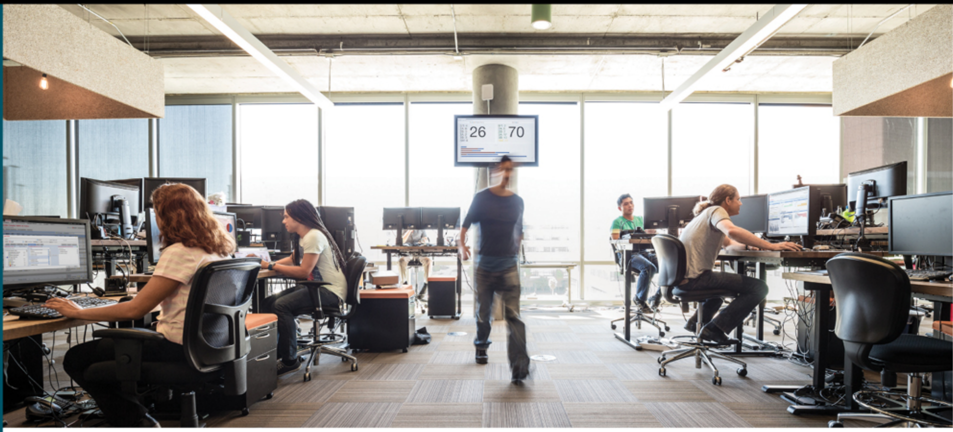

IP Multicast, Volume II

Advanced Multicast Concepts and Large-Scale Multicast Design

Josh Loveless, CCIE® No. 16638
Ray Blair, CCIE® No. 7050
Arvind Durai, CCIE® No. 7016

ciscopress.com

IP Multicast, Volume II
Advanced Multicast Concepts and Large-Scale Multicast Design

Josh Loveless, CCIE No. 16638
Ray Blair, CCIE No. 7050
Arvind Durai, CCIE No. 7016

Cisco Press
800 East 96th Street
Indianapolis, Indiana 46240 USA

```
 ipv6 address 2001:192:168::2/128
 ipv6 enable
 ipv6 ospf 65000 area 0
!
interface Ethernet0/0
 no ip address
 ipv6 address 2001:192:168:21::2/64
 ipv6 enable
 ipv6 ospf 65000 area 0
!
interface Ethernet0/1
 no ip address
 ipv6 address 2001:192:168:32::2/64
 ipv6 enable
 ipv6 ospf 65000 area 0
!
interface Ethernet0/2
 no ip address
 ipv6 address 2001:192:168:52::2/64
 ipv6 enable
 ipv6 ospf 65000 area 0
!
ipv6 router ospf 65000
 router-id 192.168.0.2


hostname R3
ipv6 unicast-routing
ipv6 multicast-routing
!
interface Loopback0
 ip address 192.168.0.3 255.255.255.255
 ipv6 address 2001:192:168::3/128
 ipv6 enable
 ipv6 ospf 65000 area 0
!
interface Ethernet0/0
 no ip address
 load-interval 30
 ipv6 address 2001:192:168:31::3/64
 ipv6 enable
 ipv6 ospf 65000 area 0
!
interface Ethernet0/1
 no ip address
```

```
 ipv6 address 2001:192:168:32::3/64
 ipv6 enable
 ipv6 ospf 65000 area 0
!
interface Ethernet0/2
 no ip address
 ipv6 address 2001:192:168:63::3/64
 ipv6 enable
 ipv6 mld join-group FF73:105:2001:192::1
 ipv6 ospf 65000 area 0
!
ipv6 router ospf 65000
 router-id 192.168.0.3
```

As you can see from the configurations in Example 1-31, there isn't anything too fancy. The highlighted commands on R1 defines R1 as the RP using the loopback 0 interface, with the **ipv6 pim rp-address 2001:192::1** command, and the second is on R3, which statically defines a join group by using the **ipv6 mld join-group FF73:105:2001:192::1** command.

Note The **ipv6 mld join-group** command should be used only temporarily, for trouble-shooting purposes only.

You may have noticed that the only router with an RP mapping is R1. Because you are embedding the RP information in the multicast message, it is not necessary to define an RP on every router.

From R2, you can watch the behavior in action by using a simple **ping** command. As shown in Example 1-32, you can ping the FF73:105:2001:192::1 address configured as a join group on R3.

Example 1-32 *Embedded RP Example*

```
R2# ping FF73:105:2001:192::1
Output Interface: ethernet0/1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FF73:105:2001:192::1, timeout is 2 seconds:
Packet sent with a source address of 2001:192:168:32::2

Reply to request 0 received from 2001:192:168:63::3, 8 ms
Reply to request 1 received from 2001:192:168:63::3, 1 ms
Reply to request 2 received from 2001:192:168:63::3, 1 ms
Reply to request 3 received from 2001:192:168:63::3, 1 ms
Reply to request 4 received from 2001:192:168:63::3, 1 ms
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
5 multicast replies and 0 errors.
```

You can see that all the ping packets received replies from R3 (2001:192:168:63::3).

On R3, you can verify the existence of both the (*, G) and (S, G) entries with the command shown in Example 1-33.

Example 1-33 *Verifying the PIM Entry with RP Mapping*

```
R3# show ipv6 mroute FF73:105:2001:192::1
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF73:105:2001:192::1), 00:13:34/00:02:55, RP 2000::1, flags: SCL
  Incoming interface: Null
  RPF nbr: ::
  Immediate Outgoing interface list:
    Ethernet0/2, Forward, 00:13:34/00:02:55

(2001:192:168:32::2, FF73:105:2001:192::1), 00:03:02/00:00:27, flags: SFT
  Incoming interface: Ethernet0/1
  RPF nbr: FE80::A8BB:CCFF:FE00:210, Registering
  Immediate Outgoing interface list:
    Tunnel0, Forward, 00:03:02/never
    Ethernet0/2, Forward, 00:03:02/00:02:35
```

Notice that multicast messages are received from interface Ethernet0/1 and sent to the destination interface, Ethernet0/2.

Even though this example is an intra-domain configuration, using IPv6 with embedded RP is a great solution for intra-AS interdomain multicasting. A single RP can be used for all domains, with the mapping function being performed by embedded RP. No additional protocols or interdomain configuration are required as the embedded RP for each group mapping propagates throughout the IPV6 network. However, this is not a very good solution for inter-AS interdomain multicast. The largest difficulty of such a design is the use of a single RP to service multiple ASs. For additional details, refer to RFC 3956.

Summary

This chapter reviews the fundamental requirements for interdomain forwarding of IP multicast flows. An understanding of PIM domains and how they are built on the three pillars of interdomain design is critical for architecting this type of forwarding. Remember that these are the three pillars:

- **The multicast control plane for source identification:** The router must know a proper path to any multicast source, either from the unicast RIB or learned (either statically or dynamically) through a specific RPF exception.
- **The multicast control plane for receiver identification:** The router must know about any legitimate receivers that have joined the group and where they are located in the network.
- **The downstream multicast control plane and MRIB:** The router must know when a source is actively sending packets for a given group. PIM–SM domains must also be able to build a shared tree from the local domain's RP, even when the source has registered to a remote RP in a different domain.

Multicast BGP, PIM, and MSDP satisfy the requirements of the three pillars. With these protocols, you should be able to configure any multidomain or interdomain network, including designs that are both internal and cross the public Internet. This chapter also reviews ways to eliminate the use of MSDP by using SSM or IPv6 embedded RP within the network.

References

RFC 3306

RFC 7371

RFC 5771

RFC 3956

RFC 4607

RFC 3446

RFC 3618

RFC 7606

RFC 4760

RFC 2283

RFC 1930

RFC 6996

This page intentionally left blank

Chapter 2

Multicast Scalability and Transport Diversification

Public cloud services are very commonly used in enterprise networks, and it is therefore important to understand how multicast messages are carried to and from a cloud service provider. Transportation of multicast messages requires consideration of several factors, especially when the cloud service providers do not support native multicast. This chapter introduces the key concepts of cloud services and explains the elements required to support multicast services.

Why Is Multicast Not Enabled Natively in a Public Cloud Environment?

Currently, public cloud environments do not have multicast enabled. This is the case because sending a packet to every host in the customer-owned cloud segment would have implications on the data plane and control plane traffic in terms of scalability of the cloud fabric. As enterprise customers add to the cloud more resources with a need for multicast, the cloud platform needs increased capacity and must meet other considerations for dynamic multicast demand. Such changes can be expensive, and speculative calculation can impact the user experience for the cloud consumers.

The calculation of changes needed becomes more complex and difficult to estimate when multitenancy is added to the equation. Allowing multicast protocol in the cloud fabric can have an impact on the performance of the cloud fabric. However, for enterprise customers that require multicast communication for business-relevant applications, it is difficult to adopt cloud services.

Enterprise Adoption of Cloud Services

Enterprise customers tend to adopt cloud services for a number of reasons, including agile provisioning, lower cost of investment, reduced operational and capital expense, closeness to geocentric user population, and speed of developing a product. Figure 2-1 illustrates the most common types of cloud models.