10:13:11/00:02:52, RP 100.100.100.1, flags: SJC
Incoming interface: LISP0.4099, RPF nbr 192.168.10.1
Outgoing interface list:
Vlan1022, Forward/Sparse, 10:13:10/00:02:52

(*, 224.0.1.40), 10:13:11/00:02:50, RP 100.100.100.1, flags: SJCL
Incoming interface: LISP0.4099, RPF nbr 192.168.10.1
Outgoing interface list:
Loopback4099, Forward/Sparse, 10:13:10/00:02:50
```

```
Pod1-Edge2#sh ip mroute vrf Campus
```

```
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD and Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.255.255.250), 02:25:57/00:02:05, RP 100.100.100.1, flags: SJC
Incoming interface: LISP0.4099, RPF nbr 192.168.10.1
Outgoing interface list:
Vlan1022, Forward/Sparse, 02:25:55/00:02:00
Vlan1021, Forward/Sparse, 02:25:57/00:02:05

(*, 224.100.100.100), 02:24:21/00:02:59, RP 100.100.100.1, flags: SJC
```



```
Incoming interface: LISPO.4099, RPF nbr 192.168.10.1
Outgoing interface list:
  Vlan1022, Forward/Sparse, 02:23:56/00:02:03
  Vlan1021, Forward/Sparse, 02:24:21/00:02:59

(*, 224.0.1.40), 02:26:17/00:01:58, RP 100.100.100.1, flags: SJCL
Incoming interface: LISPO.4099, RPF nbr 192.168.10.1
Outgoing interface list:
  Loopback4099, Forward/Sparse, 02:26:17/00:01:58
```

```
Pod1-Edge2#sh ip mfib vrf Campus 224.100.100.100/32 verbose
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
             ET - Data Rate Exceeds Threshold, K - Keepalive
             DDE - Data Driven Event, HW - Hardware Installed
             ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
             MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
             MS - MoFRR Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
               NS - Negate Signalling, SP - Signal Present,
               A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
               MA - MFIB Accept, A2 - Accept backup,
               RA2 - MRIB Accept backup, MA2 - MFIB Accept backup
```

```
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
```

```
Other counts: Total/RPF failed/Other drops
```

```
I/O Item Counts: FS Pkt Count/PS Pkt Count
```

```
VRF Campus
```

```
(* ,224.100.100.100) Flags: C K HW
```

```
0x96 OIF-IC count: 0, OIF-A count: 1
```

```
SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
HW Forwarding: 60/0/100/0, Other: 0/0/0 -----> Hardware Forwarding counters
```

```
LISPO.4099 Flags: RA A MA NS -----> A = Accept = Incoming interface
```

```
Vlan1021 Flags: RF F NS -----> F indicating traffic to be forwarded over Vlan1021
```

```
CEF: Adjacency with MAC: 01005E6464640000C9FF45C0800
```

```
Pkts: 0/0
```

```
Vlan1022 Flags: RF F NS -----> F indicating traffic to be forwarded over Vlan1022
```

```
CEF: Adjacency with MAC: 01005E6464640000C9FF45D0800
```

```
Pkts: 0/0
```

After initiating multicast stream (224.100.100.100):

```
Pod1-Edge2#sh ip mroute vrf Campus 224.100.100.100
```

```
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry
```

```
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
```

```
Timers: Uptime/Expires
```

```
Interface state: Interface, Next-Hop or VCD, State/Mode
```

```
(* , 224.100.100.100), 02:28:38/stopped, RP 100.100.100.1, flags: SJC
```

```
Incoming interface: LISPO.4099, RPF nbr 192.168.10.1
```

```
Outgoing interface list:
```

```
  Vlan1022, Forward/Sparse, 02:28:13/00:02:43
```

```
  Vlan1021, Forward/Sparse, 02:28:38/00:02:40
```

```
(172.16.12.100, 224.100.100.100), 00:01:12/00:01:47, flags: JT
```

```
Incoming interface: LISPO.4099, RPF nbr 192.168.10.3
```

```
Outgoing interface list:
```

```
  Vlan1022, Forward/Sparse, 00:01:12/00:02:43
```

```
  Vlan1021, Forward/Sparse, 00:01:12/00:02:40
```

```
Pod1-Edge2#sh ip mroute vrf Campus 224.100.100.100 count
```

```
Use "show ip mfib count" to get better response time for a large number of mroutes.
```



IP Multicast Statistics

4 routes using 5184 bytes of memory
3 groups, 0.33 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 224.100.100.100, Source count: 1, Packets forwarded: 83, Packets received: 83
RP-tree: Forwarding: 61/0/99/0, Other: 61/0/0
Source: 172.16.12.100/32, Forwarding: 22/0/60/0, Other: 22/0/0
Pod1-Edge2#

The very first packet is forwarded through the Shared Tree and by default when the Last Hop Router (Edge2) receives the packet, it initiates the switch to Source Tree.

```
Frame 50: 124 bytes on wire (992 bits), 124 bytes captured (992 bits) on Ethernet II, Src: 00:42:5a:38:b3:56 (00:42:5a:38:b3:56), Dst: 70:6b:b9:29:ee:d6 (70:6b:b9:29:ee:d6)
  Internet Protocol Version 4, Src: 192.168.10.130, Dst: 192.168.10.150
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes
  Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
  Total Length: 110
  Identification: 0x0001 (1)
  Flags: 0x00
  Fragment offset: 0
  Time to live: 125
  Protocol: UDP (17)
  Header checksum: 0xa715 [validation disabled]
  Source: 192.168.10.130
  Destination: 192.168.10.150
  [Source GeoIP: Unknown]
  [Destination GeoIP: Unknown]
  User Datagram Protocol, Src Port: 29812 (29812), Dst Port: 4789 (4789)
  Virtual eXtensible Local Area Network
    Flags: 0x0848, Don't Learn, VXLAN Network ID (VNI), Policy Applied
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 4099
    Reserved: 0
  Ethernet II, Src: 00:a3:d1:45:aa:56 (00:a3:d1:45:aa:56), Dst: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38)
  Internet Protocol Version 4, Src: 172.16.12.100, Dst: 224.100.100.100
  Internet Control Message Protocol
```

First packet forwarded through the Shared Tree via RP using unicast

L3VNI

Multicast stream encapsulated in VxLAN header

In order to initiate the switchover to the Source Tree, Edge2 router needs to know how to get to the Multicast Source host (172.16.12.100) – it sends LISP MAP-Request for 172.16.12.100 towards the MS/MR.

```
Frame 54: 170 bytes on wire (1360 bits), 170 bytes captured (1360 bits) on Ethernet II, Src: 70:6b:b9:29:ee:d6 (70:6b:b9:29:ee:d6), Dst: 00:42:5a:38:b3:56 (00:42:5a:38:b3:56)
  Internet Protocol Version 4, Src: 192.168.10.150, Dst: 192.168.10.1
  User Datagram Protocol, Src Port: 4342 (4342), Dst Port: 4342 (4342)
  Locator/ID Separation Protocol
    1000 .... = Type: Encapsulated Control Message (8)
    .... 0... = S bit (LISP-SEC capable): Not set
    .... .0.. = D bit (DDT-originated): Not set
    .... ..00 0000 0000 0000 0000 0000 0000 = Reserved bits: 0x00000000
  Internet Protocol Version 4, Src: 172.16.12.100, Dst: 172.16.12.100
  User Datagram Protocol, Src Port: 4342 (4342), Dst Port: 4342 (4342)
  Locator/ID Separation Protocol
    0001 .... = Type: Map-Request (1)
  Flags: 0x000012
  Reserved bits: 0x000000
  ITR-RLOC Count: 0
  Record Count: 1
  Nonce: 0xf0a850a03b11d116
  Source EID AFI: LISP Canonical Address Format (LCAF) (16387)
  Source EID: [4099] 100.100.100.4
  ITR-RLOC 1: 192.168.10.4
  Map-Request Record 1: [4099] 172.16.12.100/32
  Map-Reply Record
```

MS replies with RLOC.



```
‣ Frame 55: 94 bytes on wire (752 bits), 94 bytes captured (752 bits)
‣ Ethernet II, Src: 00:42:5a:38:b3:56 (00:42:5a:38:b3:56), Dst: 70:6b:b9:29:ee:d6 (70:6b:b9:29:ee:d6)
‣ Internet Protocol Version 4, Src: 192.168.10.130, Dst: 192.168.10.4
‣ User Datagram Protocol, Src Port: 4342 (4342), Dst Port: 4342 (4342)
‣ Locator/ID Separation Protocol
  0010 .... = Type: Map-Reply (2)
  .... 0... = P bit (Probe): Not set
  .... .0.. = E bit (Echo-Nonce locator reachability algorithm enabled): Not set
  .... ..0. = S bit (LISP-SEC capable): Not set
  .... ...0 0000 0000 0000 0000 = Reserved bits: 0x000000
  Record Count: 1
  Nonce: 0xf0a850a03b11d116
  ‣ Mapping Record 1, EID Prefix: [4099] 172.16.12.100/32, TTL: 1440, Action: No-Action, Not Authoritative
    Record TTL: 1440
    Locator Count: 1
    EID Mask Length: 32
    000. .... = Action: No-Action (0)
    ...0 .... = Authoritative bit: Not set
    .... .000 0000 0000 = Reserved: 0x0000
    0000 .... = Reserved: 0x0000
    .... 0000 0000 0000 = Mapping Version: 0
    EID Prefix AFI: LISP Canonical Address Format (LCAF) (16387)
  ‣ EID Prefix: [4099] 172.16.12.100
  ‣ Locator Record 1, RLOC: 192.168.10.3, Reachable, Priority/Weight: 10/10, Multicast Priority/Weight: 10/10
```

Edge2 (LHR) sends PIM join encapsulated in the unicast VxLAN header towards Edge2, where source resides.

```
‣ Frame 56: 128 bytes on wire (1024 bits), 128 bytes captured (1024 bits)
‣ Ethernet II, Src: CiscoInc 9f:ee:d6 (00:00:0c:9f:ee:d6), Dst: 00:42:5a:38:b3:56 (00:42:5a:38:b3:56)
‣ Internet Protocol Version 4, Src: 192.168.10.146, Dst: 192.168.10.3
‣ User Datagram Protocol, Src Port: 28013 (28013), Dst Port: 4789 (4789)
‣ Virtual eXtensible Local Area Network
‣ Ethernet II, Src: 70:6b:b9:29:ee:e4 (70:6b:b9:29:ee:e4), Dst: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38)
‣ Internet Protocol Version 4, Src: 100.100.100.4, Dst: 224.0.0.13
‣ Protocol Independent Multicast
```

Followed by PIM prune towards the RP (leaving Shared-Tree):

```
‣ Frame 57: 136 bytes on wire (1088 bits), 136 bytes captured (1088 bits)
‣ Ethernet II, Src: CiscoInc 9f:ee:d6 (00:00:0c:9f:ee:d6), Dst: 00:42:5a:38:b3:56 (00:42:5a:38:b3:56)
‣ Internet Protocol Version 4, Src: 192.168.10.150, Dst: 192.168.10.1
‣ User Datagram Protocol, Src Port: 28013 (28013), Dst Port: 4789 (4789)
‣ Virtual eXtensible Local Area Network
‣ Ethernet II, Src: 70:6b:b9:29:ee:d6 (70:6b:b9:29:ee:d6), Dst: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38)
‣ Internet Protocol Version 4, Src: 100.100.100.4, Dst: 224.0.0.13
‣ Protocol Independent Multicast
```

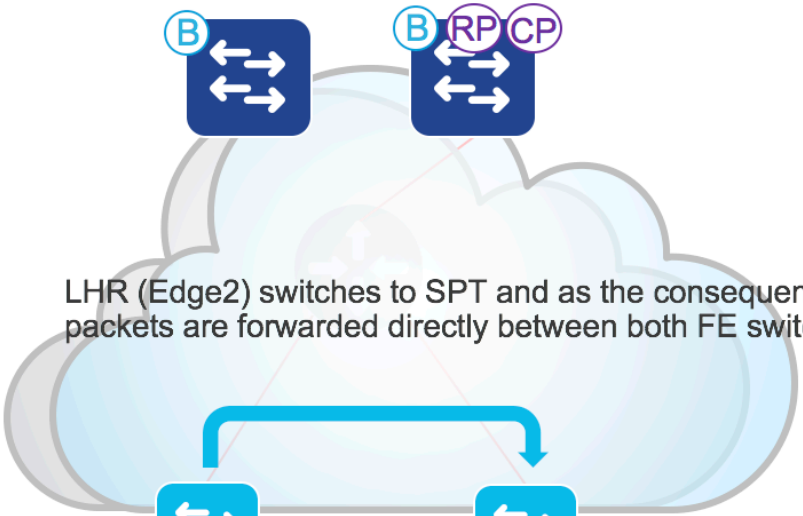
The traffic is then forwarded directly between two Edge switches using Source Tree:

```


> Frame 68: 124 bytes on wire (992 bits), 124 bytes captured (992 bits)
> Ethernet II, Src: 00:42:5a:38:b3:56 (00:42:5a:38:b3:56), Dst: 70:6b:b9:29:ee:e4 (70:6b:b9:29:ee:e4)
< Internet Protocol Version 4, Src: 192.168.10.142, Dst: 192.168.10.146
  0100 .... = Version: 4
  .... 0101 = Header Length: 20 bytes
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
  Total Length: 110
  Identification: 0x451b (17691)
  > Flags: 0x02 (Don't Fragment)
  Fragment offset: 0
  Time to live: 126
  Protocol: UDP (17)
  > Header checksum: 0x20f3 [validation disabled]
  Source: 192.168.10.142
  Destination: 192.168.10.146
  [Source GeoIP: Unknown]
  [Destination GeoIP: Unknown]
  > User Datagram Protocol, Src Port: 65485 (65485), Dst Port: 4789 (4789)
  < Virtual eXtensible Local Area Network
    > Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
    Group Policy ID: 4
    < VXLAN Network Identifier (VNI): 4099
    Reserved: 0
  > Ethernet II, Src: CiscoInc_9f:00:00 (00:00:0c:9f:00:00), Dst: Ipv4mcast_64:64:64 (01:00:5e:64:64:64)
  < Internet Protocol Version 4, Src: 172.16.12.100, Dst: 224.100.100.100
  < Internet Control Message Protocol
  
```

Unicast traffic using head-end replication on Edge1

Encapsulated multicast stream in the VxLAN header

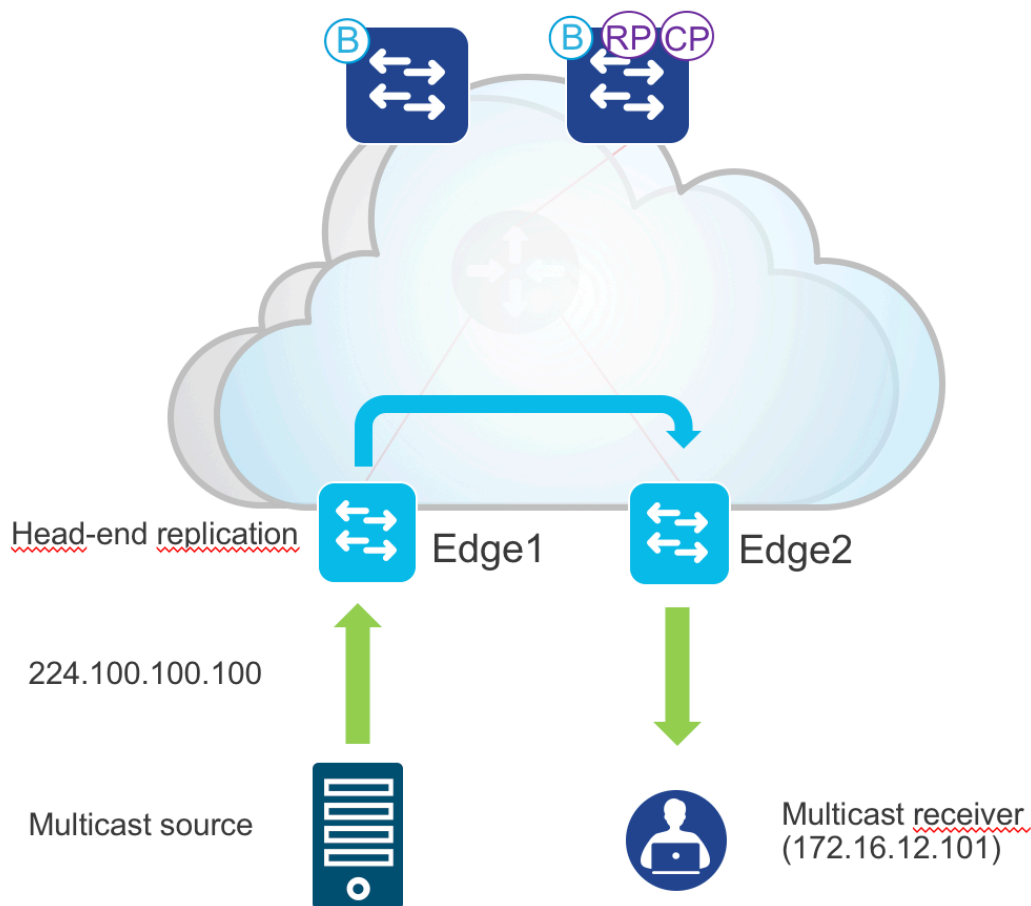


LHR (Edge2) switches to SPT and as the consequence the packets are forwarded directly between both FE switches.



PIM SSM with multicast source and receivers within the fabric

This scenario describes the situation where source and receiver are located within the Fabric domain on two separate Fabric Edge switches. Both are sitting in the same VLAN and using IGMP v3 and PIM SSM.



In order to use PIM SSM it is required to have IGMP v3 enabled on the VLAN interface facing the receivers. This is enabled by default during multicast provisioning. There is no additional configuration required.

```
Pod1-Edge2#sh run int vlan 1021 | i version
ip igmp version 3
```

```
Pod1-Edge2#sh ip igmp vrf Campus int vlan 1021 | i version
Current IGMP host version is 3
Current IGMP router version is 3
```



Pod1-Edge2#

```
*Jan 31 20:06:02.733: IGMP(17): Received v3 Report for 1 group on Vlan1021 from 172.16.11.101
*Jan 31 20:06:02.733: IGMP(17): Received Group record for group 224.40.40.40, mode 5 from
172.16.11.101 for 1 sources
*Jan 31 20:06:02.733: IGMP(17): WAVL Insert group: 224.40.40.40 interface: Vlan1021 Successful
*Jan 31 20:06:02.733: IGMP(17): Create source 172.16.12.100
*Jan 31 20:06:02.733: IGMP(17): Updating expiration time on (172.16.12.100,224.40.40.40) to 180 secs
*Jan 31 20:06:02.733: IGMP(17): Setting source flags 4 on (172.16.12.100,224.40.40.40)
*Jan 31 20:06:02.733: IGMP(17): MRT Add/Update Vlan1021 for (*,224.40.40.40) by 0
Pod1-Edge2#
```