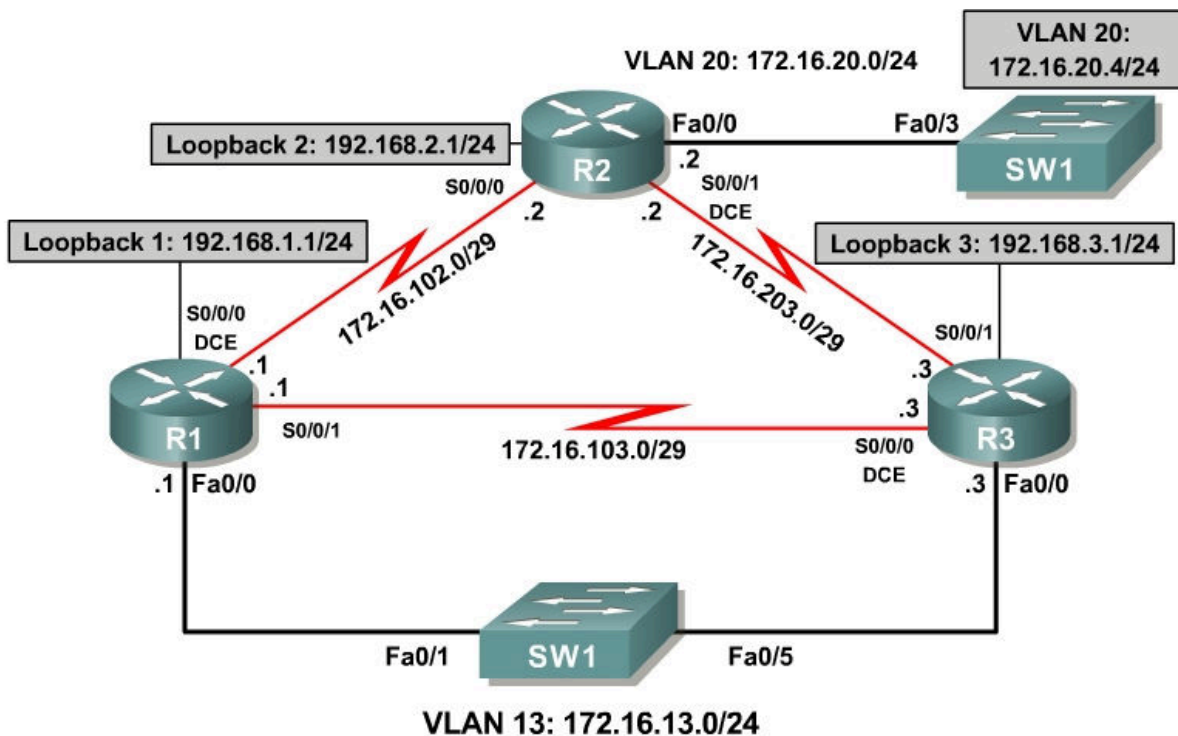# Lab 7-3 Routing IP Multicast with PIM Sparse Mode

Learning Objectives

- Implement and verify PIM-SM operation and adjacencies
- Implement and verify the use of a static rendezvous point
- Observe the shared tree and source tree
- Debug the PIM-SM operation to discover SPT shared tree to shortest-path tree transition

Topology Diagram

Scenario

Your successful assistance in configuring IP multicast routing has allowed CrushASodaCan.com to become a community action organization with a wide base of interested citizens. Because of your help, the organization can now preview their TV commercials via the network before releasing them to the public. CrushASodaCan.com also projects high growth rates in terms of volunteer workers using their network. Unfortunately, the flooding and pruning with PIM-DM that you configured for them cannot handle the new demands being made on the network. So they have hired you again as a consultant to

solve this. You decide to implement PIM Sparse Mode (PIM-SM) to create a subscription-based multicast topology in CrushASodaCan.com's network.

Step 1: Load Initial Configurations

Begin with the final configurations from Lab 7.2.

```
R1:
!
hostname R1
!
interface Loopback1
 ip address 192.168.1.1 255.255.255.0
 ip pim dense-mode
 ip igmp join-group 232.32.32.32
!
interface FastEthernet0/0
 ip address 172.16.13.1 255.255.255.0
 ip pim dense-mode
 no shutdown
!
interface Serial0/0/0
 bandwidth 64
 ip address 172.16.102.1 255.255.255.248
 ip pim dense-mode
 clock rate 64000
 no shutdown
!
interface Serial0/0/1
 bandwidth 64
 ip address 172.16.103.1 255.255.255.248
 ip pim dense-mode
 no shutdown
!
router eigrp 1
 network 172.16.0.0
 network 192.168.0.0 0.0.255.255
 auto-summary
!
end
```

```
R2:
!
hostname R2
!
interface Loopback2
 ip address 192.168.2.1 255.255.255.0
 ip pim dense-mode
 ip igmp join-group 232.32.32.32
!
interface FastEthernet0/0
 ip address 172.16.20.2 255.255.255.0
 ip pim dense-mode
 no shutdown
!
interface Serial0/0/0
 bandwidth 64
 ip address 172.16.102.2 255.255.255.248
 ip pim dense-mode
 no shutdown
!
```

```
interface Serial0/0/1
 bandwidth 128
 ip address 172.16.203.2 255.255.255.248
 ip pim dense-mode
 clock rate 128000
 no shutdown
!
router eigrp 1
 network 172.16.0.0
 network 192.168.0.0 0.0.255.255
 auto-summary
!
end
```

R3:
```
!
hostname R3
!
interface Loopback3
 ip address 192.168.3.1 255.255.255.0
 ip pim dense-mode
 ip igmp join-group 232.32.32.32
!
interface FastEthernet0/0
 ip address 172.16.13.3 255.255.255.0
 ip pim dense-mode
 no shutdown
!
interface Serial0/0/0
 bandwidth 64
 ip address 172.16.103.3 255.255.255.248
 ip pim dense-mode
 clock rate 64000
 no shutdown
!
interface Serial0/0/1
 bandwidth 128
 ip address 172.16.203.3 255.255.255.248
 ip pim dense-mode
 no shutdown
!
router eigrp 1
 network 172.16.0.0
 network 192.168.0.0 0.0.255.255
 auto-summary
!
end
```

```
SW1# show run
!
hostname SW1
!
interface FastEthernet0/1
 switchport access vlan 13
 switchport mode access
!
interface FastEthernet0/3
 switchport access vlan 20
 switchport mode access
!
interface FastEthernet0/5
 switchport access vlan 13
 switchport mode access
```

```
!
interface Vlan20
 ip address 172.16.20.4 255.255.255.0
 no shutdown
!
ip default-gateway 172.16.20.2
!
end
```

Verify that each of the loopback interfaces subscribed to the multicast group using the **show ip igmp groups** command.

```
R1# show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface           Uptime    Expires   Last Reporter
Group Accounted
232.32.32.32     Loopback1           00:05:45  00:01:59  192.168.1.1
224.0.1.40       Loopback1           00:05:45  00:01:59  192.168.1.1

R2# show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface           Uptime    Expires   Last Reporter
Group Accounted
232.32.32.32     Loopback2           00:05:12  00:02:37  192.168.2.1
224.0.1.40       Loopback2           00:05:12  00:02:40  192.168.2.1

R3# show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface           Uptime    Expires   Last Reporter
Group Accounted
232.32.32.32     Loopback3           00:05:27  00:02:23  192.168.3.1
224.0.1.40       Loopback3           00:05:27  00:02:21  192.168.3.1
```

Step 3: Implement PIM-SM

Planning is important when creating PIM-SM networks. Because traffic is not simply flooded out all PIM interfaces as it is in dense mode networks, you make sure to place the rendezvous point (RP) in a central location.

CrushASodaCan.com's management decided the RP should be placed on R1. This location allows you to observe the cutover from the shared tree to the source tree, because R2 and R3 have shorter paths to 172.16.20.4 than those through R1.

Before enabling PIM-SM on the interfaces, set the static RP address as R1's Loopback1 address using the **ip pim rp-address** *rp-address* [*access-list*] command. You can either assign a router to be the global RP for all multicast groups or have it be on a per-group basis using an access list. In this case, map R1 as the RP only for the 232.32.32.32 multicast group. This command must be used on all routers running PIM-SM.

```
R1# conf t
R1(config)# access-list 32 permit 232.32.32.32
R1(config)# ip pim rp-address 192.168.1.1 32

R2# conf t
R2(config)# access-list 32 permit 232.32.32.32
```

```
R2(config)# ip pim rp-address 192.168.1.1

R3# conf t
R3(config)# access-list 32 permit 232.32.32.32
R3(config)# ip pim rp-address 192.168.1.1
```

What is the purpose of the RP in PIM-SM?

RPs provide authoritative joining points between the multicast sources and group members in a multicast network. RPs create a common point for multicast routers to root shared trees.

Given the fact that an RP can be set for specific groups, can you have more than one static RP in a multicast network? Explain.

Yes. You may have numerous static RPs in a network. Each group could potentially have a unique RP.

PIM-SM has more redundant ways of configuring RPs and mapping agents. Notice in this example that each multicast router has a static RP address pointing to a specific router on the network and a static mapping agent of itself. A mapping agent is simply a device that provides a mapping of a group to an RP. You have statically the group-to-RP mapping on every router, so each router acts as its own mapping agent.

In PIM-DM, all (S, G) entries are present in *every* IP multicast routing table on the multicast network while the source is broadcasting. If you were using many groups or many sources for each group, the size of the multicast routing table would dramatically increase from the one in Lab 7.2.

PIM-SM rendezvous points introduce multicast sources by acting as the network receptionist that knows the location of all multicast sources. PIM-SM creates a shared distribution tree represented by (*, G) in each multicast routing table. This shared tree is calculated using the RPF upstream interface to the RP for that group. Thus, the shared tree is essentially a shortest-path tree to the RP address.

PIM-SM reduces the amount of multicast state in the network by using (*, G) shared tree entries on routers through which no sources are sending traffic for a given group G. When a node in the topology subscribes to a multicast group, that node's upstream multicast router registers with the RP and forwards traffic down the shared tree to the receiver until it discovers a shorter path to the source in its unicast routing table. At that point, the multicast router prunes the stream from the shared tree and begins receiving the stream via its shortest unicast path to the source represented by its (S, G) entry.

PIM-SM is designed for sparsely populated multicast networks in which it is not necessary to flood multicast traffic to every subnet. Sparse mode is a general-purpose multicast mode that should be used in most circumstances.

As you did in Lab 7.2, issue a repeated ping from SW1 to generate the (S, G) state in PIM-DM before you apply PIM-SM to the interfaces.

```
SW1# ping
Protocol [ip]:
Target IP address: 232.32.32.32
Repeat count [1]: 100
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 232.32.32.32, timeout is 2 seconds:

Reply to request 0 from 172.16.20.2, 8 ms
Reply to request 0 from 172.16.102.1, 36 ms
Reply to request 0 from 172.16.203.3, 36 ms
Reply to request 1 from 172.16.20.2, 4 ms
Reply to request 1 from 172.16.102.1, 28 ms
Reply to request 1 from 172.16.203.3, 28 ms
Reply to request 2 from 172.16.20.2, 8 ms
Reply to request 2 from 172.16.102.1, 32 ms
Reply to request 2 from 172.16.203.3, 32 ms
```

Display the multicast routing table on each router with the **show ip mroute** command.

```
R1# 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,
       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
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 232.32.32.32), 01:42:56/00:03:04, RP 192.168.1.1, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Dense, 01:42:26/00:00:00
    Serial0/0/1, Forward/Dense, 01:42:27/00:00:00
    Serial0/0/0, Forward/Dense, 01:42:28/00:00:00
    Loopback1, Forward/Dense, 01:42:56/00:00:00

(172.16.20.4, 232.32.32.32), 00:01:20/00:02:59, flags: LT
  Incoming interface: FastEthernet0/0, RPF nbr 172.16.13.3
  Outgoing interface list:
    Loopback1, Forward/Dense, 00:01:20/00:00:00
    Serial0/0/0, Prune/Dense, 00:01:20/00:01:39
    Serial0/0/1, Prune/Dense, 00:01:20/00:01:39

(*, 224.0.1.40), 00:57:27/00:02:49, RP 0.0.0.0, flags: DCL
```

```
        Incoming interface: Null, RPF nbr 0.0.0.0
        Outgoing interface list:
          Serial0/0/1, Forward/Dense, 00:57:27/00:00:00
          Serial0/0/0, Forward/Dense, 00:57:35/00:00:00
          FastEthernet0/0, Forward/Dense, 00:57:35/00:00:00
          Loopback1, Forward/Dense, 00:57:35/00:00:00

    R2# 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,
           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
    Outgoing interface flags: H - Hardware switched, A - Assert winner
     Timers: Uptime/Expires
     Interface state: Interface, Next-Hop or VCD, State/Mode

    (*, 232.32.32.32), 01:43:06/stopped, RP 192.168.1.1, flags: SJCLF
      Incoming interface: Serial0/0/1, RPF nbr 172.16.203.3
      Outgoing interface list:
        Serial0/0/0, Forward/Dense, 01:43:02/00:00:00
        Loopback2, Forward/Dense, 01:43:06/00:00:00

    (172.16.20.4, 232.32.32.32), 00:01:24/00:02:59, flags: LFT
      Incoming interface: FastEthernet0/0, RPF nbr 0.0.0.0
      Outgoing interface list:
        Serial0/0/1, Forward/Dense, 00:01:26/00:00:00
        Loopback2, Forward/Dense, 00:01:26/00:00:00
        Serial0/0/0, Forward/Dense, 00:01:26/00:00:00, A

    (*, 224.0.1.40), 00:55:42/00:02:41, RP 0.0.0.0, flags: DCL
      Incoming interface: Null, RPF nbr 0.0.0.0
      Outgoing interface list:
        Serial0/0/1, Forward/Dense, 00:55:42/00:00:00
        Serial0/0/0, Forward/Dense, 00:55:42/00:00:00
        Loopback2, Forward/Dense, 00:55:42/00:00:00

    R3# 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,
           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
    Outgoing interface flags: H - Hardware switched, A - Assert winner
     Timers: Uptime/Expires
     Interface state: Interface, Next-Hop or VCD, State/Mode

    (*, 232.32.32.32), 01:43:06/00:03:26, RP 192.168.1.1, flags: SJCL
      Incoming interface: FastEthernet0/0, RPF nbr 172.16.13.1
      Outgoing interface list:
        Serial0/0/1, Forward/Dense, 01:42:37/00:00:00
        Serial0/0/0, Forward/Dense, 01:43:04/00:00:00
        Loopback3, Forward/Dense, 01:43:06/00:00:00

    (172.16.20.4, 232.32.32.32), 00:01:26/00:02:58, flags: LT
      Incoming interface: Serial0/0/1, RPF nbr 172.16.203.2
      Outgoing interface list:
```

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

```
    FastEthernet0/0, Forward/Dense, 00:01:28/00:00:00, A
    Loopback3, Forward/Dense, 00:01:28/00:00:00
    Serial0/0/0, Forward/Dense, 00:01:28/00:00:00, A

(*, 224.0.1.40), 01:43:07/00:02:30, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Dense, 01:42:35/00:00:00
    Serial0/0/1, Forward/Dense, 01:42:38/00:00:00
    Serial0/0/0, Forward/Dense, 01:43:20/00:00:00
    Loopback3, Forward/Dense, 01:43:21/00:00:00
```

Notice that R1 has 0.0.0.0 as the RPF neighbor of the (*, G) entry, indicating that it is the RP for the 232.32.32.32 multicast group. R2 and R3 listen to multicast traffic on the shared tree coming from their RPF neighbor for that (*, G) entry in their multicast routing tables.

Enable PIM-SM on all interfaces using the **ip pim sparse-mode** command in interface configuration mode.

```
R1(config)# interface loopback 1
R1(config-if)# ip pim sparse-mode
R1(config-if)# interface fastethernet 0/0
R1(config-if)# ip pim sparse-mode
R1(config-if)# interface serial 0/0/0
R1(config-if)# ip pim sparse-mode
R1(config-if)# interface serial 0/0/1
R1(config-if)# ip pim sparse-mode

R2(config)# interface loopback 2
R2(config-if)# ip pim sparse-mode
R2(config-if)# interface fastethernet 0/0
R2(config-if)# ip pim sparse-mode
R2(config-if)# interface serial 0/0/0
R2(config-if)# ip pim sparse-mode
R2(config-if)# interface serial 0/0/1
R2(config-if)# ip pim sparse-mode

R3(config)# interface loopback 1
R3(config-if)# ip pim sparse-mode
R3(config-if)# interface fastethernet 0/0
R3(config-if)# ip pim sparse-mode
R3(config-if)# interface serial 0/0/0
R3(config-if)# ip pim sparse-mode
R3(config-if)# interface serial 0/0/1
R3(config-if)# ip pim sparse-mode
```

Now that multicast routing and PIM-SM are enabled on R2, do you think a ping from SW1 to the 232.32.32.32 group will receive any replies? Explain.

Yes. ICMP echo replies will be sent from each router because PIM-SM is active throughout the network and there are subscribers to the 232.32.32.32 group.

Verify your answer by pinging the multicast group from SW1. You will not necessarily receive replies from the IP address of the interface on which the

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

영원한 우리땅, 독도
Dokdo forever
belongs to Korea

multicast packet was received but rather from whichever interface on the responding router encapsulated the return packet.

```
SW1# ping 232.32.32.32

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 232.32.32.32, timeout is 2 seconds:

Reply to request 0 from 172.16.20.2, 4 ms
Reply to request 0 from 172.16.102.1, 32 ms
Reply to request 0 from 172.16.203.3, 32 ms
```

Based on your understanding of IP unicast routing, why are the ICMP echo replies sent from the interfaces shown above?

The interface where the ICMP replies originate is determined by the unicast routing table. The outgoing packets are labeled with the IP address of the interface out which packets will flow to 172.16.20.4.

Step 4: Verify PIM Adjacencies

In this step, we explore PIM adjacencies and how PIM functions over various Layer 2 media.

Issue the **show ip pim neighbors** command to display all adjacent PIM routers.

```
R1# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface               Uptime/Expires     Ver   DR
Address                                                            Prio/Mode
172.16.13.3       FastEthernet0/0         00:07:22/00:01:19 v2    1 / DR S
172.16.102.2      Serial0/0/0             00:07:23/00:01:22 v2    1 / S
172.16.103.3      Serial0/0/1             00:07:23/00:01:29 v2    1 / S

R2# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface               Uptime/Expires     Ver   DR
Address                                                            Prio/Mode
172.16.102.1      Serial0/0/0             00:08:27/00:01:20 v2    1 / S
172.16.203.3      Serial0/0/1             00:08:29/00:01:22 v2    1 / S

R3# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface               Uptime/Expires     Ver   DR
Address                                                            Prio/Mode
172.16.13.1       FastEthernet0/0         00:08:28/00:01:23 v2    1 / S
172.16.103.1      Serial0/0/0             00:08:59/00:01:19 v2    1 / S
172.16.203.2      Serial0/0/1             00:08:32/00:01:39 v2    1 / S
```

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

PIM-DM uses a designated router (DR) as the source of IGMPv1 queries. PIM-SM also supports this type of behavior because it is required in the IGMPv1 protocol.

Which router is the IGMP querier for the 172.16.13.0/24 subnet?

R1 is the DR for the 172.16.13.0/24 subnet and therefore functions as the IGMP querier.

Display information about PIM-enabled interfaces with the **show ip pim interface detail** command.

```
R1# show ip pim interface detail

Address          Interface         Ver/    Nbr    Query   DR      DR
                                   Mode    Count  Intvl   Prior
192.168.1.1      Loopback1         v2/S    0      30      1       192.168.1.1
172.16.13.1      FastEthernet0/0   v2/S    1      30      1       172.16.13.3
172.16.102.1     Serial0/0/0       v2/S    1      30      1       0.0.0.0
172.16.103.1     Serial0/0/1       v2/S    1      30      1       0.0.0.0

R2# show ip pim interface detail

Address          Interface         Ver/    Nbr    Query   DR      DR
                                   Mode    Count  Intvl   Prior
192.168.2.1      Loopback2         v2/S    0      30      1       192.168.2.1
172.16.20.2      FastEthernet0/0   v2/S    0      30      1       172.16.20.2
172.16.102.2     Serial0/0/0       v2/S    1      30      1       0.0.0.0
172.16.203.2     Serial0/0/1       v2/S    1      30      1       0.0.0.0

R3# show ip pim interface detail

Address          Interface         Ver/    Nbr    Query   DR      DR
                                   Mode    Count  Intvl   Prior
192.168.3.1      Loopback3         v2/S    0      30      1       192.168.3.1
172.16.13.3      FastEthernet0/0   v2/S    1      30      1       172.16.13.3
172.16.103.3     Serial0/0/0       v2/S    1      30      1       0.0.0.0
172.16.203.3     Serial0/0/1       v2/S    1      30      1       0.0.0.0
```

The serial interfaces use the default DR address 0.0.0.0 as the DR for the interface. Because a multicast is received by either 0 or 1 remote routers on a serial segment, PIM does need not to set up a complex neighbor relationship.

The Ver/Mode column displays the PIM version and mode running on each interface. S refers to sparse mode. D is used for dense mode, and SD indicates the hybrid sparse-dense mode.

Step 5: Verify Multicast Routing Operation

On each router, use the **mrinfo** command to view information about the connected multicast-enabled routers.

```
R1# mrinfo
```

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

```
172.16.13.1 [version  12.4] [flags: PMA]:
  192.168.1.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
  172.16.13.1 -> 172.16.13.3 [1/0/pim]
  172.16.102.1 -> 172.16.102.2 [1/0/pim]
  172.16.103.1 -> 172.16.103.3 [1/0/pim]

R2# mrinfo
172.16.20.2 [version  12.4] [flags: PMA]:
  192.168.2.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
  172.16.20.2 -> 0.0.0.0 [1/0/pim/querier/leaf]
  172.16.102.2 -> 172.16.102.1 [1/0/pim]
  172.16.203.2 -> 172.16.203.3 [1/0/pim]

R3# mrinfo
172.16.13.3 [version  12.4] [flags: PMA]:
  192.168.3.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
  172.16.13.3 -> 172.16.13.1 [1/0/pim/querier]
  172.16.103.3 -> 172.16.103.1 [1/0/pim]
  172.16.203.3 -> 172.16.203.2 [1/0/pim]
```

Each router realizes that the loopback interfaces are topological leaves in which
PIM will never establish an adjacency with any other routers. These routers also
record the neighboring multicast router addresses and the multicast routing
protocols they utilize.

As discussed in the previous step, IGMPv1 requires the multicast routing
protocol to elect a querier on a multi-access network.

Based on the information in the output of the **mrinfo** command, which interface
is the querier on VLAN 13?

R1's Fast Ethernet 0/0 interface with IP address 172.16.13.1 is the IGMP
querier for the 172.16.13.0/24 subnet.

Generate a stream of multicast data to the group by issuing an extended ping
from SW1 with a repeat count of 100.

```
SW1# ping
Protocol [ip]:
Target IP address: 232.32.32.32
Repeat count [1]: 100
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 232.32.32.32, timeout is 2 seconds:

Reply to request 0 from 172.16.20.2, 8 ms
Reply to request 0 from 172.16.102.1, 36 ms
Reply to request 0 from 172.16.203.3, 36 ms
Reply to request 1 from 172.16.20.2, 4 ms
Reply to request 1 from 172.16.102.1, 28 ms
Reply to request 1 from 172.16.203.3, 28 ms
Reply to request 2 from 172.16.20.2, 8 ms
```

```
Reply to request 2 from 172.16.102.1, 32 ms
Reply to request 2 from 172.16.203.3, 32 ms
...
```

On each of the routers, you should see that PIM and IGMP have communicated to install the 232.32.32.32 multicast group in the multicast routing table. Verify this with the **show ip mroute** command on each router.

```
R1# 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,
       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
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 232.32.32.32), 02:26:04/00:03:21, RP 192.168.1.1, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse, 02:25:34/00:03:21
    Serial0/0/1, Forward/Sparse, 02:25:34/00:00:00
    Serial0/0/0, Forward/Sparse, 02:25:35/00:00:00
    Loopback1, Forward/Sparse, 02:26:04/00:02:43

(172.16.20.4, 232.32.32.32), 00:00:14/00:02:59, flags: T
  Incoming interface: FastEthernet0/0, RPF nbr 172.16.13.3
  Outgoing interface list:
    Loopback1, Forward/Sparse, 00:00:14/00:02:45
    Serial0/0/1, Forward/Sparse, 00:00:14/00:02:45

(*, 224.0.1.40), 01:40:34/00:02:45, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0/1, Forward/Sparse, 01:40:34/00:00:00
    Serial0/0/0, Forward/Sparse, 01:40:34/00:00:00
    FastEthernet0/0, Forward/Sparse, 01:40:36/00:00:00
    Loopback1, Forward/Sparse, 01:40:36/00:02:44

R2# 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,
       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
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 232.32.32.32), 02:26:14/stopped, RP 192.168.1.1, flags: SJCLF
  Incoming interface: Serial0/0/1, RPF nbr 172.16.203.3
  Outgoing interface list:
    Serial0/0/0, Forward/Sparse, 02:26:10/00:00:00
    Loopback2, Forward/Sparse, 02:26:14/00:02:32
```

```
(172.16.20.4, 232.32.32.32), 00:00:19/00:03:28, flags: LFT
  Incoming interface: FastEthernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0/1, Forward/Sparse, 00:00:20/00:03:09, A
    Loopback2, Forward/Sparse, 00:00:20/00:02:39
    Serial0/0/0, Forward/Sparse, 00:00:20/00:02:39, A

(*, 224.0.1.40), 01:38:49/00:02:34, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0/1, Forward/Sparse, 01:38:49/00:00:00
    Serial0/0/0, Forward/Sparse, 01:38:49/00:00:00
    Loopback2, Forward/Sparse, 01:38:49/00:02:34

R3# 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,
       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
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 232.32.32.32), 02:26:16/00:02:47, RP 192.168.1.1, flags: SJCL
  Incoming interface: FastEthernet0/0, RPF nbr 172.16.13.1
  Outgoing interface list:
    Serial0/0/1, Forward/Sparse, 00:30:25/00:02:32
    Loopback3, Forward/Sparse, 02:26:16/00:02:23

(172.16.20.4, 232.32.32.32), 00:00:24/00:03:29, flags: LT
  Incoming interface: Serial0/0/1, RPF nbr 172.16.203.2
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse, 00:00:25/00:03:06
    Loopback3, Forward/Sparse, 00:00:25/00:02:34

(*, 224.0.1.40), 02:26:17/00:02:22, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback3, Forward/Sparse, 02:26:17/00:02:22
```

Multicast forwarding state refers to the current set of (S, G) and (*, G) entries in the multicast routing table. In general, "state" is defined as a stored set of properties at a given point in time. In unicast tables, the state for a destination network is referred to as the route to that network. A unicast IP route is essentially stored in the unicast routing table as a destination network address and mask, along with the next-hop routing information.

Multicast routing tables contain significantly more information than unicast routing tables because a multicast routing table must track many time-based variables for sources and group members. For instance, a multicast router must track the time when the last packet from a source is received, the outgoing interface list, and the RPF neighbors for each source-based and shared tree it maintains. These variables are not stored in a PIM topology table but in the

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

영원한 우리땅, 독도
Dokdo forever
belongs to Korea

multicast routing table natively. Multicast routing state is subject to quick changes based on source registrations and group subscriptions.

In PIM-DM, the IP multicast routing table indicates whether the PIM interfaces negotiated to flood multicast data to that group out an interface or pruned (S, G) multicast traffic from being sent out that interface.

In the multicast tables generated above, notice that PIM-SM does not refer to interfaces as pruned.
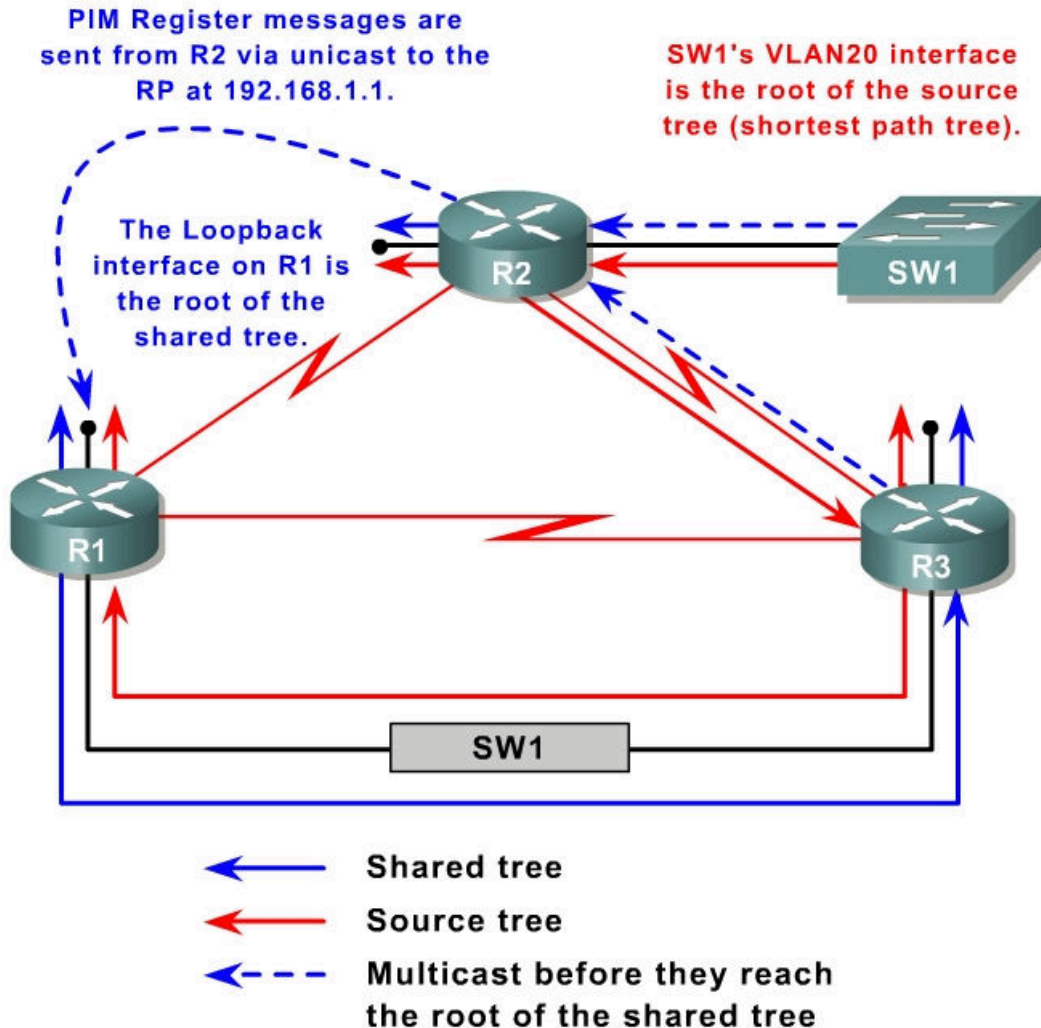
Why does PIM-SM not need to record the interfaces over which the router will prune (S, G) streams from flowing?

PIM-DM acts by flooding and pruning. Thus it maintains an outgoing interface list of all of the multicast interfaces facing away from the upstream source. PIM-DM blocks traffic to unnecessary subnets by pruning those interfaces from the source tree (shortest path tree). PIM-SM only sends multicast traffic out interfaces that have requested it. Therefore, PIM-SM's default state is to not send multicast traffic out of any of its interfaces. When requests are received to forward traffic, it *adds* those interfaces to the outgoing interface list. When requests are received to stop forwarding traffic, PIM-SM *removes* those interfaces from the outgoing interface list. PIM-SM does not need to record in the outgoing interface list any interfaces that do not request traffic because its default state is to not forward that traffic.

When a PIM-SM router receives an explicit join message from IGMP on an interface, it adds that interface to the outgoing interface list of the source-based tree. The source-based tree is also known as the shortest-path tree and is represented in the routing table by an (S, G) entry. Since PIM-SM is a registration-based multicast service, it simply drops from the list the interfaces to which it does not need to send data.

Based on the output of the **show ip mroute** command shown above:
- Trace the (*, G) shared tree on the diagram below. Be sure to reference the RPF interfaces for the (*, G) traffic.
- Trace the (172.16.20.4, 232.32.32.32) source-based tree.

PIM Register messages are sent from R2 via unicast to the RP at 192.168.1.1.

SW1's VLAN20 interface is the root of the source tree (shortest path tree).

The Loopback interface on R1 is the root of the shared tree.

Shared tree

Source tree

Multicast before they reach the root of the shared tree

From which PIM neighbor will each loopback interface receive data?
- R1 Loopback 1: R3's Fast Ethernet 0/0 interface, 172.16.13.3
- R2 Loopback 2: PIM will forward (172.16.20.4, 232.32.32.32) traffic directly to Loopback 2.
- R3 Loopback 3: R2's Serial 0/0/1 interface, 172.16.203.2

Step 6: Verify PIM-SM Registration and SPT Cutover

You are interested in viewing how sources register with the RP. You are also interested in exploring pruning data from the shared tree and the cutover to the source tree. Wait for the multicast stream from 172.16.20.4 to complete and for the (S, G) state in the multicast routing tables to expire. Then issue the **debug ip igmp** and **debug ip pim** commands on all routers.

PIM relies on the unicast IP routing table to construct shortest-path trees from the sources to the multicast subscribers. PIM interfaces send control messages to determine which neighbor is closer to the source in terms of the unicast routing information on each neighbor. PIM-SM neighbors elect a particular

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

영원한 우리땅, 독도
Dokdo forever belongs to Korea

router on the subnet as the forwarder for that (S, G) pair and then prune that (S, G) pair from being forwarded by any other routers on the subnet.

Display the unicast routing table entry for 172.16.20.4 on each router.

```
R1# show ip route 172.16.20.4
Routing entry for 172.16.20.0/24
  Known via "eigrp 1", distance 90, metric 20517120, type internal
  Redistributing via eigrp 1
  Last update from 172.16.13.3 on FastEthernet0/0, 03:05:58 ago
  Routing Descriptor Blocks:
  * 172.16.13.3, from 172.16.13.3, 03:05:58 ago, via FastEthernet0/0
      Route metric is 20517120, traffic share count is 1
      Total delay is 20200 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2

R2# show ip route 172.16.20.4
Routing entry for 172.16.20.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via eigrp 1
  Routing Descriptor Blocks:
  * directly connected, via FastEthernet0/0
      Route metric is 0, traffic share count is 1

R3# show ip route 172.16.20.4
Routing entry for 172.16.20.0/24
  Known via "eigrp 1", distance 90, metric 20514560, type internal
  Redistributing via eigrp 1
  Last update from 172.16.203.2 on Serial0/0/1, 03:06:09 ago
  Routing Descriptor Blocks:
  * 172.16.203.2, from 172.16.203.2, 03:06:09 ago, via Serial0/0/1
      Route metric is 20514560, traffic share count is 1
      Total delay is 20100 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Recall from the previous lab: What does PIM use to select the forwarder for a multicast routing entry on a point-to-point link?

PIM Assert messages are used to select the forwarder for each (S, G) pair on point-to-point links.

Begin the same extended ping on SW1 to the multicast group. Examine the debugging messages shown on R1.

```
R1# debug ip pim
*Nov  6 16:10:08.879: PIM(0): Received v2 Register on FastEthernet0/0 from
172.16.203.2
*Nov  6 16:10:08.879:      for 172.16.20.4, group 232.32.32.32
*Nov  6 16:10:08.883: PIM(0): Insert (172.16.20.4,232.32.32.32) join in nbr
172.16.13.3's queue
*Nov  6 16:10:08.883: PIM(0): Forward decapsulated data packet for
232.32.32.32 on FastEthernet0/0
*Nov  6 16:10:08.883: PIM(0): Forward decapsulated data packet for
232.32.32.32 on Serial0/0/1
```

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

```
*Nov  6 16:10:08.883: PIM(0): Forward decapsulated data packet for
232.32.32.32 on Serial0/0/0
*Nov  6 16:10:08.883: PIM(0): Forward decapsulated data packet for
232.32.32.32 on Loopback1
... continued below ...
```

Initially, R2 begins encapsulating multicast packets from 172.16.20.4 in unicast
PIMv2 register packets and sends those unicast messages to the RP. This
packet is sent via R2's shortest path to 192.168.1.1 and received on R1's
FastEthernet0/0 interface.

The RP sends a PIM control message to the downstream neighbor in the
shared tree informing R3 to record (172.16.20.4, 232.32.32.32) state in its
multicast routing table. R1 then forwards the newly decapsulated *multicast*
packets to 232.32.32.32 down the shared tree.

```
... continued from above ...
R1#
*Nov  6 16:10:08.883: PIM(0): Building Join/Prune packet for nbr 172.16.13.3
*Nov  6 16:10:08.883: PIM(0): Adding v2 (172.16.20.4/32, 232.32.32.32), S-bit
Join
*Nov  6 16:10:08.883: PIM(0): Send v2 join/prune to 172.16.13.3
(FastEthernet0/0)
*Nov  6 16:10:08.907: PIM(0): Received v2 Assert on Serial0/0/0 from
172.16.102.2
*Nov  6 16:10:08.907: PIM(0): Assert metric to source 172.16.20.4 is [0/0]
*Nov  6 16:10:08.907: PIM(0): We lose, our metric [90/20517120]
*Nov  6 16:10:08.907: PIM(0): Prune Serial0/0/0/224.0.0.2 from
(172.16.20.4/32, 232.32.32.32) - deleted
... continued below ...
```

Because Loopback1 has received the first multicast packet from SW1, R1 cuts
over to the SPT. R1 builds and sends a join packet for R3 because R3 is its
RPF neighbor for source 172.16.20.4. As soon as PIM routers accumulate state
about the (S, G) pair, they can begin to cut over to the source-based tree.

By default, when a PIM router with a connected multicast leaf node receives the
first packet from a source to a group via the shared tree, it cuts over to the
shortest path tree (SPT). You can also manual configure this as a bandwidth
threshold using the **ip pim spt-threshold** command. When the (S, G) stream
reaches the threshold bandwidth (in kbps), the PIM router switches to the SPT.

The cutover process is initiated on R2 and R3 by the first (172.16.20.4,
232.32.32.32) multicast packet. R1 and R2 compare administrative distances
and metrics regarding their shortest path to the source, and R2, the winner,
sends a prune message to R1 indicating that it should not send traffic out
Serial0/0/0 to 232.32.32.32. R2's SPT state process completes.

```
... continued from above ...
R1#
*Nov  6 16:10:10.879: PIM(0): Received v2 Register on FastEthernet0/0 from
172.16.203.2
*Nov  6 16:10:10.879:        for 172.16.20.4, group 232.32.32.32
```

영원한 우리땅, 독도
Dokdo forever
belongs to Korea

```
*Nov  6 16:10:10.879: PIM(0): Forward decapsulated data packet for
232.32.32.32 on FastEthernet0/0
*Nov  6 16:10:10.879: PIM(0): Forward decapsulated data packet for
232.32.32.32 on Serial0/0/1
*Nov  6 16:10:10.879: PIM(0): Forward decapsulated data packet for
232.32.32.32 on Serial0/0/0
*Nov  6 16:10:10.879: PIM(0): Forward decapsulated data packet for
232.32.32.32 on Loopback1
*Nov  6 16:10:10.891: PIM(0): Received v2 Join/Prune on FastEthernet0/0 from
172.16.13.3, to us
*Nov  6 16:10:10.895: PIM(0): Prune-list: (172.16.20.4/32, 232.32.32.32) RPT-
bit set
... continued below ...
```

R2 continues encapsulating multicast data in unicast packets and sending them
to the RP via R3.

Since all the multicast receivers have now obtained (S, G) SPT state for
(172.16.20.4, 232.32.32.32), R3 decides it does not need R1 to continue
sending this traffic down the shared tree. R3 sends a prune message to R1 via
the Ethernet link between them.

```
... continued from above ...
R1#
*Nov  6 16:10:10.903: PIM(0): Received v2 Assert on Serial0/0/0 from
172.16.102.2
*Nov  6 16:10:10.903: PIM(0): Assert metric to source 172.16.20.4 is [0/0]
*Nov  6 16:10:10.979: PIM(0): Insert (172.16.20.4,232.32.32.32) join in nbr
172.16.13.3's queue
*Nov  6 16:10:10.979: PIM(0): Building Join/Prune packet for nbr 172.16.13.3
*Nov  6 16:10:10.979: PIM(0): Adding v2 (172.16.20.4/32, 232.32.32.32), S-bit
Join
*Nov  6 16:10:10.979: PIM(0): Send v2 join/prune to 172.16.13.3
(FastEthernet0/0)
*Nov  6 16:10:12.879: PIM(0): Received v2 Register on FastEthernet0/0 from
172.16.203.2
*Nov  6 16:10:12.879:       for 172.16.20.4, group 232.32.32.32
*Nov  6 16:10:12.879: PIM(0): Send v2 Register-Stop to 172.16.203.2 for
172.16.20.4, group 232.32.32.32
```

The first two messages indicate that R2 sent a periodic Assert message on the
serial link between them. The Insert, Building, Adding, and Send messages
represent R1 acting as a standard multicast router and sending a join message
to its upstream SPT neighbor which requests (S, G) traffic to continue.

R2 sends another Register packet to the RP with an encapsulated multicast
packet. R1, acting as the RP, checks its routing table for outgoing interfaces on
the shared tree. Since all downstream branches of the shared tree have been
pruned, indicating that R2 and R3 have switched over to the SPT, R1 sends a
Register-Stop message to R2 telling it to stop forwarding multicast packets as
unicasts to the RP.

At this point, all multicast routers are listening directly to the source via their
shortest path to 172.16.20.4, and the current receivers do not need the RP or
the shared tree because they are listening directly to the source via the SPT. It

영원한 우리땅, 독도
Dokdo forever
belongs to Korea

would not be unusual if the router acting as the RP was still a part of the SPT, as it is in this case, because R1's Loopback1 interface subscribes to the 232.32.32.32 group. However, R1 is acting simply as a subscriber to the group and no longer as the RP.

Why does R2 initially encapsulate the multicast data in a unicast packet?

R2 cannot allow the traffic directly into the multicast network, or PIM-SM would degrade into PIM-DM. Traffic first needs to be sent to the rendezvous point at which other routers have rooted their shared trees. Traffic will then be sent down the shared tree until all sources have switched over to the source-based tree. In order to send the traffic to the rendezvous point without sending it through the multicast network (thereby creating multicast state), PIM-SM sends the data as a unicast to the RP.

According to the debugging output shown above, what happens to the multicast data once it reaches the RP?

At the RP, the unicast packets are de-encapsulated and the packets are sent as a multicast down the shared tree.

If R2 simply forwarded the data to all receivers via multicast packets, does that sound similar to another multicast routing protocol or procedure?

If R2 simply forwarded the data to all connected receivers using multicast packets, the PIM-SM network would degrade to PIM-DM.

Conclusion

Why is PIM-SM more efficient?

PIM-SM only stores multicast routing information about the shared trees and about source trees for which it has had some request. PIM-SM is also more efficient than PIM-DM because it does not utilize the amount of bandwidth on leaf networks that have no interest in the multicast traffic.

Why did the routers cut over to the SPT?

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

Routers cut over to the SPT to minimize the delay between the sources and the group members. This effectively prunes the shared tree out of the branch of the source tree between the source and that particular group member.

Issue the **show ip mroute** command on all routers. Compare the output to the output in Step 3 before you implemented PIM-SM.

```
R1# show ip mroute
<output omitted>

(*, 232.32.32.32), 04:27:40/stopped, RP 192.168.1.1, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse, 04:27:10/00:02:45
    Serial0/0/1, Forward/Sparse, 04:27:11/00:00:00
    Serial0/0/0, Forward/Sparse, 04:27:12/00:00:00
    Loopback1, Forward/Sparse, 04:27:40/00:02:04

(172.16.20.4, 232.32.32.32), 00:00:05/00:02:58, flags: T
  Incoming interface: FastEthernet0/0, RPF nbr 172.16.13.3
  Outgoing interface list:
    Loopback1, Forward/Sparse, 00:00:05/00:02:54
    Serial0/0/1, Forward/Sparse, 00:00:05/00:02:54

(*, 224.0.1.40), 03:42:11/00:02:01, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0/1, Forward/Sparse, 03:42:11/00:00:00
    Serial0/0/0, Forward/Sparse, 03:42:11/00:00:00
    FastEthernet0/0, Forward/Sparse, 03:42:12/00:00:00
    Loopback1, Forward/Sparse, 03:42:12/00:02:00

R2# show ip mroute
<output omitted>

(*, 232.32.32.32), 04:28:00/stopped, RP 192.168.1.1, flags: SJCLF
  Incoming interface: Serial0/0/1, RPF nbr 172.16.203.3
  Outgoing interface list:
    Serial0/0/0, Forward/Sparse, 04:27:56/00:00:00
    Loopback2, Forward/Sparse, 04:28:00/00:02:49

(172.16.20.4, 232.32.32.32), 00:00:19/00:03:18, flags: LFT
  Incoming interface: FastEthernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0/1, Forward/Sparse, 00:00:20/00:03:09, A
    Loopback2, Forward/Sparse, 00:00:20/00:02:48
    Serial0/0/0, Forward/Sparse, 00:00:20/00:02:39, A

(*, 224.0.1.40), 03:40:35/00:02:43, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0/1, Forward/Sparse, 03:40:35/00:00:00
    Serial0/0/0, Forward/Sparse, 03:40:35/00:00:00
    Loopback2, Forward/Sparse, 03:40:35/00:02:43

R3# show ip mroute
<output omitted>

(*, 232.32.32.32), 04:28:01/00:02:47, RP 192.168.1.1, flags: SJCL
```

영원한 우리땅, 독도
Dokdo forever
belongs to Korea

```
    Incoming interface: FastEthernet0/0, RPF nbr 172.16.13.1
    Outgoing interface list:
      Serial0/0/1, Forward/Sparse, 02:32:10/00:02:46
      Loopback3, Forward/Sparse, 04:28:01/00:02:37

(172.16.20.4, 232.32.32.32), 00:00:23/00:03:28, flags: LT
  Incoming interface: Serial0/0/1, RPF nbr 172.16.203.2
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse, 00:00:24/00:03:07
    Loopback3, Forward/Sparse, 00:00:24/00:02:36

(*, 224.0.1.40), 04:28:02/00:02:31, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback3, Forward/Sparse, 04:28:02/00:02:31
```

Does the implementation of PIM-SM in CrushASodaCan.com's multicast network save bandwidth compared to the PIM-DM implementation? Explain.

Yes. Bandwidth has been saved because the flooding and pruning inherent in PIM-DM does not take place.

Although not part of this lab activity, could the implementation of PIM-SM also save memory by decreasing multicast routing state? Explain.

Normally, the implementation of PIM-SM indeed does decrease multicast state. Since only state about shared and source trees that pass through the local router need be stored, all other multicast state does not get created. The only (*, G) and (S, G) entries on a router running PIM-SM are for the trees that pass through that router.

Final Configurations

```
R1# show run
!
hostname R1
!
interface Loopback1
 ip address 192.168.1.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 232.32.32.32
!
interface FastEthernet0/0
 ip address 172.16.13.1 255.255.255.0
 ip pim sparse-mode
 no shutdown
!
interface Serial0/0/0
 bandwidth 64
 ip address 172.16.102.1 255.255.255.248
 ip pim sparse-mode
 clock rate 64000
```

```
 no shutdown
!
interface Serial0/0/1
 bandwidth 64
 ip address 172.16.103.1 255.255.255.248
 ip pim sparse-mode
 no shutdown
!
router eigrp 1
 network 172.16.0.0
 network 192.168.0.0 0.0.255.255
 auto-summary
!
ip pim rp-address 192.168.1.1 32
!
access-list 32 permit 232.32.32.32
!
end

R2# show run
!
hostname R2
!
interface Loopback2
 ip address 192.168.2.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 232.32.32.32
!
interface FastEthernet0/0
 ip address 172.16.20.2 255.255.255.0
 ip pim sparse-mode
 no shutdown
!
interface Serial0/0/0
 bandwidth 64
 ip address 172.16.102.2 255.255.255.248
 ip pim sparse-mode
 no shutdown
!
interface Serial0/0/1
 bandwidth 128
 ip address 172.16.203.2 255.255.255.248
 ip pim sparse-mode
 clock rate 128000
 no shutdown
!
router eigrp 1
 network 172.16.0.0
 network 192.168.0.0 0.0.255.255
 auto-summary
!
ip pim rp-address 192.168.1.1 32
!
access-list 32 permit 232.32.32.32
!
end

R3# show run
!
hostname R3
!
interface Loopback3
 ip address 192.168.3.1 255.255.255.0
```

CCNP: Building Scalable Internetworks v5.0 - Lab 7-3   Copyright © 2006, Cisco Systems, Inc

```
 ip pim sparse-mode
 ip igmp join-group 232.32.32.32
!
interface FastEthernet0/0
 ip address 172.16.13.3 255.255.255.0
 ip pim sparse-mode
 no shutdown
!
interface Serial0/0/0
 bandwidth 64
 ip address 172.16.103.3 255.255.255.248
 ip pim sparse-mode
 clock rate 64000
 no shutdown
!
interface Serial0/0/1
 bandwidth 128
 ip address 172.16.203.3 255.255.255.248
 ip pim sparse-mode
 no shutdown
!
router eigrp 1
 network 172.16.0.0
 network 192.168.0.0 0.0.255.255
 auto-summary
!
ip pim rp-address 192.168.1.1 32
!
access-list 32 permit 232.32.32.32
!
end

SW1# show run
!
hostname SW1
!
interface FastEthernet0/1
 switchport access vlan 13
 switchport mode access
!
interface FastEthernet0/3
 switchport access vlan 20
 switchport mode access
!
interface FastEthernet0/5
 switchport access vlan 13
 switchport mode access
!
interface Vlan20
 ip address 172.16.20.4 255.255.255.0
 no shutdown
!
ip default-gateway 172.16.20.2
!
end
```