

# Configure and Troubleshoot Multicast Service Reflection on Catalyst 8000

## Contents

---

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Topology](#)

[Configuration](#)

[Initial Verification](#)

[Troubleshooting](#)

[Related Information](#)

---

## Introduction

This document describes the process to troubleshoot Multicast Service Reflection feature on Catalyst 8000.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- Knowledge of Basic Multicast Routing
- Knowledge of Cisco IOS® XE command line

### Components Used

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

- Router Catalyst 8500
- Cisco IOS XE 17.9.4

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.

## Background Information

The purpose of this document is to demonstrate the basic configuration and troubleshooting of the Multicast Service Reflection feature to validate a Multicast-to-Multicast destination translation scenario with a Cisco Catalyst 8000.


Multicast Service Reflection is a versatile feature that can be deployed at various points within a network to

meet specific multicast requirements. It helps enhance the efficiency and reliability of multicast services in complex network environments as it allows users to translate received multicast or unicast destination addresses to multicast or unicast addresses that conform to their internal addressing policy.

Multicast Service Reflection has the ability to translate both source and destination addresses in the IPv4 header. It identifies the multicast traffic streams (associated with specific multicast groups and sources) that need to be translated (based on the predefined configuration) and when the traffic reaches the reflection point, the router doing the translation selectively reflects these packets.

One of the core components of this feature is the Virtual Interface (Vif). This Vif interface is a logical interface that can statically join the group that needs to be reflected in order to build the multicast tree. In other words, this Vif interface acts as multicast receiver for the original stream and once the traffic is forwarded to the Vif interface, the packets are reflected and the source/destination addresses are translated to the new desired IPv4 addresses.

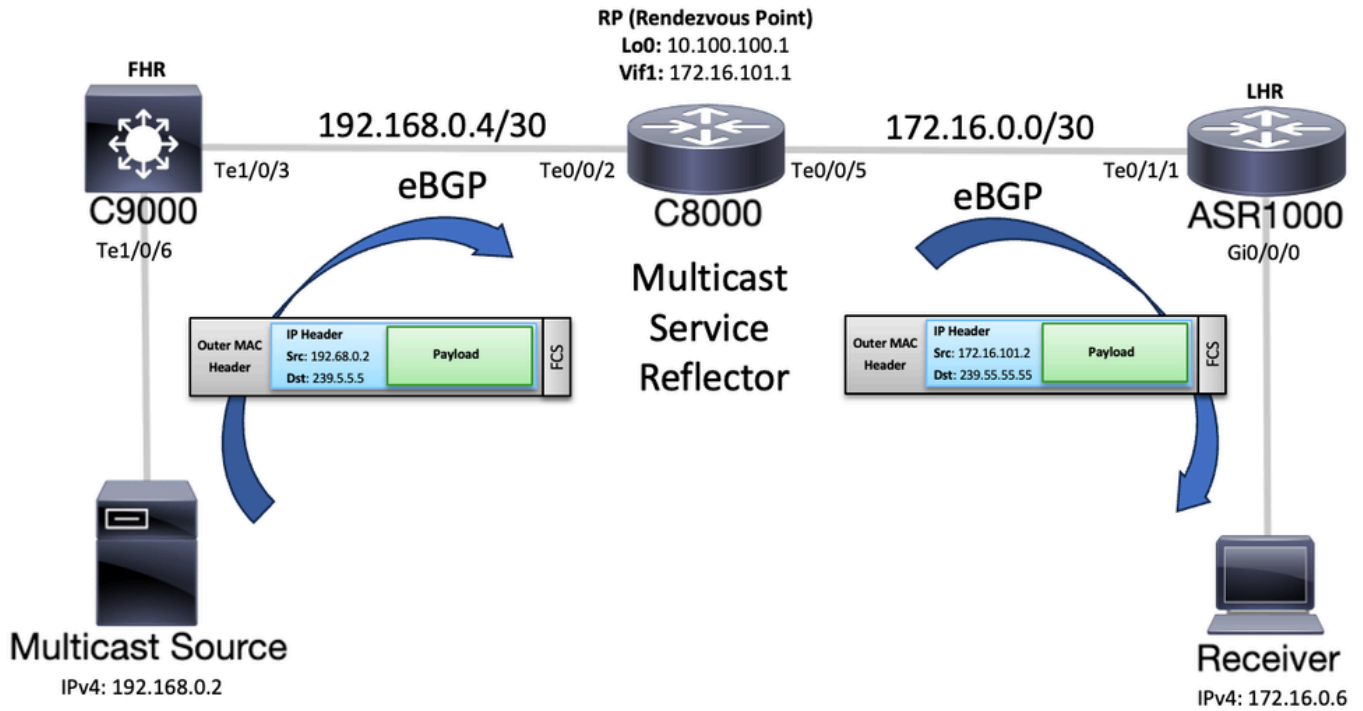
---



**Note:** To implement this feature your router must have at least a **network-advantage** license installed, otherwise an error can be seen when trying to configure the Vif interface.

---

## Topology



Topology Diagram

## Configuration

To configure the feature in this setup multicast routing needs to be enabled across the path and basic IPv4 connectivity between the Multicast Source and the Receiver is required as well. In this scenario, there is a Catalyst 9K switch (C9000) acting as a First Hop Router (FHR), then the Catalyst 8K router (C8000) that is the Multicast Service Reflector and also the Rendezvous Point (RP) for the different multicast groups and an ASR1K router (ASR1000) is the Last Hop Router (LHR) where the receiver is connected. For this example, external Border Gateway Protocol (eBGP) is used to interconnect these devices, however there is no requirement to implement a particular routing protocol.



**Note:** This example makes use of only one RP for the various multicast streams, for more information on scenarios where different RPs are used for the multicast domains review the section [Configuration Examples for Multicast Service Reflection](#) from the IP Multicast Configuration Guide, Cisco IOS XE 17.x

---

## FHR (C9000)

To start the configuration at the FHR, enable multicast routing, add the RP IPv4 address and configure Protocol Independent Multicast Sparse Mode (PIM-SM) on interfaces connected to the Multicast Source and to the Multicast Service Reflector router (C8000).

```
<#root>
```

```
C9000#
```

```
show run | i multicast|rp-address
```

```
ip multicast-routing
```

```
class-map match-any system-cpp-police-multicast-end-station
```

```
class-map match-any system-cpp-police-multicast
```

```
ip pim rp-address 10.100.100.1
```

```
C9000#
```

```
show run interface Ten 1/0/6
```

```
Building configuration...
```

```
Current configuration : 116 bytes
```

```
!
```

```
interface TenGigabitEthernet1/0/6
```

```
no switchport
```

```
ip address 192.168.0.1 255.255.255.252
```

```
ip pim sparse-mode
```

```
end
```

```
C9000#
```

```
show run interface Ten 1/0/3
```

```
Building configuration...
```

```
Current configuration : 116 bytes
```

```
!
```

```
interface TenGigabitEthernet1/0/3
```

```
no switchport
```

```
ip address 192.168.0.5 255.255.255.252
```

```
ip pim sparse-mode
```

```
end
```

The C9000 switch is doing an eBGP adjacency to the C8000 router to advertise the connected routes.

```
<#root>
```

```
C9000#
```

```
show run | section router bgp
```

```
router bgp 65003
```

```
bgp log-neighbor-changes
```

```
redistribute connected
```

```
neighbor 192.168.0.6 remote-as 65001
```

## **Multicast Service Reflector Router (C8000)**

At the Catalyst 8K router, for the initial configuration enable multicast routing, add the appropriate RP IPv4 address and also configure PIM-SM on the interfaces connected to C9000 and ASR1000 devices as well as

in the Loopback0 interface.

```
<#root>
```

```
Mcast_SR#
```

```
show run | i multicast|rp-address  
ip multicast-routing distributed  
ip pim rp-address 10.100.100.1
```

```
Mcast_SR#
```

```
show run interface Ten 0/0/2
```

```
Building configuration...
```

```
Current configuration : 156 bytes
```

```
!
```

```
interface TenGigabitEthernet0/0/2  
description to C9000  
ip address 192.168.0.6 255.255.255.252
```

```
ip pim sparse-mode
```

```
no negotiation auto  
cdp enable  
end
```

```
Mcast_SR#
```

```
show run interface Ten 0/0/5
```

```
Building configuration...
```

```
Current configuration : 157 bytes
```

```
!
```

```
interface TenGigabitEthernet0/0/5  
description to ASR1000  
ip address 172.16.0.1 255.255.255.252
```

```
ip pim sparse-mode
```

```
no negotiation auto  
cdp enable  
end
```

```
Mcast_SR#
```

```
show run interface loopback 0
```

```
Building configuration...
```

```
Current configuration : 88 bytes
```

```
!
```

```
interface Loopback0  
ip address 10.100.100.1 255.255.255.255
```

```
ip pim sparse-mode
```

```
end
```

This router has two eBGP neighbors, 192.168.0.5 neighbor is the C9000 switch sending the prefixes to reach the Multicast Source and the second neighbor is the ASR1000 router used to reach the Receiver. Notice that this router is advertising its local Loopback0 and also advertises the Vif subnet.

```
<#root>
```

```
Mcast_SR#
```

```
show run | section router bgp
```

```
router bgp 65001
  bgp log-neighbor-changes
```

```
network 10.100.100.1 mask 255.255.255.255
network 172.16.101.0 mask 255.255.255.252
```

```
neighbor 172.16.0.2 remote-as 65002
neighbor 192.168.0.5 remote-as 65003
```

Now, it is time to configure the Vif interface, which is the one that helps to translate the multicast group from 239.5.5.5 to 239.55.55.55. This logical interface needs to have an IP address and this subnet needs to be routable, additionally PIM-SM must be configured. To enable the Multicast Service Reflection feature, the command **ip service reflect <interface\_id> destination <destination\_ip1> to <destination\_ip2> mask-len <subnet\_mask\_length> source <source\_ip>** is used, where the *interface\_id* is the interface receiving the original multicast traffic, the *destination\_ip1* is the original IPv4 address that needs to be translated (239.5.5.5), and *destination\_ip2* is the new IPv4 address that is expected after the translation (239.55.55.55). Last but not least, the command **ip igmp static-group <group-address>** is required in order to make this interface a statically connected member of the original multicast group.

```
<#root>
```

```
Mcast_SR#
```

```
configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
Mcast_SR(config)#
```

```
interface Vif1
```

```
Mcast_SR(config-if)#
```

```
ip address 172.16.101.1 255.255.255.252
```

```
Mcast_SR(config-if)#
```

```
ip pim sparse-mode
```

```
Mcast_SR(config-if)#
```

```
ip service reflect TenGigabitEthernet0/0/2 destination 239.5.5.5 to 239.55.55.55 mask-len 32 source 172.
```


```
Mcast_SR(config-if)#
```

```
ip igmp static-group 239.5.5.5
```

```
Mcast_SR(config-if)#
```

```
end
```

---



**Note:** Interface "Vif1" is strictly associated with the default Virtual Route Forwarding (VRF) for the Multicast Service Reflection feature, if a different number for this interface is selected a message can be displayed into the console similar to "Only Vif1 can be used for Service Reflect in global table".

---

The next output shows the Vif interface after the configuration has been applied, notice that the interface operation becomes UP/UP and the output feature is SERVICE REFLECT with the help of the **show ip interface <interface>** command:

```
<#root>
```

```
Mcast_SR#
```

```
show run interface Vif1
```



Building configuration...

Current configuration : 229 bytes

```
!  
interface Vif1  
 ip address 172.16.101.1 255.255.255.252  
 ip pim sparse-mode  
 ip service reflect TenGigabitEthernet0/0/2 destination 239.5.5.5 to 239.55.55.55 mask-len 32 source 172.16.101.1  
 ip igmp static-group 239.5.5.5  
end
```

Mcast\_SR#

**show ip interface vif1**

**Vif1 is up, line protocol is up**

```
Internet address is 172.16.101.1/30  
Broadcast address is 255.255.255.255  
Address determined by non-volatile memory  
MTU is 1514 bytes  
Helper address is not set  
Directed broadcast forwarding is disabled  
Multicast reserved groups joined: 224.0.0.1 224.0.0.2 224.0.0.22 224.0.0.13  
Outgoing Common access list is not set  
Outgoing access list is not set  
Inbound Common access list is not set  
Inbound access list is not set  
Proxy ARP is enabled  
Local Proxy ARP is disabled  
Security level is default  
Split horizon is enabled  
ICMP redirects are always sent  
ICMP unreachable are always sent  
ICMP mask replies are never sent  
IP fast switching is enabled  
IP Flow switching is disabled  
IP CEF switching is enabled  
IP CEF switching turbo vector  
IP Null turbo vector  
Associated unicast routing topologies:  
  Topology "base", operation state is UP  
Associated multicast routing topologies:  
  Topology "base", operation state is UP  
IP multicast fast switching is enabled  
IP multicast distributed fast switching is disabled  
IP route-cache flags are Fast, CEF  
Router Discovery is disabled  
IP output packet accounting is disabled  
IP access violation accounting is disabled  
TCP/IP header compression is disabled  
RTP/IP header compression is disabled  
Probe proxy name replies are disabled  
Policy routing is disabled  
Network address translation is disabled  
BGP Policy Mapping is disabled  
Input features: MCI Check
```

**Output features: SERVICE REFLECT**

```
IPv4 WCCP Redirect outbound is disabled  
IPv4 WCCP Redirect inbound is disabled
```

```
IPv4 WCCP Redirect exclude is disabled
IP Clear Dont Fragment is disabled
```

## LHR (ASR1000)

For the configuration at the LHR, in the same way, enable multicast routing and PIM-SM on the interfaces connected to Multicast Service Reflector router and towards the Receiver and also make sure to configure the RP IPv4 address:

```
<#root>
```

```
ASR1000#
```

```
show run | i multicast|rp-address
```

```
ip multicast-routing distributed
ip pim rp-address 10.100.100.1
```

```
ASR1000#
```

```
show run interface Ten 0/1/1
```

```
Building configuration...
```

```
Current configuration : 133 bytes
```

```
!
interface TenGigabitEthernet0/1/1
 ip address 172.16.0.2 255.255.255.252
```

```
ip pim sparse-mode
```

```
no negotiation auto
 cdp enable
end
```

```
ASR1000#
```

```
show run interface GigabitEthernet0/0/0
```

```
Building configuration...
```

```
Current configuration : 127 bytes
```

```
!
interface GigabitEthernet0/0/0
 ip address 172.16.0.5 255.255.255.252
```

```
ip pim sparse-mode
```

```
negotiation auto
 cdp enable
end
```

For the IP connectivity, this router is also using BGP to exchange the connected prefixes:

```
<#root>
```

```
ASR1000#
```

```
show run | section router bgp
```

```
router bgp 65002  
  bgp log-neighbor-changes
```

```
redistribute connected
```

```
  neighbor 172.16.0.1 remote-as 65001
```

## Initial Verification

As an initial verification, you can start checking the basic IP connectivity, route exchange between devices and PIM status. The next output was taken from the Catalyst 9K and shows the routes installed towards the Receiver (172.16.0.6), RP Loopback0 (10.100.100.1), and Vif (172.16.101.1). It can also be seen that there is a PIM adjacency with the C8000 router and that the RP mapping is also pointing to the 10.100.100.1 address.

```
<#root>
```

```
C9000#
```

```
show ip route 172.16.0.6
```

```
Routing entry for 172.16.0.4/30  
  Known via "bgp 65003", distance 20, metric 0  
  Tag 65001, type external  
  Last update from 192.168.0.6 2d22h ago  
  Routing Descriptor Blocks:  
  * 192.168.0.6, from 192.168.0.6, 2d22h ago  
    opaque_ptr 0x7FD797D4BA18  
    Route metric is 0, traffic share count is 1  
    AS Hops 2  
    Route tag 65001  
    MPLS label: none
```

```
C9000#
```

```
ping 172.16.0.6
```

```
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 172.16.0.6, timeout is 2 seconds:  
!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
C9000#
```

```
show ip route 10.100.100.1
```

```
Routing entry for 10.100.100.1/32  
  Known via "bgp 65003", distance 20, metric 0  
  Tag 65001, type external  
  Last update from 192.168.0.6 2d22h ago  
  Routing Descriptor Blocks:  
  * 192.168.0.6, from 192.168.0.6, 2d22h ago  
    opaque_ptr 0x7FD797D4B8D8  
    Route metric is 0, traffic share count is 1  
    AS Hops 1
```

```
Route tag 65001
MPLS label: none
```

```
C9000#
```

```
ping 10.100.100.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.100.100.1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
C9000#
```

```
show ip route 172.16.101.1
```

```
Routing entry for 172.16.101.0/30
```

```
Known via "bgp 65003", distance 20, metric 0
```

```
Tag 65001, type external
```

```
Last update from 192.168.0.6 2d00h ago
```

```
Routing Descriptor Blocks:
```

```
* 192.168.0.6, from 192.168.0.6, 2d00h ago
```

```
opaque_ptr 0x7FD797D4B8D8
```

```
Route metric is 0, traffic share count is 1
```

```
AS Hops 1
```

```
Route tag 65001
```

```
MPLS label: none
```

```
C9000#
```

```
ping 172.16.101.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 172.16.101.1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
C9000#
```

```
show ip pim 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 |
|------------------|-------------------------|----------------|-----|--------------|
| 192.168.0.6      | TenGigabitEthernet1/0/3 | 2d22h/00:01:29 | v2  | 1 / DR S P G |

```
C9000#
```

```
show ip pim interface
```

| Address     | Interface               | Ver/ Mode | Nbr Count | Query Intvl | DR Prior | DR          |
|-------------|-------------------------|-----------|-----------|-------------|----------|-------------|
| 192.168.0.5 | TenGigabitEthernet1/0/3 | v2/S      | 1         | 30          | 1        | 192.168.0.6 |
| 192.168.0.1 | TenGigabitEthernet1/0/6 | v2/S      | 0         | 30          | 1        | 192.168.0.1 |

```
C9000#
```

```
show ip pim rp mapping
```

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

```
Group(s): 224.0.0.0/4, Static
```

```
RP: 10.100.100.1
```

```
(?)
```

The Multicast Service Reflector shows both addresses 172.16.101.1 (Vif) and 10.100.100.1 (Loopback0) as directly connected, and routes to Source and Receiver being learned through BGP:

```
<#root>
```

```
Mcast_SR#
```

```
show ip route 10.100.100.1
```

```
Routing entry for 10.100.100.1/32
```

```
Known via "connected", distance 0, metric 0 (connected, via interface)
```

```
Advertised by bgp 65001
```

```
Routing Descriptor Blocks:
```

```
* directly connected, via Loopback0
```

```
Route metric is 0, traffic share count is 1
```

```
Mcast_SR#
```

```
show ip route 172.16.101.1
```

```
Routing entry for 172.16.101.1/32
```

```
Known via "connected", distance 0, metric 0 (connected)
```

```
Routing Descriptor Blocks:
```

```
* directly connected, via Vif1
```

```
Route metric is 0, traffic share count is 1
```

```
Mcast_SR#
```

```
show ip route 172.16.0.6
```

```
Routing entry for 172.16.0.4/30
```

```
Known via "bgp 65001", distance 20, metric 0
```

```
Tag 65002, type external
```

```
Last update from 172.16.0.2 5d20h ago
```

```
Routing Descriptor Blocks:
```

```
* 172.16.0.2, from 172.16.0.2, 5d20h ago
```

```
opaque_ptr 0x7F841CBD53D0
```

```
Route metric is 0, traffic share count is 1
```

```
AS Hops 1
```

```
Route tag 65002
```

```
MPLS label: none
```

```
Mcast_SR#
```

```
show ip route 192.168.0.2
```

```

Routing entry for 192.168.0.0/30
  Known via "bgp 65001", distance 20, metric 0
  Tag 65003, type external
  Last update from 192.168.0.5 5d20h ago
  Routing Descriptor Blocks:
  * 192.168.0.5, from 192.168.0.5, 5d20h ago
    opaque_ptr 0x7F841CBD5290
    Route metric is 0, traffic share count is 1
    AS Hops 1
    Route tag 65003
    MPLS label: none

```

Router C8000 is also showing two PIM neighbors with four interfaces enabled with PIM-SM (including Vif and Loopback0), the correct RP mapping is also displayed and with the use of command **show ip multicast** you can verify the Multicast Service Reflect capabilities of the device:

```
<#root>
```

```
Mcast_SR#
```

```
show ip pim 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 |
|------------------|-------------------------|----------------|-----|--------------|
| 192.168.0.5      | TenGigabitEthernet0/0/2 | 5d20h/00:01:34 | v2  | 1 / S P G    |
| 172.16.0.2       | TenGigabitEthernet0/0/5 | 5d20h/00:01:15 | v2  | 1 / DR S P G |

```
Mcast_SR#
```

```
show ip pim interface
```

| Address      | Interface               | Ver/ Mode | Nbr Count | Query Intvl | DR Prior | DR           |
|--------------|-------------------------|-----------|-----------|-------------|----------|--------------|
| 10.100.100.1 | Loopback0               | v2/S      | 0         | 30          | 1        | 10.100.100.1 |
| 192.168.0.6  | TenGigabitEthernet0/0/2 | v2/S      | 1         | 30          | 1        | 192.168.0.6  |
| 172.16.0.1   | TenGigabitEthernet0/0/5 | v2/S      | 1         | 30          | 1        | 172.16.0.2   |
| 172.16.101.1 | Vif1                    | v2/S      | 0         | 30          | 1        | 0.0.0.0      |

```
Mcast_SR#
```

```
show ip pim rp mapping
```

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

```
Group(s): 224.0.0.0/4, Static
```

```
RP: 10.100.100.1
```

```
(?)
```

Mcast\_SR#

**show ip multicast**

```
Multicast Routing: enabled
Multicast Multipath: disabled
Multicast Route limit: No limit
Multicast Fallback group mode: Sparse
Number of multicast boundaries configured with filter-atorp option: 0
MoFRR: Disabled
```

**Multicast Service-Reflect Capabilities:**

```
    Multicast to Multicast  Unicast to Multicast  Multicast to Unicast
```

At the LHR, similar verifications are completed. The next output taken from the ASR1000 router shows the routes installed for the translated Multicast Source 172.16.101.2 (this is important to note as at this point the multicast stream packet is received with a new source and destination IPv4 header) and for the RP IPv4 address 10.100.100.1. Notice that LHR also has reachability to the original Multicast Source 192.168.0.2, however at this point in the topology is more important to have reachability to the Vif subnet (172.16.101.0/30):

<#root>

ASR1000#

**show ip route 10.100.100.1**

```
Routing entry for 10.100.100.1/32
  Known via "bgp 65002", distance 20, metric 0
  Tag 65001, type external
  Last update from 172.16.0.1 2d22h ago
  Routing Descriptor Blocks:
  * 172.16.0.1, from 172.16.0.1, 2d22h ago
    opaque_ptr 0x7F2BFA855B80
    Route metric is 0, traffic share count is 1
    AS Hops 1
    Route tag 65001
    MPLS label: none
```

ASR1000#

**ping 10.100.100.1**

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.100.100.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

ASR1000#

**show ip route 172.16.101.1**

```
Routing entry for 172.16.101.0/30
  Known via "bgp 65002", distance 20, metric 0
  Tag 65001, type external
  Last update from 172.16.0.1 2d00h ago
  Routing Descriptor Blocks:
  * 172.16.0.1, from 172.16.0.1, 2d00h ago
```

```
opaque_ptr 0x7F2BFA855B80
Route metric is 0, traffic share count is 1
AS Hops 1
Route tag 65001
MPLS label: none
```

```
ASR1000#
```

```
ping 172.16.101.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 172.16.101.1, timeout is 2 seconds:
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
ASR1000#
```

```
show ip pim 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 |
|------------------|-------------------------|----------------|-----|--------------|
| 172.16.0.1       | TenGigabitEthernet0/1/1 | 2d22h/00:01:24 | v2  | 1 / S P G    |

```
ASR1000#
```

```
show ip pim interface
```

| Address    | Interface               | Ver/ Mode | Nbr Count | Query Intvl | DR Prior | DR         |
|------------|-------------------------|-----------|-----------|-------------|----------|------------|
| 172.16.0.5 | GigabitEthernet0/0/0    | v2/S      | 0         | 30          | 1        | 172.16.0.5 |
| 172.16.0.2 | TenGigabitEthernet0/1/1 | v2/S      | 1         | 30          | 1        | 172.16.0.2 |

```
ASR1000#
```

```
show ip route 192.168.0.2
```

```
Routing entry for 192.168.0.0/30
```

```
Known via "bgp 65002", distance 20, metric 0
```

```
Tag 65001, type external
```

```
Last update from 172.16.0.1 5d21h ago
```

```
Routing Descriptor Blocks:
```

```
* 172.16.0.1, from 172.16.0.1, 5d21h ago
```

```
opaque_ptr 0x7F2BFA855E00
```

```
Route metric is 0, traffic share count is 1
```

```
AS Hops 2
```

```
Route tag 65001
```

```
MPLS label: none
```

```
ASR1000#
```

```
ping 192.168.0.2 source 172.16.0.5
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 192.168.0.2, timeout is 2 seconds:
Packet sent with a source address of 172.16.0.5
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```



# Troubleshooting

After the Multicast Source start sending the stream for group 239.5.5.5, and while enabling **debug ip pim** on the different devices across the topology, it can be confirmed that the FHR registers the Multicast Source 192.168.0.2 to the RP and adds the interface TenGigabitEthernet1/0/3 into the Outgoing Interface List (OIL).

```
<#root>
```

```
C9000#
```

```
*Oct 16 16:23:01.817: PIM(0)[default]: Re-check RP 10.100.100.1 into the (*, 239.5.5.5) entry
```

```
*Oct 16 16:23:01.817: PIM(0)[default]: Building Triggered (*,G) Join / (S,G,RP-bit) Prune message for 239.5.5.5
```

```
*Oct 16 16:23:01.817: PIM(0)[default]: Adding register encap tunnel (Tunnel0) as forwarding interface of
```

```
*Oct 16 16:23:01.839: PIM(0)[default]: Received v2 Join/Prune on TenGigabitEthernet1/0/3 from 192.168.0.6
```

```
*Oct 16 16:23:01.839: PIM(0)[default]: Join-list: (192.168.0.2/32, 239.5.5.5), S-bit set
```

```
*Oct 16 16:23:01.839: PIM(0)[default]: MIDB Add TenGigabitEthernet1/0/3/192.168.0.6 to (192.168.0.2, 239.5.5.5)
```

```
*Oct 16 16:23:01.839: PIM(0)[default]: Join to 0.0.0.0 on TenGigabitEthernet1/0/6 for (192.168.0.2, 239.5.5.5)
```

```
*Oct 16 16:23:03.869: PIM(0)[default]: Received v2 Join/Prune on TenGigabitEthernet1/0/3 from 192.168.0.6
```

```
*Oct 16 16:23:03.869: PIM(0)[default]: Join-list: (192.168.0.2/32, 239.5.5.5), S-bit set
```

```
*Oct 16 16:23:03.869: PIM(0)[default]: MIDB Update TenGigabitEthernet1/0/3/192.168.0.6 to (192.168.0.2, 239.5.5.5)
```

```
*Oct 16 16:23:05.818: PIM(0)[default]: Received v2 Register-Stop on TenGigabitEthernet1/0/3 from 10.100.100.1
```

```
*Oct 16 16:23:05.818: PIM(0)[default]: for source 192.168.0.2, group 239.5.5.5
```

```
*Oct 16 16:23:05.818: PIM(0)[default]: Removing register encap tunnel (Tunnel0) as forwarding interface of
```

```
*Oct 16 16:23:05.818: PIM(0)[default]: Clear Registering flag to 10.100.100.1 for (192.168.0.2/32, 239.5.5.5)
```

```
*Oct 16 16:23:11.997: PIM(0)[default]: Send v2 join/prune to 192.168.0.6 (TenGigabitEthernet1/0/3)
```

```
C9000#
```

```
show ip mroute 239.5.5.5
```

```
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,
```

```
* - determined by Assert, # - iif-starg configured on rpf intf,
```

```
e - encap-helper tunnel flag, l - LISP decap ref count contributor
```

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

```
t - LISP transit group
```

```
Timers: Uptime/Expires
```

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

```
(*, 239.5.5.5), 00:24:22/stopped, RP 10.100.100.1, flags: SPF
```

```
Incoming interface: TenGigabitEthernet1/0/3, RPF nbr 192.168.0.6
```

```
Outgoing interface list: Null
```

(192.168.0.2, 239.5.5.5), 00:24:22/00:03:11, flags: FT

Incoming interface: TenGigabitEthernet1/0/6, RPF nbr 0.0.0.0  
Outgoing interface list:

TenGigabitEthernet1/0/3, Forward/Sparse, 00:24:22/00:02:47, flags:

At the same time, the Multicast Service Reflector router (that is also the RP for both groups) receives the registration for the S,G (192.168.0.2, 239.5.5.5) and immediately detects that this group falls under the reflection rule previously defined. It then adds the TenGigabitEthernet0/0/2 as the incoming interface for the 239.5.5.5 stream, and later registers the new S,G (172.16.101.2, 239.55.55.55) to itself. The router adds Vif1 as the incoming interface for the 239.55.55.55 stream and places TenGigabitEthernet0/0/5 (connecting to ASR1000) into the OIL as shown in the debug and the **show ip mroute** command output:

<#root>

Mcast\_SR#

```
*Oct 16 15:46:11.758: PIM(0)[default]: Received v2 Register on TenGigabitEthernet0/0/2 from 192.168.0.5
*Oct 16 15:46:11.758:      for 192.168.0.2, group 239.5.5.5

*Oct 16 15:46:11.758: PIM(0)[default]: Adding register decap tunnel (Tunnel1) as accepting interface of
*Oct 16 15:46:11.758: MSR(0)[default]: Add 239.5.5.5 to all the SR rules
*Oct 16 15:46:11.758: MSR: Found Grp idx 553648130 for rule 239.5.5.5->239.55.55.55
*Oct 16 15:46:11.758: MSR: Found Grp idx 553648130 for rule 239.5.5.5->239.55.55.55

*Oct 16 15:46:11.758: PIM(0)[default]: Insert (192.168.0.2,239.5.5.5) join in nbr 192.168.0.5's queue
*Oct 16 15:46:11.758: PIM(0)[default]: Building Join/Prune packet for nbr 192.168.0.5
*Oct 16 15:46:11.758: PIM(0)[default]: Adding v2 (192.168.0.2/32, 239.5.5.5), S-bit Join
*Oct 16 15:46:11.758: PIM(0)[default]: Send v2 join/prune to 192.168.0.5 (TenGigabitEthernet0/0/2)
*Oct 16 15:46:12.929: PIM(0)[default]: Building Periodic (*,G) Join / (S,G,RP-bit) Prune message for 239.5.5.5
*Oct 16 15:46:12.929: PIM(0)[default]: rp our address for group 239.55.55.55
*Oct 16 15:46:13.251: PIM(0)[default]: Received v2 Join/Prune on TenGigabitEthernet0/0/5 from 172.16.101.2
*Oct 16 15:46:13.251: PIM(0)[default]: Join-list: (*, 239.55.55.55), RPT-bit set, WC-bit set, S-bit set
*Oct 16 15:46:13.252: PIM(0)[default]: MIDB Update TenGigabitEthernet0/0/5/172.16.101.2 to (*, 239.55.55.55)
*Oct 16 15:46:13.732: PIM(0)[default]: Received v2 Register on TenGigabitEthernet0/0/2 from 192.168.0.5
*Oct 16 15:46:13.732:      for 192.168.0.2, group 239.5.5.5
*Oct 16 15:46:13.732: PIM(0)[default]: Removing register decap tunnel (Tunnel1) as accepting interface of
*Oct 16 15:46:13.732: PIM(0)[default]: Installing TenGigabitEthernet0/0/2 as accepting interface for (192.168.0.2, 239.5.5.5)
*Oct 16 15:46:13.732: PIM(0)[default]: Prune to 0.0.0.0 on  for (192.168.0.2, 239.5.5.5), Ignored.

*Oct 16 15:46:13.773: PIM(0)[default]: Adding register decap tunnel (Tunnel1) as accepting interface of
*Oct 16 15:46:13.773: PIM(0)[default]: Adding register encap tunnel (Tunnel0) as forwarding interface of
*Oct 16 15:46:13.773: PIM(0)[default]: Removing register decap tunnel (Tunnel1) as accepting interface of
*Oct 16 15:46:13.773: PIM(0)[default]: Installing Vif1 as accepting interface for (172.16.101.2, 239.55.55.55)

*Oct 16 15:46:13.788: PIM(0)[default]: Insert (192.168.0.2,239.5.5.5) join in nbr 192.168.0.5's queue
*Oct 16 15:46:13.788: PIM(0)[default]: Building Join/Prune packet for nbr 192.168.0.5
*Oct 16 15:46:13.788: PIM(0)[default]: Adding v2 (192.168.0.2/32, 239.5.5.5), S-bit Join
*Oct 16 15:46:13.788: PIM(0)[default]: Send v2 join/prune to 192.168.0.5 (TenGigabitEthernet0/0/2)
*Oct 16 15:46:13.794: PIM(0)[default]: Received v2 Register on Vif1 from 10.100.100.1
*Oct 16 15:46:13.794:      for 172.16.101.2, group 239.55.55.55
*Oct 16 15:46:13.794: PIM(0)[default]: Send v2 Register-Stop to 10.100.100.1 for 172.16.101.2, group 239.55.55.55
*Oct 16 15:46:13.794: PIM(0)[default]: Received v2 Register-Stop on Loopback0 from 10.100.100.1
```

```
*Oct 16 15:46:13.794: PIM(0)[default]: for source 172.16.101.2, group 239.55.55.55
*Oct 16 15:46:13.794: PIM(0)[default]: Removing register encap tunnel (Tunnel0) as forwarding interface
*Oct 16 15:46:13.794: PIM(0)[default]: Clear Registering flag to 10.100.100.1 for (172.16.101.2/32, 239
*Oct 16 15:46:13.809: PIM(0)[default]: Received v2 Join/Prune on TenGigabitEthernet0/0/5 from 172.16.0.
*Oct 16 15:46:13.809: PIM(0)[default]: Join-list: (172.16.101.2/32, 239.55.55.55), S-bit set

*Oct 16 15:46:13.809: PIM(0)[default]: MIB Update TenGigabitEthernet0/0/5/172.16.0.2 to (172.16.101.2,
```

Mcast\_SR#

show ip mroute

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,  
\* - determined by Assert, # - iif-starg configured on rpf intf,  
e - encap-helper tunnel flag, l - LISP decap ref count contributor  
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join  
t - LISP transit group

Timers: Uptime/Expires

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

```
(* , 239.5.5.5), 00:50:36/stopped, RP 10.100.100.1, flags: SJC
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
Vif1, Forward/Sparse, 00:50:36/00:00:23, flags:
```

```
(192.168.0.2, 239.5.5.5), 00:10:08/00:02:36, flags: T
Incoming interface: TenGigabitEthernet0/0/2, RPF nbr 192.168.0.5
Outgoing interface list:
Vif1, Forward/Sparse, 00:10:08/00:01:51, flags:
```


```
(* , 239.55.55.55), 00:50:21/00:03:11, RP 10.100.100.1, flags: SF
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
TenGigabitEthernet0/0/5, Forward/Sparse, 00:50:21/00:03:11, flags:
```

```
(172.16.101.2, 239.55.55.55), 00:10:06/00:03:23, flags: FT
Incoming interface: Vif1, RPF nbr 0.0.0.0
Outgoing interface list:
TenGigabitEthernet0/0/5, Forward/Sparse, 00:10:06/00:03:16, flags:
```

```
(* , 224.0.1.40), 00:50:36/00:03:28, RP 10.100.100.1, flags: SJCL
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
TenGigabitEthernet0/0/2, Forward/Sparse, 00:50:13/00:03:28, flags:
TenGigabitEthernet0/0/5, Forward/Sparse, 00:50:28/00:03:10, flags:
```

Loopback0, Forward/Sparse, 00:50:36/00:02:17, flags:

---



**Note:** Notice that to be able to observe the Multicast Service Reflection rules being used in the debug output, the **debug ip multicast service-reflect** needs to be enabled. Be aware that this particular debug was introduced in Cisco IOS-XE release 17.9.1a, for additional information refer to [Feature Information for Multicast Service Reflection](#) section from the IP Multicast Configuration Guide, Cisco IOS XE 17.x

---

From the previous **show ip mroute** output it can be seen that there are \*,G and S,G entries for both the original and new multicast groups. Additionally, the use of the **show ip mroute <group> count** command can give an idea of the packets received and forwarded for each multicast stream, this output is also useful as it helps to identify the traffic forwarding by the confirmation of the incrementing counters:

```
<#root>
```

```
Mcast_SR#
```

```
show ip mroute 239.5.5.5 count
```

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

IP Multicast Statistics

5 routes using 5796 bytes of memory  
3 groups, 0.66 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: 239.5.5.5, Source count: 1, Packets forwarded: 4053, Packets received: 4056

RP-tree: Forwarding: 11/0/100/0, Other: 12/1/0  
Source: 192.168.0.2/32, Forwarding: 4042/0/113/0, Other: 4044/2/0

Mcast\_SR#

show ip mroute 239.5.5.5 count

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

IP Multicast Statistics

5 routes using 5796 bytes of memory  
3 groups, 0.66 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: 239.5.5.5, Source count: 1, Packets forwarded: 4058, Packets received: 4061

RP-tree: Forwarding: 11/0/100/0, Other: 12/1/0  
Source: 192.168.0.2/32, Forwarding: 4047/0/113/0, Other: 4049/2/0

Mcast\_SR#

show ip mroute 239.55.55.55 count

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

IP Multicast Statistics

5 routes using 5796 bytes of memory  
3 groups, 0.66 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: 239.55.55.55, Source count: 1, Packets forwarded: 4046, Packets received: 4046

RP-tree: Forwarding: 0/0/0/0, Other: 0/0/0  
Source: 172.16.101.2/32, Forwarding: 4046/0/114/0, Other: 4046/0/0

Mcast\_SR#

show ip mroute 239.55.55.55 count

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

IP Multicast Statistics

5 routes using 5796 bytes of memory  
3 groups, 0.66 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: 239.55.55.55, Source count: 1, Packets forwarded: 4051, Packets received: 4051

RP-tree: Forwarding: 0/0/0/0, Other: 0/0/0  
Source: 172.16.101.2/32, Forwarding: 4051/0/114/0, Other: 4051/0/0

If the S,G entries are not being created, it is important to determine if the multicast packets are being received at expected interface. A couple of options to confirm this situation is with use of an Embedded Packet Capture (EPC) or with a Packet Trace, in the next output an EPC was taken in the incoming (Te0/0/2) and outgoing (Te0/0/5) interfaces from the C8000 router. In this working scenario you can observe the original multicast stream 239.5.5.5 incoming and the new 239.55.55.55 stream egressing the router:

<#root>

Mcast\_SR#

```
monitor capture msr buffer size 10 match any interface tenGigabitEthernet 0/0/2 in
```

Mcast\_SR#

```
monitor capture msr start
```

```
Started capture point : msr  
*Oct 16 17:59:06.986: %BUFCAP-6-ENABLE: Capture Point msr enabled.
```

Mcast\_SR#

```
monitor capture msr stop
```

```
Stopped capture point : msr
```

Mcast\_SR#s

```
*Oct 16 17:59:25.699: %BUFCAP-6-DISABLE: Capture Point msr disabled
```

Mcast\_SR#

```
show monitor capture msr buffer brief
```

```
-----  
#   size  timestamp      source                destination          dscp  protocol  
-----  
0  114    0.000000    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
  
1  114    2.000000    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
  
2   68    2.979961    192.168.0.5          -> 224.0.0.13          48 CS6  PIM  
3  114    4.000000    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
4  114    6.000000    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
5  114    8.000000    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
6  114   10.000000    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
7  114   12.000000    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
8  114   14.001007    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
9  114   16.001999    192.168.0.2          -> 239.5.5.5           0 BE   ICMP  
10 114   18.001007    192.168.0.2          -> 239.5.5.5           0 BE   ICMP
```

Mcast\_SR#

```
monitor capture msr buffer size 10 match any interface tenGigabitEthernet0/0/5 out
```

Mcast\_SR#

```
monitor capture msr start
```

```
Started capture point : msr
```

```
*Oct 16 18:07:26.846: %BUFCAP-6-ENABLE: Capture Point msr enabled.
```

```
Mcast_SR#
```

```
monitor capture msr stop
```

```
Stopped capture point : msr
```

```
*Oct 16 18:07:50.360: %BUFCAP-6-DISABLE: Capture Point msr disabled.
```

```
Mcast_SR#
```

```
show monitor capture msr buffer brief
```

| #  | size | timestamp | source            | destination          | dscp   | protocol |
|----|------|-----------|-------------------|----------------------|--------|----------|
| 0  | 73   | 0.000000  | 172.16.0.1        | -> 172.16.0.2        | 48 CS6 | TCP      |
| 1  | 114  | 0.172040  | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 2  | 114  | 2.173047  | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 3  | 114  | 4.172040  | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 4  | 114  | 6.173047  | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 5  | 114  | 8.172040  | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 6  | 114  | 10.172040 | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 7  | 114  | 12.173047 | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 8  | 114  | 14.175046 | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 9  | 114  | 16.176038 | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 10 | 114  | 18.176038 | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 11 | 46   | 19.318029 | 172.16.0.1        | -> 224.0.0.1         | 48 CS6 | IGMP     |
| 12 | 114  | 20.177045 | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 13 | 72   | 20.629017 | 172.16.0.1        | -> 224.0.0.13        | 48 CS6 | PIM      |
| 14 | 114  | 22.178037 | 172.16.101.2      | -> 239.55.55.55      | 0 BE   | ICMP     |
| 15 | 430  | 23.273033 | 90:77:EE:3C:80:05 | -> 01:00:0C:CC:CC:CC | --     | LLC      |

A Packet Trace with an extended Access Control List (ACL) filtering on the 239.5.5.5 was also implemented. From the next output it can be observed that the packets are being received at the interface Te0/0/2 and initially consumed and copied internally to later be used as input packets into Vif1 interface and then forwarded out translated throughout interface Te0/0/5:

```
<#root>
```

```
Mcast_SR#
```

```
configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Mcast_SR(config)#
```

```
ip access-list extended original-mcast
```

```
Mcast_SR(config-ext-nacl)#
```

```
permit ip any host 239.5.5.5
```

```
Mcast_SR(config-ext-nacl)#
```

```
end
```

Mcast\_SR#

\*Oct 16 18:31:49.187: %SYS-5-CONFIG\_I: Configured from console by console

Mcast\_SR#

debug platform packet-trace packet 16

Please remember to turn on 'debug platform condition start' for packet-trace to work  
Mcast\_SR#

debug platform condition ipv4 access-list original-mcast ingress

Mcast\_SR#

debug platform condition start

Mcast\_SR#

show platform packet-trace packet all

Packet: 0                    CBUG ID: 0

Summary

Input        : TenGigabitEthernet0/0/2  
Output       : <none>  
State        : CONS Packet Consumed Silently  
Timestamp

Start        : 521286121898055 ns (10/16/2023 18:33:39.36879 UTC)  
Stop         : 521286122522268 ns (10/16/2023 18:33:39.37503 UTC)

Path Trace

Feature: IPV4(Input)

Input        : TenGigabitEthernet0/0/2  
Output       : <unknown>  
Source       : 192.168.0.2  
Destination : 239.5.5.5  
Protocol     : 1 (ICMP)

Packet: 1                    CBUG ID: 0

Summary

Input        : TenGigabitEthernet0/0/2

Output       : <none>

State        : CONS Packet Consumed Silently

Timestamp

Start        : 521286121980895 ns (10/16/2023 18:33:39.36962 UTC)  
Stop         : 521286122618968 ns (10/16/2023 18:33:39.37600 UTC)

Path Trace

Feature: PACKET\_COPY

Original packet number: 0

Feature: IPV4(Input)

Input        : TenGigabitEthernet0/0/2  
Output       : <unknown>

Source       : 192.168.0.2  
Destination : 239.5.5.5

Protocol     : 1 (ICMP)

Packet: 2                    CBUG ID: 0

Summary



Input : Vif1  
Output : TenGigabitEthernet0/0/5

State : FWD

Timestamp

Start : 521286122038938 ns (10/16/2023 18:33:39.37020 UTC)

Stop : 521286122694795 ns (10/16/2023 18:33:39.37676 UTC)

Path Trace

Feature: PACKET\_COPY

Original packet number: 1

Feature: IPV4(Input)

Input : Vif1

Output : <unknown>

Source : 172.16.101.2

Destination : 239.55.55.55

Protocol : 1 (ICMP)

It is relevant to also check the Multicast Routing Information Base (MRIB) and Multicast Forwarding Information Base (MFIB) at the Catalyst 8K. Notice the hardware installed (HW) flags as well as the hardware forwarding counters incrementing with ideally very few Reverse Path Forwarding (RPF) failures or drops seen. Also note the forward (F) flag associated to the egress interface (Te0/0/5) towards ASR1000 router:

<#root>

Mcast\_SR#

show ip mrib route

IP Multicast Routing Information Base

Entry flags: L - Domain-Local Source, E - External Source to the Domain,  
C - Directly-Connected Check, S - Signal, IA - Inherit Accept, D - Drop  
ET - Data Rate Exceeds Threshold, K - Keepalive, DDE - Data Driven Event  
ME - MoFRR ECMP Flow based, MNE - MoFRR Non-ECMP Flow based,  
MP - Primary MoFRR Non-ECMP Flow based entry,  
e - Encap helper tunnel flag

Interface flags: F - Forward, A - Accept, IC - Internal Copy,  
NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,  
II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,  
LD - Local Disinterest, MD - mCAC Denied, MI - mLDP Interest  
A2 - MoFRR ECMP Backup Accept

(\* ,224.0.0.0/4) Flags: C

(\* ,224.0.1.40) RPF nbr: 0.0.0.0 Flags: C

TenGigabitEthernet0/0/2 Flags: F NS

TenGigabitEthernet0/0/5 Flags: F NS

Loopback0 Flags: F IC NS

Tunnel1 Flags: A

(\* ,239.5.5.5) RPF nbr: 0.0.0.0 Flags: C

Vif1 Flags: F NS SR(172.16.101.2, 239.55.55.55) Index(2, 2)

Tunnel Flags: A

(\* ,239.55.55.55) RPF nbr: 0.0.0.0 Flags: C  
TenGigabitEthernet0/0/5 Flags: F NS  
Tunnel Flags: A

(172.16.101.2,239.55.55.55) RPF nbr: 0.0.0.0 Flags:  
Vif1 Flags: A  
TenGigabitEthernet0/0/5 Flags: F NS

(192.168.0.2,239.5.5.5) RPF nbr: 192.168.0.5 Flags:  
TenGigabitEthernet0/0/2 Flags: A NS  
Vif1 Flags: F NS SR(172.16.101.2, 239.55.55.55) Index(2, 2)

Mcast\_SR#

show ip mfib 192.168.0.2 239.5.5.5

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,  
e - Encap helper tunnel flag.

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: HW Pkt Count/FS Pkt Count/PS Pkt Count Egress Rate in pps  
Default

(192.168.0.2,239.5.5.5) Flags: HW

SW Forwarding: 1/0/100/0, Other: 0/0/0

HW Forwarding: 6530/0/114/0, Other: 2/2/0

TenGigabitEthernet0/0/2 Flags: A

Vif1, SR/(172.16.101.2,239.55.55.55):553648130 Flags: F NS

Pkts: 0/0/1 Rate: 0 pps

Mcast\_SR#

show ip mfib 172.16.101.2 239.55.55.55

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,  
e - Encap helper tunnel flag.  
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: HW Pkt Count/FS Pkt Count/PS Pkt Count Egress Rate in pps  
Default

(172.16.101.2,239.55.55.55) Flags: HW

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

HW Forwarding: 6550/0/114/0, Other: 0/0/0

Vif1 Flags: A

TenGigabitEthernet0/0/5 Flags: F NS

Pkts: 0/0/0 Rate: 0 pps

Mcast\_SR#

show adjacency vif1 detail

| Protocol | Interface | Address   |
|----------|-----------|---|
| IP       | Vif1      | point2point(8)<br>13270 packets, 657502 bytes<br>epoch 0<br>sourced in sev-epoch 0<br>empty encap string<br>P2P-ADJ |

Mcast\_SR#

show platform software ip rp active mfib group 239.5.5.5/32

Route flags:

S - Signal; C - Directly connected;

IA - Inherit A Flag; L - Local;

BR - Bidir route

\*, 239.5.5.5/32 --> OBJ\_INTF\_LIST (0x15)

Obj id: 0x15, Flags: C  
OM handle: 0x3480200a50

Mcast\_SR#

show platform software ip rp active mfib group 239.55.55.55/32

Route flags:

S - Signal; C - Directly connected;

IA - Inherit A Flag; L - Local;

BR - Bidir route

\*, 239.55.55.55/32 --> OBJ\_INTF\_LIST (0x36)

Obj id: 0x36, Flags: C  
OM handle: 0x34801fa988

Mcast\_SR#

show platform software mlist rp active index 0x15

Multicast List entries

OCE Flags:  
NS - Negate Signalling; IC - Internal copy;  
A - Accept; F - Forward;

| OCE        | Type          | OCE Flags | Interface |
|------------|---------------|-----------|-----------|
| 0xf8000029 | OBJ_ADJACENCY | NS, F     | Vif1      |
| 0xf8000186 | OBJ_ADJACENCY | A         | Tunnel1   |

Mcast\_SR#

show platform software mlist rp active index 0x36

Multicast List entries

OCE Flags:  
NS - Negate Signalling; IC - Internal copy;  
A - Accept; F - Forward;

| OCE        | Type          | OCE Flags | Interface               |
|------------|---------------|-----------|-------------------------|
| 0xf80000d1 | OBJ_ADJACENCY | NS, F     | TenGigabitEthernet0/0/5 |
| 0xf8000186 | OBJ_ADJACENCY | A         | Tunnel1                 |

For the RPF check, it is important to verify this across the different devices in the multicast path.

<#root>

C9000#

show ip rpf 192.168.0.2

RPF information for ? (192.168.0.2)

RPF interface: TenGigabitEthernet1/0/6

RPF neighbor: ? (192.168.0.2) -

directly connected

RPF route/mask: 192.168.0.0/30  
RPF type: multicast (connected)  
Doing distance-preferred lookups across tables

RPF topology: ipv4 multicast base

Mcast\_SR#

show ip rpf 192.168.0.2

RPF information for ? (192.168.0.2)

RPF interface: TenGigabitEthernet0/0/2

RPF neighbor: ? (192.168.0.5)

RPF route/mask: 192.168.0.0/30

RPF type: unicast (bgp 65001)

Doing distance-preferred lookups across tables

RPF topology: ipv4 multicast base, originated from ipv4 unicast base

Mcast\_SR#

show ip rpf 172.16.101.2

RPF information for ? (172.16.101.2)

RPF interface: Vif1

RPF neighbor: ? (172.16.101.2) -

directly connected

RPF route/mask: 172.16.101.0/30

RPF type: multicast (connected)

Doing distance-preferred lookups across tables

RPF topology: ipv4 multicast base

ASR1000#

show ip rpf 172.16.101.2

RPF information for ? (172.16.101.2)

RPF interface: TenGigabitEthernet0/1/1

RPF neighbor: ? (172.16.0.1)

RPF route/mask: 172.16.101.0/30

RPF type: unicast (bgp 65002)

Doing distance-preferred lookups across tables

RPF topology: ipv4 multicast base, originated from ipv4 unicast base

Finally, you can confirm the traffic being received and forwarded at the LHR, the first capture shows how ASR1000 received the packets from the new stream (172.16.101.2, 239.55.55.55) where the IPv4 addresses were translated as expected. The second capture shows the traffic being forwarded to the Receiver.

<#root>

ASR1000#

monitor capture LHR buffer size 10 match any interface TenGigabitEthernet 0/1/1 in

ASR1000#

monitor capture LHR start

Started capture point : LHR

\*Oct 17 15:16:06.192: %BUFCAP-6-ENABLE: Capture Point LHR enabled.show moni

ASR1000#

show monitor capture LHR buffer brief

```
-----  
#   size  timestamp      source           destination      dscp  protocol  
-----  
0   73    0.000000    172.16.0.1      -> 172.16.0.2      48 CS6  TCP  
  
1  114    0.050992    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
2  114    2.051984    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
  
3   54    2.896993    172.16.0.1      -> 172.16.0.2      48 CS6  TCP  
4  114    4.051984    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
5   72    4.967989    172.16.0.1      -> 224.0.0.13      48 CS6  PIM  
6  114    6.052991    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
7  114    8.053983    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
8  114   10.053983    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
9  114   12.053983    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
10 114   14.055997    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
11 114   16.056989    172.16.101.2    -> 239.55.55.55    0 BE   ICMP  
12 114   18.055997    172.16.101.2    -> 239.55.55.55    0 BE   ICMP
```

ASR1000#

monitor capture LHR stop

Stopped capture point : LHR

ASR1000#

\*Oct 17 15:16:32.029: %BUFCAP-6-DISABLE: Capture Point LHR disabled.

ASR1000#

monitor capture TO\_RECEIVER buffer size 10 match any interface GigabitEthernet 0/0/0 out

ASR1000#

monitor capture TO\_RECEIVER start

Started capture point : TO\_RECEIVER

\*Oct 17 15:19:00.309: %BUFCAP-6-ENABLE: Capture Point TO\_RECEIVER enabled.ni

ASR1000#

show monitor capture TO\_RECEIVER buffer brief

```
-----  
#   size  timestamp      source           destination      dscp  protocol
```

```

-----
 0 114 0.000000 172.16.101.2 -> 239.55.55.55 0 BE ICMP
 1 46 0.276994 172.16.0.5 -> 224.0.0.1 48 CS6 IGMP

 2 114 1.999009 172.16.101.2 -> 239.55.55.55 0 BE ICMP
 3 114 4.000000 172.16.101.2 -> 239.55.55.55 0 BE ICMP

 4 46 5.550027 172.16.0.5 -> 224.0.1.40 48 CS6 IGMP
 5 114 5.999009 172.16.101.2 -> 239.55.55.55 0 BE ICMP
 6 114 8.000000 172.16.101.2 -> 239.55.55.55 0 BE ICMP
 7 114 10.001007 172.16.101.2 -> 239.55.55.55 0 BE ICMP
 8 114 12.001999 172.16.101.2 -> 239.55.55.55 0 BE ICMP
 9 114 14.003006 172.16.101.2 -> 239.55.55.55 0 BE ICMP
10 114 16.003998 172.16.101.2 -> 239.55.55.55 0 BE ICMP

```

ASR1000#

monitor capture TO\_RECEIVER stop

Stopped capture point : TO\_RECEIVER

\*Oct 17 15:19:24.938: %BUFCAP-6-DISABLE: Capture Point TO\_RECEIVER disabled.

If for some reason the \*,G is missing at some point, you can validate the Internet Group Management Protocol (IGMP) status with the commands **show ip igmp membership <group>** and **show ip igmp groups <group>**. In this scenario, it is expected to see an entry for the translated multicast group 239.55.55.55.

<#root>

ASR1000#

show ip igmp membership 239.55.55.55

Flags: A - aggregate, T - tracked

L - Local, S - static, V - virtual, R - Reported through v3

I - v3lite, U - Urd, M - SSM (S,G) channel

1,2,3 - The version of IGMP, the group is in

Channel/Group-Flags:

/ - Filtering entry (Exclude mode (S,G), Include mode (G))

Reporter:

<mac-or-ip-address> - last reporter if group is not explicitly tracked

<n>/<m> - <n> reporter in include mode, <m> reporter in exclude

| Channel/Group  | Reporter   | Uptime | Exp.  | Flags | Interface |
|----------------|------------|--------|-------|-------|-----------|
| *,239.55.55.55 | 172.16.0.6 | 1w0d   | 02:40 | 2A    | Gi0/0/0   |

ASR1000#

show ip igmp groups 239.55.55.55

IGMP Connected Group Membership

| Group Address | Interface            | Uptime   | Expires  | Last Reporter | Group Accounted |
|---------------|----------------------|----------|----------|---------------|-----------------|
| 239.55.55.55  | GigabitEthernet0/0/0 | 01:11:56 | 00:02:49 | 172.16.0.6    |                 |

In addition, at the LHR a verification of the mroute can be used to confirm that the interface towards the Receiver is added into the OIL and the use of an mtrace to confirm the reverse path to both sources.

```
<#root>
```

```
ASR1000#
```

```
show ip mroute 239.55.55.55
```

```
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,  
* - determined by Assert, # - iif-starg configured on rpf intf,  
e - encap-helper tunnel flag, l - LISP Decap Refcnt Contributor
```

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

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

```
(* , 239.55.55.55), 23:28:12/stopped, RP 10.100.100.1, flags: SJC  
Incoming interface: TenGigabitEthernet0/1/1, RPF nbr 172.16.0.1  
Outgoing interface list:  
GigabitEthernet0/0/0, Forward/Sparse, 00:01:52/00:02:48
```

```
(172.16.101.2, 239.55.55.55), 00:10:11/00:02:40, flags: JT
```

```
Incoming interface: TenGigabitEthernet0/1/1, RPF nbr 172.16.0.1  
Outgoing interface list:
```

```
GigabitEthernet0/0/0, Forward/Sparse, 00:01:52/00:02:48
```

```
ASR1000#
```

```
mtrace 172.16.101.2
```

```
Type escape sequence to abort.
```

```
Mtrace from 172.16.101.2 to 172.16.0.2 via RPF
```

```
From source (?) to destination (?)
```

```
Querying full reverse path...
```

```
0 172.16.0.2  
-1 172.16.0.2 ==> 172.16.0.2 PIM/MBGP [172.16.101.0/30]  
-2 172.16.0.1 ==> 172.16.101.1 PIM_MT [172.16.101.0/30]  
-3 172.16.101.2
```

```
ASR1000#
```

```
mtrace 192.168.0.2
```

```
Type escape sequence to abort.
```

```
Mtrace from 192.168.0.2 to 172.16.0.2 via RPF
```



From source (?) to destination (?)

Querying full reverse path...

```
0 172.16.0.2
-1 172.16.0.2 ==> 172.16.0.2 PIM/MBGP [192.168.0.0/30]
-2 172.16.0.1 ==> 192.168.0.6 PIM/MBGP [192.168.0.0/30]
-3 192.168.0.5 ==> 192.168.0.1 PIM_MT [192.168.0.0/30]
-4 192.168.0.2
```

## Related Information

- [Implementing Multicast Service Reflection](#)
- [Multicast Service Reflection using PIM-SM on IOS-XE : Multicast to Unicast](#)