

Multicast Observation Lab – Multiple Routers (PIM-Dense Mode)

Objective

The objective is to observe how **interconnected multicast routers** use **PIM Dense Mode** to distribute multicast traffic. The focus should be on the changes in the multicast routing table. This is an observation lab. No troubleshooting is required.

Topology Overview

RTR1, RTR2, and RTR3 are multicast-enabled routers. Source and Host A are routers with IP routing disabled. They will serve as PCs. Host A is configured to join group 239.1.1.1.




Figure 1 – Topology

Task 1 – View Group Membership

- a. On RTR1 and RTR2 view the IGMP group membership information and multicast routing table.

```
RTR# show ip igmp groups
```

```
RTR# show ip mroute
```

Observe: In the IGMP group table, RTR2 has an entry for group 239.1.1.1 associated with interface Gi0/0. Gi0/0 is connected to Host A, the interested receiver. In the multicast routing table, RTR2 has created entry (*, 239.1.1.1) and populated the OIL with interface Gi0/0. RTR2 is aware of a interested receiver for group 239.1.1.1.

RTR1's IGMP group table and multicast routing table contain no entries for group 239.1.1.1. RTR1 is unaware of any interested receivers.

```
RTR2#sh ip igmp groups
IGMP Connected Group Membership
Group Address    Interface          Uptime    Expires    Last Reporter  Group Accounted
239.1.1.1        GigabitEthernet0/0 00:29:09  00:02:59   172.16.2.2
224.0.1.40       GigabitEthernet0/0 00:29:10  00:02:57   172.16.2.1
RTR2#
```

Figure 2 – RTR2 show ip igmp groups output

```
RTR1#sh ip igmp groups
IGMP Connected Group Membership
Group Address    Interface          Uptime    Expires    Last Reporter  Group Accounted
224.0.1.40       GigabitEthernet0/3 00:29:54  00:02:08   44.44.44.1
RTR1#
```

Figure 3 – RTR1 show ip igmp groups output

```
(*, 239.1.1.1), 00:29:32/00:02:36, RP 0.0.0.0, flags: DC
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
GigabitEthernet0/0, Forward/Dense, 00:29:32/stopped
```

Figure 4 – RTR2 show ip mroute output

- b. On Source, begin sending multicast traffic to group 239.1.1.1. On RTR1 and RTR2 view mroute tables.

```
Source# ping 239.1.1.1 re 2
```

```
RTR# show ip mroute
```

Observe: Host A does not respond to the ping. RTR1 has an (S,G) multicast routing entry (44.44.44.44, 239.1.1.1) with an incoming interface of Gi0/3. RTR1 is aware of the multicast stream from source 44.44.44.44 to group 239.1.1.1.

RTR1 also has a (*, 239.1.1.1) entry; however, no interfaces are recorded in the OIL. With no interfaces in the OIL, there is nowhere to forward multicast traffic received for group 239.1.1.1.

RTR2 has no (44.44.44.44, 239.1.1.1) multicast routing entry, indicating that it is currently unaware of the multicast stream from source 44.44.44.44 to group 239.1.1.1.

```
(*, 239.1.1.1), 00:00:06/stopped, RP 0.0.0.0, flags: DP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(44.44.44.44, 239.1.1.1), 00:00:06/00:02:53, flags: PT
  Incoming interface: GigabitEthernet0/3, RPF nbr 0.0.0.0
  Outgoing interface list: Null
```

Figure 5 – RTR1 show ip mroute output

Engineer Insight: In the previous labs, a single router was responsible for forwarding multicast traffic from the source to the receiver. The source and receiver were directly connected to the same router. Recall that a receiver uses the Internet Group Management Protocol (IGMP) to inform its default gateway that it wants to join, or subscribe to, a multicast group.

The dynamics change in this lab with the introduction of interconnected multicast routers. The source and receiver are now connected to different multicast routers. The problem becomes how to efficiently forward multicast traffic to several receivers on different networks that are represented by one multicast group. For multicast traffic to be forwarded, a router must know where interested receivers are located. IGMP is used only

for communication between hosts and their directly connected multicast router. This explains why RTR2 has an IGMP entry for the multicast group 239.1.1.1 while RTR1 does not. RTR1 has no directly connected receivers using IGMP to express interest in group 239.1.1.1.

A separate protocol is required for multicast routers to exchange the information needed to establish **multicast forwarding state** between a router connected to the multicast source, called the **First Hop Router (FHR)**, and a router connected to the receivers, called the **Last Hop Router (LHR)**.

Protocol Independent Multicast (PIM) is the multicast routing protocol used between multicast routers to establish and maintain multicast forwarding state. Forwarding state provides the control-plane information, or instructions, needed to forward multicast traffic downstream from the FHR to one or more LHRs and ultimately to all interested receivers. This forwarding state is represented by multicast routing entries, such as (S,G). For the sake of building a mental model, think of an (S,G) entry as representing the existence of a multicast stream and the interfaces associated with forwarding that stream. Think of a (*,G) entry as representing receiver interest in a multicast group along with the associated interfaces.

PIM operates in two primary modes: **Dense Mode (PIM-DM)** and **Sparse Mode (PIM-SM)**. This lab focuses on PIM-DM. PIM-SM is outside the scope of this lab and will be covered in a future lab.

Task 2 – Configure PIM-DM

- a. On RTR1, view the IGMP information for Gi0/1.

```
RTR1# show ip igmp interface gi0/1
```

Observe: Multicast and IGMP are disabled on the interface.

```
RTR1#sh ip igmp interface gi0/1
GigabitEthernet0/1 is up, line protocol is up
Internet address is 172.16.1.1/24
IGMP is disabled on interface
Multicast routing is disabled on interface
Multicast TTL threshold is 0
No multicast groups joined by this system
RTR1#
```

Figure 6 – RTR1 show ip igmp interface Gi0/1 output

- b. On RTR1, RTR2, and RTR3 configure PIM-DM on all inter-router interfaces. View IGMP information for RTR1's Gi0/1.

```
RTR(config-if)# ip pim dense-mode
```

Observe: Multicast and IGMP are enabled on the interface.

```
RTR1#sh ip igmp interface gi0/1
GigabitEthernet0/1 is up, line protocol is up
Internet address is 172.16.1.1/24
IGMP is enabled on interface
Current IGMP host version is 2
Current IGMP router version is 2
IGMP query interval is 60 seconds
IGMP configured query interval is 60 seconds
IGMP querier timeout is 120 seconds
IGMP configured querier timeout is 120 seconds
IGMP max query response time is 10 seconds
Last member query count is 2
Last member query response interval is 1000 ms
Inbound IGMP access group is not set
IGMP activity: 0 joins, 0 leaves
Multicast routing is enabled on interface
Multicast TTL threshold is 0
Multicast designated router (DR) is 172.16.1.1 (this system)
IGMP querying router is 172.16.1.1 (this system)
No multicast groups joined by this system
RTR1#
```

Figure 7 – RTR1 show ip igmp interface Gi0/1 output

Engineer Insight: To enable Multicast and IGMP on an interface, enable PIM-DM or PIM-SM. In the previous labs, this was pre-configured. **PIM is not enabled by default.**

- c. On Source, begin sending multicast traffic to group 239.1.1.1. On RTR1, RTR2, and RTR3 view mroute tables.

```
Source# ping 239.1.1.1 re 2
```

```
RTR# show ip mroute
```

Observe: Host A responds to the ping. Recall that RTR1 previously had a (*, 239.1.1.1) entry with a NULL OIL. The OIL now contains the interfaces connected to RTR2 and RTR3, indicating that RTR1 has forwarding interfaces for multicast traffic destined for group 239.1.1.1.

RTR2 creates forwarding state for (44.44.44.44, 239.1.1.1) and populates the OIL with interface Gi0/0. RTR3 creates forwarding state for (44.44.44.44, 239.1.1.1) and (*, 239.1.1.1).

```
(*, 239.1.1.1), 00:01:14/stopped, RP 0.0.0.0, flags: D
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
  GigabitEthernet0/2, Forward/Dense, 00:00:35/stopped
  GigabitEthernet0/1, Forward/Dense, 00:00:48/stopped
(44.44.44.44, 239.1.1.1), 00:01:14/00:01:45, flags: T
Incoming interface: GigabitEthernet0/3, RPF nbr 0.0.0.0
Outgoing interface list:
  GigabitEthernet0/2, Prune/Dense, 00:00:16/00:02:43
  GigabitEthernet0/1, Forward/Dense, 00:00:48/stopped
```

Figure 8 – RTR1 show ip mroute output

```
(*, 239.1.1.1), 00:57:24/stopped, RP 0.0.0.0, flags: DC
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
  GigabitEthernet0/1, Forward/Dense, 00:01:34/stopped
  GigabitEthernet0/0, Forward/Dense, 00:57:24/stopped
(44.44.44.44, 239.1.1.1), 00:01:02/00:01:57, flags: T
Incoming interface: GigabitEthernet0/1, RPF nbr 172.16.1.1
Outgoing interface list:
  GigabitEthernet0/0, Forward/Dense, 00:01:02/stopped
```

Figure 9 – RTR2 show ip mroute output

```
(*, 239.1.1.1), 00:01:48/stopped, RP 0.0.0.0, flags: D
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
  GigabitEthernet0/2, Forward/Dense, 00:01:48/stopped
(44.44.44.44, 239.1.1.1), 00:01:48/00:01:11, flags: PT
Incoming interface: GigabitEthernet0/2, RPF nbr 192.168.1.1
Outgoing interface list: Null
```

Figure 10 – RTR3 show ip mroute output

Questions:

- On RTR2, Gi0/1 is in the OIL of (*, 239.1.1.1) but not in the OIL for (44.44.44.44, 239.1.1.1). Why?
- Based on *Figure 10*, what forwarding action will RTR3 take?

Engineer Insight: PIM-DM uses a **flood-and-prune** model to distribute multicast traffic. When a multicast source first begins transmitting, each router that receives the multicast traffic forwards it out all PIM-DM-enabled interfaces except the incoming interface. This explains why RTR2's interface Gi0/1 is in the (*, 239.1.1.1) OIL but not the (44.44.44.44, 239.1.1.1) OIL.

In PIM-DM, RTR1 assumes that interested receivers may be reachable through all PIM-DM-enabled interfaces and creates a (*, G) entry to represent this state of potential receiver interest. RTR1's (*, 239.1.1.1) entry includes both inter-router interfaces in its OIL because both interfaces are enabled for PIM-DM. As a result, multicast traffic is initially forwarded to both RTR2 and RTR3. RTR3 receives the multicast traffic even though it has no interested receivers.

RTR3 informs RTR1 that it does not require the multicast traffic by sending a PIM Prune message. Evidence of this is visible in *Figure 8*, where Gi0/2 in the OIL of entry (44.44.44.44, 239.1.1.1) displays the **Prune/Forward** state. As a result, multicast traffic is temporarily no longer forwarded out Gi0/2.

After receiving the PIM Prune message, RTR1 places Gi0/2 in a prune state and temporarily stops forwarding multicast traffic out Gi0/2. The prune state eventually expires, requiring RTR3 to periodically send additional PIM Prune messages to renew the prune state and to continue preventing multicast traffic from being forwarded.