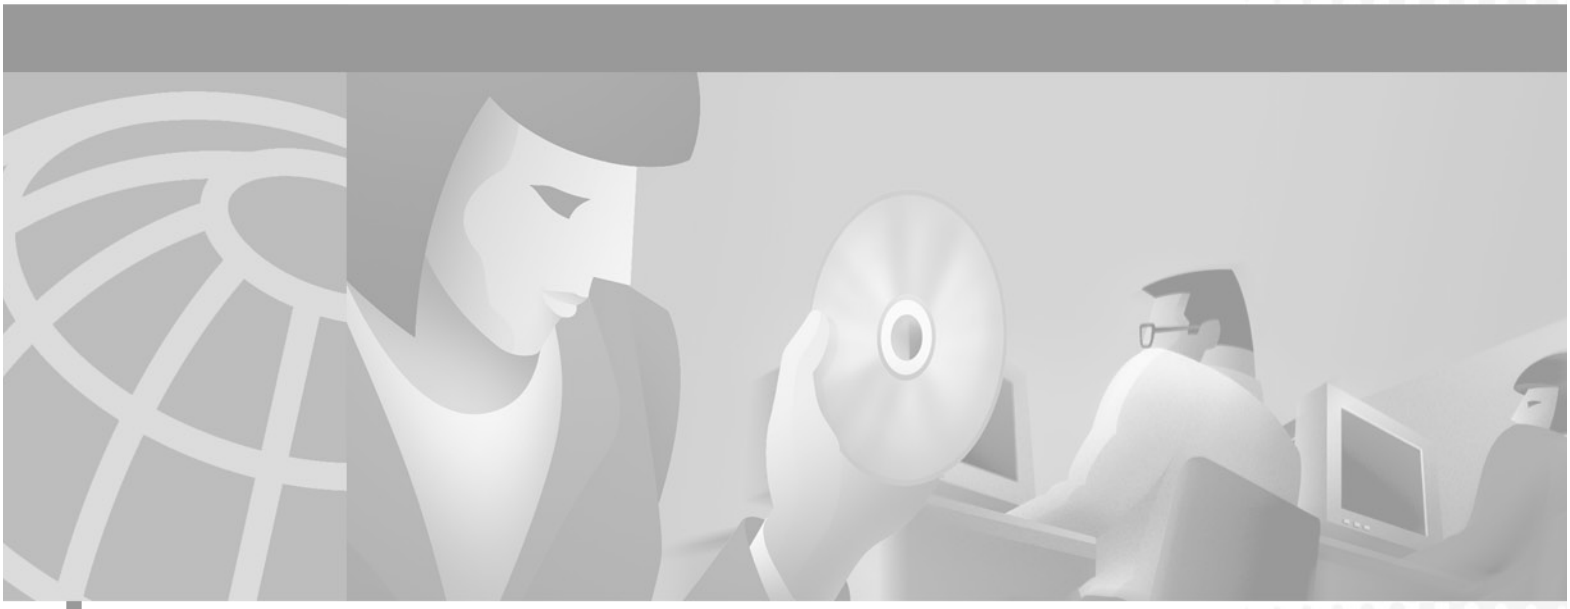

Cisco AVVID Network Infrastructure IP Multicast Design

Solutions Reference Network Design
March, 2003

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number: 956651

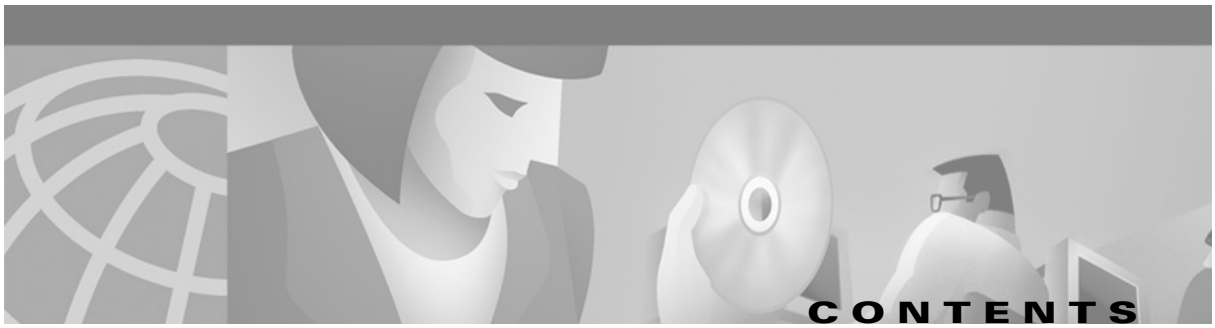

CONTENTS

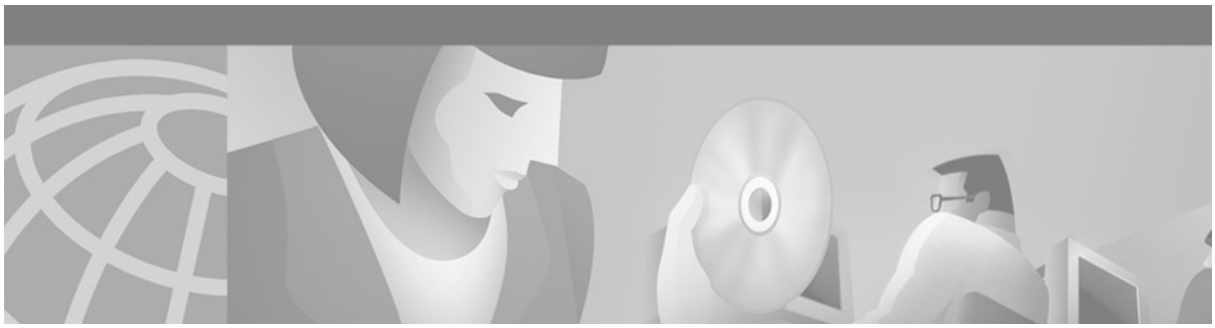

About this Document

This document presents an overview of AVVID IP multicast design and implementation.

Intended Audience

This document is intended for use by the Enterprise Systems Engineer (SE) or customer who may be unfamiliar with the deployment choices available to an AVVID Enterprise customer for IP multicast.

Document Organization

This document contains the following chapters:

Chapter or Appendix	Description
Chapter 1, “IP Multicast Overview”	Provides an overview of IP multicast design.
Chapter 2, “IP Multicast in a Campus Network”	Provides tips and recommendations for deploying IP multicast in a campus network.
Chapter 3, “IP Multicast in a Wireless LAN”	Provides tips and recommendations for deploying IP multicast in a wireless LAN.
Chapter 4, “IP Multicast in a Data Center”	Provides tips and recommendations for deploying IP multicast in a data center.
Chapter 5, “IP Multicast in a WAN”	Provides tips and recommendations for deploying IP multicast in a WAN.
Chapter 6, “IP Multicast in a Site-to-Site VPN”	Provides tips and recommendations for deploying IP multicast in a site-to-site VPN.
Chapter 7, “Multicast Music-on-Hold and IP/TV Configurations”	Provides the reference configurations for Multicast Music-on-Hold and IP/TV as used in the examples within the other chapters.
Chapter 8, “Security, Timers, and Traffic Engineering in IP Multicast Networks”	Provides recommendations for implementing security with IP multicast.
Chapter 9, “Managing IP Multicast”	Provides recommendations for managing IP multicast.

Note This document contains product and configuration information that is complete at the publish date. Subsequent product introductions may modify recommendations made in this document.

Document Conventions

This guide uses the following conventions to convey instructions and information:

Table 1 Document Conventions

Convention	Description
boldface font	Commands and keywords.
italic font	Variables for which you supply values.
[ ]	Keywords or arguments that appear within square brackets are optional.
{x | y | z}	A choice of required keywords appears in braces separated by vertical bars. You must select one.
`screen font`	Examples of information displayed on the screen.
`boldface screen font`	Examples of information you must enter.
< >	Nonprinting characters, for example passwords, appear in angle brackets.
[ ]	Default responses to system prompts appear in square brackets.

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Timesaver Means *the described action saves time*. You can save time by performing the action described in the paragraph.

Tips Means *the following information will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Obtaining Documentation

The following sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following URL:

http://www.cisco.com

Translated documentation is available at the following URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

  http://www.cisco.com/go/subscription

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

If you are reading Cisco product documentation on Cisco.com, you can submit technical comments electronically. Click the **Fax** or **Email** option under the "Leave Feedback" at the bottom of the Cisco Documentation home page.

You can e-mail your comments to bug-doc@cisco.com.

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Cisco Systems
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you to

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

You can self-register on Cisco.com to obtain customized information and service. To access Cisco.com, go to the following URL:

http://www.cisco.com

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available through the Cisco TAC: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Inquiries to Cisco TAC are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

Which Cisco TAC resource you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

The Cisco TAC Web Site allows you to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to the following URL:

http://www.cisco.com/tac

All customers, partners, and resellers who have a valid Cisco services contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to the following URL to register:

http://www.cisco.com/register/

If you cannot resolve your technical issues by using the Cisco TAC Web Site, and you are a Cisco.com registered user, you can open a case online by using the TAC Case Open tool at the following URL:

http://www.cisco.com/tac/caseopen

If you have Internet access, it is recommended that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses issues that are classified as priority level 1 or priority level 2; these classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer will automatically open a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to the following URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled; for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). In addition, please have available your service agreement number and your product serial number.

CHAPTER 1

IP Multicast Overview

IP multicast allows for a streamlined approach to data delivery whenever multiple hosts need to receive the same data at the same time. For example:

- When configured for IP multicast services, Music-on-Hold (MoH) can stream the same audio file to multiple IP phones without the overhead of duplicating that stream one time for each phone on hold.
- IP/TV allows for the streaming of audio, video, and slides to thousands of receivers simultaneously across the network. High-rate IP/TV streams that would normally congest a low-speed WAN link can be filtered to remain on the local campus network.

Multicast vs. Unicast

Multicast behaves differently than unicast. Unicast allows users to gain access to their "own" stream of data. The drawback to unicast is its inefficiency in distributing the same stream of data to multiple users. When a data stream from a single source is sent to many receivers using a unicast transmission, a load is created not only on the source but also on the network itself. The stream must be copied once for each user across the network. Figure 1-1 illustrates this difference.

Figure 1-1 Unicast vs. Multicast

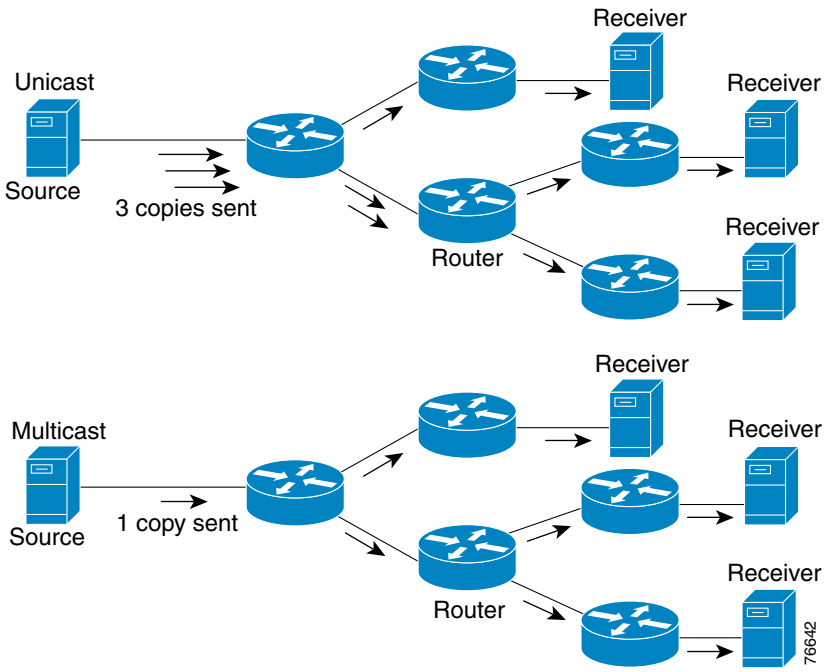

IP multicast traffic is UDP based and, as such, has some less-than-desirable characteristics. For example, it does not detect packet loss and, due to the lack of a windowing mechanism, it does not react to congestion. To compensate for this, applications and network devices can be configured to classify, queue, and provision multicast traffic using QoS. QoS virtually eliminates dropped packets and minimizes delay and delay variation for multicast streams. Thus, these limitations of IP multicast are not an issue.

Multicast MoH natively sends the audio music streams with a classification of DiffServ Code-Point equal to Expedited Forwarding (DSCP=EF). The classification values can be used to identify MoH traffic for preferential treatment when QoS policies are applied. Multicast streaming video is recommended to be classified at DSCP CS4.

Multicast Addressing

IP multicast uses the Class D range of IP addresses (224.0.0.0 through 239.255.255.255). Within the IP multicast Class D address range, there are a number of addresses reserved by the Internet Assigned Numbers Authority (IANA). These addresses are reserved for well-known multicast protocols and applications, such as routing protocol hellos.

For multicast addressing, there are generally two types of addresses as follows:

- Well known addresses designated by IANA
  - Packets using the following Reserved Link Local Addresses (also called the Local Network Control Block [224.0.0.0 - 224.0.0.255]) are sent throughout the local subnet only and are transmitted with TTL=1.

    Note The addresses listed below are just a few of the many addresses in the Link Local Address space.

    224.0.0.1—Sent to all systems on a subnet.

    224.0.0.2—Sent to all routers on a subnet.

    224.0.0.5—Sent to all OSPF routers.

    224.0.0.6—Sent to all OSPF DRs.

    224.0.0.9—Sent to all RIPv2 routers.

    224.0.0.10—Sent to all IGRP routers.

    224.0.0.13—Sent to all PIMv2 routers.

    224.0.0.22—Sent to all IGMPv3 devices.
  - Packets using the following Internetwork Control Block (224.0.1.0 - 224.0.1.255) addresses are also sent throughout the network.

    Note The addresses listed below are just a few of the many addresses in the Internetwork Control Block.

    224.0.1.39—Cisco-RP-Announce (Auto-RP)

    224.0.1.40— Cisco-RP-Discovery (Auto-RP)
- Administratively scoped addresses (239.0.0.0 - 239.255.255.255). For more information, see RFC 2365.

Tip For more information about multicast addresses, see http://www.iana.org/assignments/multicast-addresses.

Administratively-scoped addresses should be constrained to a local group or organization. They are used in a private address space and are not used for global Internet traffic. “Scoping” can be implemented to restrict groups with a given address or range from being forwarded to certain areas of the network.

Organization-local and site-local scopes are defined scopes that fall into the administratively scoped address range.

- Organization-local scope (239.192.0.0 - 239.251.255.255)—Regional or global applications that are used within a private enterprise network.
- Site-local scope (239.255.0.0 - 239.255.255.255)—Local applications that are isolated within a site/region and blocked on defined boundaries.

Scoping group addresses to applications allows for easy identification and control of each application.

The addressing used in this chapter reflects the organization-local scope and site-local scope ranges found in the administratively scoped address range.

For illustration purposes, the examples in this chapter implement IP/TV and MoH in an IP multicast environment. Table 1-1 lists the example address ranges used in these examples.

Table 1-1 Design Guide IP Multicast Address Assignment for Multicast Music-on-Hold and IP/TV

Application	Multicast Groups /22	Address Range	Scope	Notes
IP/TV High-Rate Traffic	239.255.0.0/16	239.255.0.0 - 239.255.255.255	Site-local	Restricted to local Campus
IP/TV Medium-Rate Traffic	239.192.248.0/22	239.192.248.0 - 239.192.251.255	Organization-local	Restricted to 768k+ Sites
IP/TV Low-Rate Traffic	239.192.244.0/22	239.192.244.0 - 239.192.247.255	Organization-local	Restricted to 256k+ Sites
Multicast Music-on-Hold	239.192.240.0/22	239.192.240.0 - 239.192.243.255	Organization-Local	No restrictions

The IP/TV streams have been separated based on the bandwidth consumption of each stream. IP/TV High-Rate traffic falls into the site-local scope (239.255.0.0/16) and is restricted to the local campus network. IP/TV Medium-Rate traffic falls into one range of the organization-local scope (239.192.248.0/22) and is restricted to sites with bandwidth of 768 Kbps or greater. IP/TV Low-Rate traffic falls into another range of the organization-local scope (239.192.244.0/22) and is restricted to sites with bandwidth of 256 Kbps or greater. Finally, multicast MoH traffic falls into yet another range of the organization-local scope (239.192.240.0/22) and has no restrictions.

This type of scoping allows multicast applications to be controlled through traffic engineering methods discussed later in this chapter.

Note The /22 networks were subnetted from the 239.192.240.0/20 range, allowing for four address classes. 239.192.252.0/22 can be used for additional applications not defined in this document.

Multicast Forwarding

IP multicast delivers source traffic to multiple receivers using the least amount of network resources as possible without placing additional burden on the source or the receivers. Multicast packets are replicated in the network by Cisco routers and switches enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols.

Multicast capable routers create “distribution trees” that control the path that IP Multicast traffic takes through the network in order to deliver traffic to all receivers. PIM uses any unicast routing protocol to build data distribution trees for multicast traffic. The two basic types of multicast distribution trees are source trees and shared trees.

- Source trees—The simplest form of a multicast distribution tree is a source tree with its root at the source and branches forming a tree through the network to the receivers. Because this tree uses the shortest path through the network, it is also referred to as a shortest path tree (SPT).
- Shared trees—Unlike source trees that have their root at the source, shared trees use a single common root placed at some chosen point in the network. This shared root is called a Rendezvous Point (RP).

PIM uses the concept of a designated router (DR). The DR is responsible for sending Internet Group Management Protocol (IGMP) Host-Query messages, PIM Register messages on behalf of sender hosts, and Join messages on behalf of member hosts.

PIM Dense Mode

PIM Dense Mode (PIM-DM) is a protocol that floods multicast packets to every PIM enabled interface on every router in the network. Because it is difficult to scale and has a propensity to stress network performance, dense mode is not optimal for most multicast applications and, therefore, not recommended.

PIM Sparse Mode

PIM Sparse Mode (PIM-SM) Version 2 is a more effective multicasting protocol than PIM-DM. PIM-SM assumes that no one on the network wants a multicast stream unless they request it via IGMP. In a PIM-SM environment, RPs act as matchmakers, matching sources to receivers. With PIM-SM, the tree is rooted at the RP not the source. When a match is established, the receiver joins the multicast distribution tree. Packets are replicated and sent down the multicast distribution tree toward the receivers.

Sparse mode's ability to replicate information at each branching transit path eliminates the need to flood router interfaces with unnecessary traffic or to clog the network with multiple copies of the same data. As a result, sparse mode is highly scalable across an enterprise network and is the multicast routing protocol of choice in the enterprise.

Note In a many-to-many deployment (many sources to many receivers), Bidir PIM is the recommended forwarding mode. Bidir PIM is outside the scope of this document. For more information on Bidir PIM, see the IP Multicast Technology Overview white paper located at: http://www.cisco.com/en/US/tech/tk648/tk363/technologies_white_paper09186a00800d6b5e.shtml

Resource Requirements

The memory impact on the router occurs when the router has to carry (*,G) state, which is the indication that a receiver has signaled an IGMP join, and (S,G), which is the indication that the Source is sending to the Group. The RP and any other router between the RP and the source are required to carry both state entries.

Note The default behavior of PIM-SM is to perform a SPT-switchover. By default, all routers will carry both states. The **spt-threshold infinity** command, described in Chapter 2, “IP Multicast in a Campus Network”, can be used to control the state.

When deciding which routers should be used as RPs, use the following to determine the memory impact on the router:

- Each (*,G) entry requires 380 bytes + outgoing interface list (OIL) overhead.
- Each (S,G) entry requires 220 bytes + outgoing interface list overhead.
- The outgoing interface list overhead is 150 bytes per OIL entry.

For example, if there are 10 groups with 6 sources per group and 3 outgoing interfaces:

of (*,G)s* x (380 + (*# of OIL entries* x 150)) = 10 x (380 + (3 x 150)) = 8300 bytes for (*,G)

***# of (S,G)s* x (220 + (*# of OIL entries* x 150)) = 60 x (220 + (3 x 150))= 40,200 bytes for (S,G)**

A total of 48,500 bytes of memory is required for the mroute table.

RP Deployment

There are several methods for deploying RPs.

- RPs can be deployed using a single, static RP. This method does not provide redundancy or load-balancing and is not recommended.
- Auto-RP is used to distribute group-to-RP mapping information and can be used alone or with Anycast RP. Auto-RP alone provides failover, but does not provide the fastest failover nor does it provide load-balancing.
- Anycast RP is used to define redundant and load-balanced RPs and can be used with static RP definitions or with Auto-RP. Anycast RP is the optimal choice as it provides the fast failover and load-balancing of the RPs.

Note In this document, the examples illustrate the most simplistic approach to Anycast RP by using locally-defined RP mappings.

Anycast RP

Anycast RP is the preferred deployment model as opposed to a single static RP deployment. It provides for fast failover of IP multicast (within milliseconds or in some cases seconds of IP Unicast routing) and allows for load-balancing.

In the PIM-SM model, multicast sources must be registered with their local RP. The router closest to a source performs the actual registration. Anycast RP provides load sharing and redundancy across RPs in PIM-SM networks. It allows two or more RPs to share the load for source registration and to act as hot backup routers for each other (multicast only). Multicast Source Discovery Protocol (MSDP) is the key protocol that makes Anycast RP possible. MSDP allows RPs to share information about active sources.

With Anycast RP, the RPs are configured to establish MSDP peering sessions using a TCP connection. When the RP learns about a new multicast source (through the normal PIM registration mechanism), the RP encapsulates the first data packet in a Source-Active (SA) message and sends the SA to all MSDP peers.

Two or more RPs are configured with the same IP address on loopback interfaces. The Anycast RP loopback address should be configured with a 32-bit mask, making it a host address. All the downstream routers are configured to “know” that the Anycast RP loopback address is the IP address of their RP. The non-RP routers will use the RP (host route) that is favored by the IP unicast route table. When an RP fails, IP routing converges and the other RP assumes the RP role for sources and receiver that were previously registered with the failed RP. New sources register and new receivers join with the remaining RP.

Auto-RP

Auto-RP automates the distribution of group-to-RP mappings in a network supporting PIM-SM. Auto-RP supports the use of multiple RPs within a network to serve different group ranges and allows configurations of redundant RPs for reliability purposes. Auto-RP allows only one RP to be active at once. Auto-RP can be used as the distribution mechanism to advertise the Anycast RP addresses previously discussed The automatic distribution of group-to-RP mappings simplifies configuration and guarantees consistency.

The Auto-RP mechanism operates using two basic components, the candidate RPs and the RP mapping agents.

- Candidate RPs advertise their willingness to be an RP via “RP-announcement” messages. These messages are periodically sent to a reserved well-known group 224.0.1.39 (CISCO-RP-ANNOUNCE).
- RP mapping agents join group 224.0.1.39 and map the RPs to the associated groups. The RP mapping agents advertise the authoritative RP-mappings to another well-known group address 224.0.1.40 (CISCO-RP-DISCOVERY). All PIM routers join 224.0.1.40 and store the RP-mappings in their private cache.

CHAPTER 2

IP Multicast in a Campus Network

This chapter discusses the basic layout needed to use IP multicast in a campus network and includes the following sections:

- Multicast Campus Deployment Recommendations
- Campus Deployment
- IP Multicast Small Campus Design
- IP Multicast Medium Campus Design
- IP Multicast Large Campus Design
- Summary

Note This chapter uses MoH and IP/TV in the examples. It does not, however, provide detailed configurations and designs for MoH and IP/TV. A basic MoH and IP/TV implementation is covered in Chapter 7, "Multicast Music-on-Hold and IP/TV Configurations."

Also, other types of IP multicast implementations, such as IP multicast for financial deployments, are not covered.

To get the most out of this chapter, the reader should understand the AVVID recommendations for the following:

- Campus design
- IP Telephony
- Content Delivery with IP/TV
- QoS
- High-Availability
- Security
- Management

Multicast Campus Deployment Recommendations

This chapter discusses the recommended and optional configurations for IP multicast campus deployment. The recommended guidelines are summarized below:

- Use IP multicast to scale streaming applications, such as MoH and IP/TV.
- Use administratively scoped addresses to differentiate multicast applications by type and bandwidth.
- Use Anycast RP when high availability and load balancing are needed.
- Understand and deploy the correct features to support filtering of non-RPF traffic in the hardware.
- Understand and correctly deploy HSRP when used with IP multicast deployment.
- Select Catalyst switches that have IGMP snooping and use CGMP in low-end switches that do not support IGMP snooping.
- Use recommended commands to ensure that the correct RPs and sources are used.
- Use IP multicast boundaries to control where certain multicast streams go.
- Use "show" commands to ensure proper operation of the multicast configurations and enable SNMP traps to log multicast events.

Campus Deployment

This section provides information for deploying the following IP multicast elements in a campus network:

- IGMP Snooping and CGMP
- Non-RPF Traffic
- HSRP

IGMP Snooping and CGMP

In addition to PIM, IP multicast uses the host signaling protocol IGMP to indicate that there are multicast receivers interested in multicast group traffic.

Internet Group Management Protocol (IGMP) snooping is a multicast constraining mechanism that runs on a Layer 2 LAN switch. IGMP snooping requires the LAN switch to examine some Layer 3 information (IGMP join/leave messages) in the IGMP packets sent between the hosts and the router. When the switch hears the "IGMP host report" message from a host for a multicast group, it adds the port number of the host to the associated multicast table entry. When the switch hears the "IGMP leave group" message from a host, the switch removes the host entry from the table.

Because IGMP control messages are sent as multicast packets, they are indistinguishable from multicast data at Layer 2. A switch running IGMP snooping must examine every multicast data packet to determine if it contains any pertinent IGMP control information. Catalyst switches that support IGMP snooping use special Application Specific Integrated Circuits (ASICs) that can perform the IGMP checks in hardware.

Optimal bandwidth management can be achieved on IGMP snooping enabled switches by enabling the IGMP Fast-Leave processing. With Fast-Leave, upon receiving an "IGMP leave group" message, the switch immediately removes the interface from its Layer 2 forwarding table entry for that multicast group. Without leave processing, the multicast group will remain in the Layer 2 forwarding table until the default IGMP timers expire and the entry is flushed.

The following example shows how to configure IGMP Fast-Leave on a Catalyst switch running Native IOS:

```
switch(config)#ip igmp snooping vlan 1 immediate-leave
```

The following example shows how to configure IGMP Fast-Leave on a Catalyst switch running Catalyst OS:

```
CatOS> (enable)set igmp fastleave enable
```

Use Fast-Leave processing *only* on VLANs where only one host is connected to each Layer 2 LAN interface. Otherwise, some multicast traffic might be dropped inadvertently. For example, if multiple hosts are attached to a Wireless LAN Access Point that connects to a VLAN where Leave processing is enabled (as shown in Figure 2-1), then Fast-Leave processing should **not** be used.

Figure 2-1 When Not to Use Fast-Leave Processing

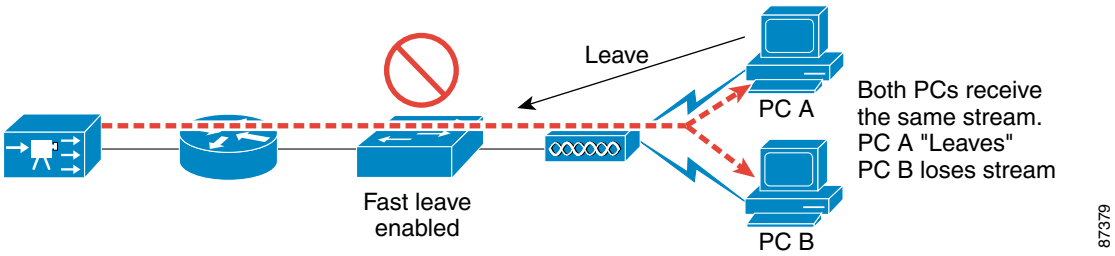

Cisco Group Management Protocol (CGMP) is a Cisco-developed protocol that allows Catalyst switches to leverage IGMP information on Cisco routers to make Layer 2 forwarding decisions. CGMP must be configured on the multicast routers and the Layer 2 switches. With CGMP, IP multicast traffic is delivered only to the Catalyst switch ports that are attached to interested receivers. All ports that have not explicitly requested the traffic will not receive it unless these ports are connected to a multicast router. Multicast router ports must receive every IP multicast data packet.

The default behavior of CGMP is to not remove multicast entries until an event, such as a spanning tree topology change, occurs or the router sends a CGMP leave message. The following example shows how to enable the CGMP client (switch) to act on actual IGMP leave messages:

```
switch(config)#cgmp leave-processing
```

Note Due to a conflict with HSRP, CGMP Leave processing is disabled by default.If HSRP hellos pass through a CGMP enabled switch, then refer to CSCdr59007 before enabling CGMP leave-processing.

Table 2-1 lists the support for CGMP and IGMP snooping in Cisco switches.

Table 2-1 Support for IGMP Snooping and/or CGMP

Model	CGMP Server	CGMP Client	IGMP Snooping
EtherSwitch Module	No	No	Yes
2950	No	No	Yes
3524-PWR	No	Yes	No
3550	Yes	No	Yes
4006/4500-SupIII/IV	Yes	Yes	Yes
6500-SupI/II	No-Switch, Yes - MSFC	No	Yes

Non-RPF Traffic

A router drops any multicast traffic received on a non-reverse path forwarding (non-RPF) interface. If there are two routers for a subnet, the DR will forward the traffic to the subnet and the non-DR will receive that traffic on its own VLAN interface. This will not be its shortest path back to the source and so the traffic will fail the RPF check. How non-RPF traffic is handled depends on the Catalyst switch platform and the version of software running (as shown in Figure 2-2).

Figure 2-2 Handling of Non-RPF Traffic

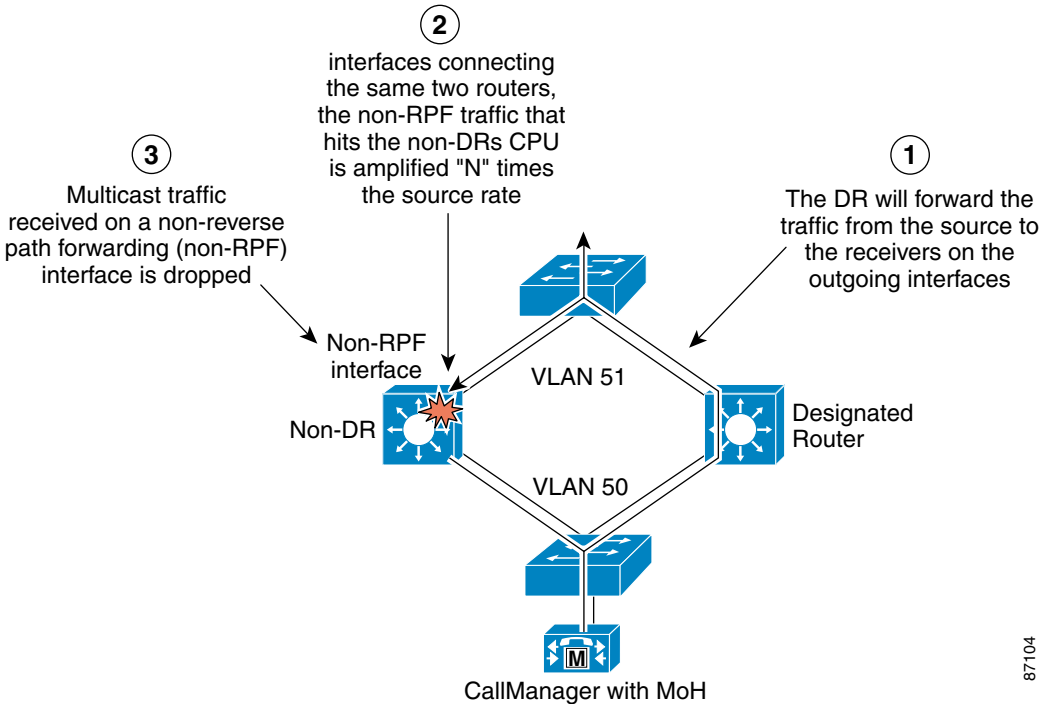

Catalyst 6500 Series

Without a method to control non-RPF traffic in hardware, the CPU on the supervisor engine will reach 99%. The Supervisor I requires a RACL on the non-DR (only). The RACL must allow only locally sourced multicast traffic on the VLAN interface. The **no ip unreachables** command must also be configured on the interface. For RACL denied traffic, if the interface does not have the **no ip unreachables** command configured, the ACL denied traffic is leaked to the MSFC at 10 pps per VLAN.

The following example shows how to enable manual RACL configuration for blocking non-RPF traffic on the non-DR router:

```
interface VLAN X
ip access-group 100 in
no ip unreachables

access-list 100 permit ip w.x.y.z 0.0.0.255 any          Local subnet addresses
access-list 100 permit ip any 224.0.0.0 0.0.0.255
access-list 100 permit ip any 224.0.1.0 0.0.0.255
access-list 100 deny   ip any 224.0.0.0 15.255.255.255
```

Versions of code for the Catalyst 6500 series Supervisor II (Catalyst OS 6.2(1) and IOS 12.1(5)EX and later) support Multicast Fast Drop (MFD) and will rate-limit non-RPF traffic by default. The **mls ip multicast non-rpf cef** command is enabled by default on the MSFC. Use the **show mls ip multicast summary** command to verify that non-RPF traffic is being rate-limited in hardware.

Tip For more information, see http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/ios121_e/swcg/mcastmls.htm#xtocid2827011.

Catalyst 4006 and 4500 with Supervisor III/IV

By default, the Supervisor III and IV enable MFD and handle non-RPF traffic in hardware. The **ip mfib fastdrop** command is used to enable this feature.

To display the MFIB table information, use the **show ip mfib** command. To display the information in the hardware tables, use the **show platform hardware** command.

Catalyst 3550

The Catalyst 3550 does not use a command to enable non-RPF traffic to be hardware filtered. In 3550 switches, the non-RPF packets reach the CPU through the RPF Failure Queue. This hardware queue is separate from other queues, which are reserved for routing protocols and Spanning-Tree protocol packets. Thus, non-RPF packets will not interfere with these critical packets. The RPF Failure Queue is of minimal depth. So if this queue is full, then subsequent packets will be dropped by the hardware itself. The CPU gives low priority and shorter time to process the packets from the RPF Failure Queue to ensure that priority is given to routing protocol packets. A limited number of packet buffers are available for the RPF Failure Queue and when all of the buffers allocated for the RPF Failure Queue are full, the software will drop the incoming packets. If the rate of the non-RPF packets is still high and if this in turn makes the software process a lot of packets within a certain period of time, then the queue is disabled and re-enabled after 50 milliseconds. This will flush the queue and give the CPU a chance to process the existing packets

To see the number of packets dropped, use the **show controller cpu-interface | include rpf** command.

HSRP

There is often some confusion about how HSRP active and standby devices route IP Unicast and IP multicast traffic. The issue appears when an HSRP active device is successfully routing IP Unicast traffic and the HSRP standby device is forwarding the IP multicast traffic. IP multicast forwarding is not influenced by which box is active or in standby.

The path that multicast traffic takes is based on the shortest path to the source or, in a shared-tree model, the shortest path to the RP. The shortest path to the source or RP is based on the route entries in the unicast routing table. In most Campus designs, the links between layers are equal cost paths. Therefore, the multicast traffic will follow the path through the DR. The DR is determined by which PIM router has the highest IP address on the shared subnet and also which has an RPF interface toward the source.

If multiple paths exist and they are **not** equal, it is possible for the DR to decide that the shortest path to the source or RP is actually back out the same VLAN that the host is on and through the non-DR router.

Figure 2-3 illustrates the possible issue with HSRP and IP multicast. In this example, it is assumed that the routes are equal-cost. Switch 1 is configured with an HSRP priority of 120 and Switch 2 is configured with a priority of 110. Therefore, HSRP will use Switch 1 as the active router. However, Switch 2 has a higher IP address than Switch 1. Therefore, PIM will use Switch 2 as the DR. The result is that Unicast traffic is sent through Switch 1 while Multicast traffic is sent through Switch 2.

Figure 2-3 HSRP and IP Multicast

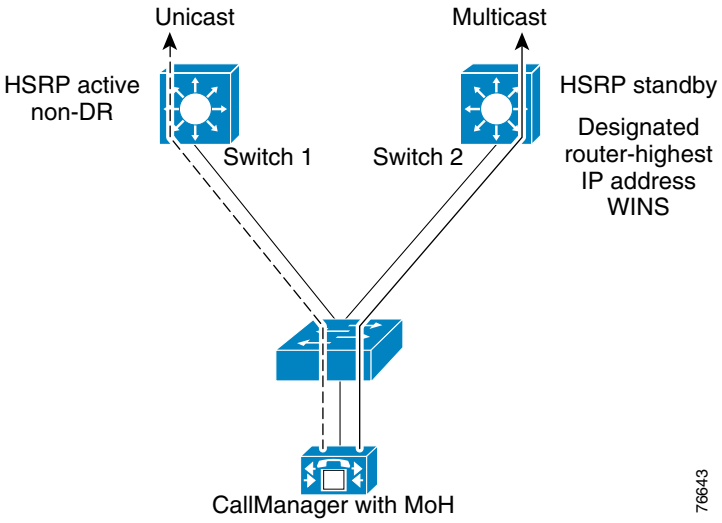

Solution

To avoid this situation, either adjust the HSRP priority to make Switch 2 the active router or change the IP addresses so that Switch 1 has the higher address. Either of these actions will make the HSRP active router and the PIM DR the same.

Note In IOS Release 12.2, the **ip pim dr-priority** command forces a router to be the DR (regardless of IP address). Not all platforms support the **ip pim dr-priority** command.

Figure 2-4 shows how changing the HSRP priority value on both switches will cause both unicast and multicast traffic will flow the same way.

Figure 2-4 HSRP and IP Multicast—HSRP Priority Changed

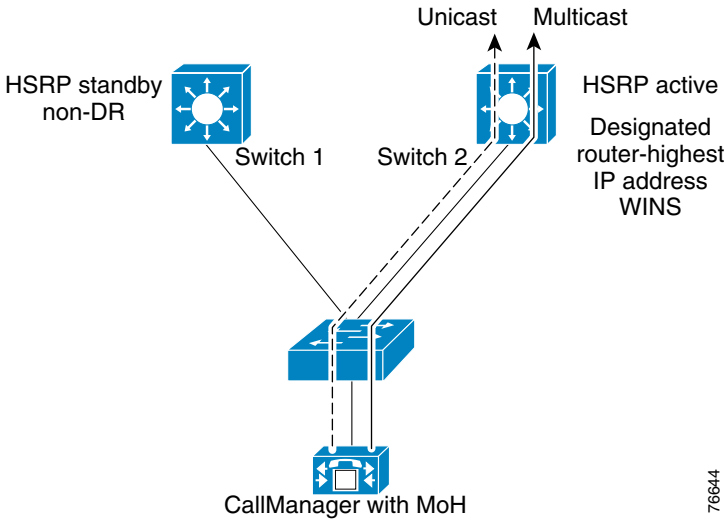

The following example shows the configuration of Switch 2:

```
interface Vlan10
 description To Server-Farm MoH Server
 ip address 10.5.10.3 255.255.255.0
 ip pim sparse-mode
 standby 10 priority 120 preempt
 standby 10 ip 10.5.10.1
```

The following example shows how to verify that the switch is the HSRP active device:

```
switch2#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active addr     Standby addr    Group addr
Vl10        10  120  P Active   local           10.5.10.3        10.5.10.1
```

The following example shows how to verify that the switch is the PIM DR for the subnet:

```
switch2#show ip pim interface vlan10

Address          Interface                Version/Mode     Nbr    Query     DR
                                                           Count  Intvl
10.5.10.3        Vlan10                   v2/Sparse         1     5       10.5.10.3
```

Trace routes should show that unicast traffic is flowing through Switch 2. The **mtrace** and **show ip multicast active** commands should show that the multicast traffic is also flowing through Switch 2.

Note In some situations, there may be a desire to load-balance over the two links from the access layer to the distribution layer. If this is the case, simply ensure that the HSRP Active router is **not** the DR for the VLAN. For this configuration, which is counter to the one recommended above, give the HSRP Active router the *lower IP* address of the two distribution switches to ensure that it is selected as the DR.

RP of Last Resort

If the active RPs are no longer available or there are no RPs configured for a specific group, the default behavior is to dense-mode flood the multicast traffic. This is called *dense-mode fallback*. Typically, after deploying the recommendations in this document, an RP will always be available. However, in the event that something happens to all of the RPs or routing instability diverts access from the RPs, it is necessary to ensure that dense-mode fallback does not occur.

To prevent dense mode flooding, on every PIM-enabled router configure an access control list (ACL) and use the **ip pim rp-address** *address group-list* command to specify an "RP of last resort." This way, if all else fails, the RP defined in this command will be used. It is recommended that a local loopback interface be configured on each router and that the address of this loopback interface be specified as the IP address of the RP of last resort. By configuring an RP of last resort, the local router will be aware of an RP and will not fallback to dense mode.

Note Do not advertise this loopback interface in the unicast routing table.

Example 2-1 RP of Last Resort (configured on every PIM-enabled router)

```
interface Loopback2
 description Garbage-CAN RP
 ip address 2.2.2.2 255.255.255.255

ip pim rp-address 2.2.2.2 1

access-list 1 deny 224.0.1.39
access-list 1 deny 224.0.1.40
access-list 1 permit any
```

This not only helps with dense-mode fallback (by ensuring that an RP is always present on the local router if the main RPs become unavailable), but this also helps to guard against rogue sources that may stream unwanted multicast traffic. This blocking of unwanted multicast sources is sometimes referred to as a "Garbage-can RP." For more information about using Garbage-can RPs, see Chapter 8, "Security, Timers, and Traffic Engineering in IP Multicast Networks."

IP Multicast Small Campus Design

This section provides a sample design for IP multicast in a small campus network. In this design, there are Layer 3 interfaces on the backbone switches for HSRP. As shown in Figure 2-5, there is only one building. The VLANs are identified as either data VLANs or voice VLANs. HSRP is configured on both backbone switches and on their links connecting to the access switches. The addressing scheme for HSRP used in this design is 10.*bulding_number.VLAN_number.role*.

For example, the address layout for VLAN 2/building 1 is:

- 10.1.2.1—HSRP address (Default Gateway)
- 10.1.2.2—IP address of standby router (4kS3-left-BB)
- 10.1.2.3—IP address of active router (4kS3-right-BB).

All hosts on VLAN 2 use 10.1.2.1 as their default gateway.

Note This type of addressing plan is used throughout this chapter.

Figure 2-5 Small Campus Design Reference Diagram

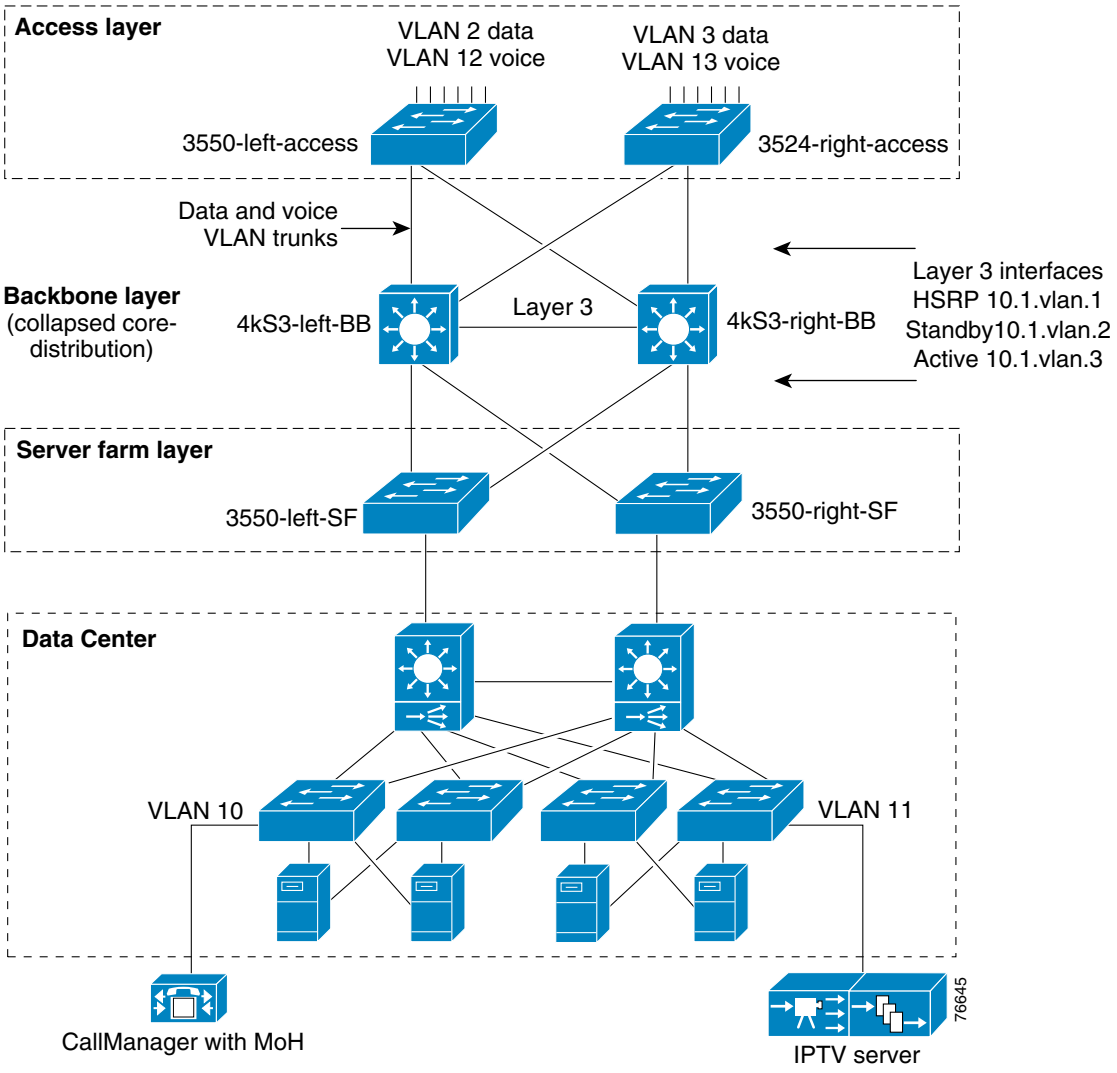

In this design:

- Both IGMP snooping and CGMP are used. There is a Catalyst 3550 (IGMP snooping) and a Catalyst 3524-PWR (CGMP) in the access layer.
- There are two Catalyst 4006 switches with Supervisor III in the collapsed core/distribution layer.
- The access-layer switches rely on the backbone switches as the RPs for PIM-SM.

Based on the size of the network and the number of sources and receivers that will be active on the network, there is no need for an advanced multicast deployment model. PIM-SIM is used on the two backbone switches to provide multicast forwarding between the Layer 3 links/VLANs in this design. Auto-RP is used to distribute group-to-RP mapping information. Additionally, simple security features are implemented to secure the network from unauthorized sources and RPs.

Core/Distribution-Layer Switch Configuration

The following example shows the configuration of "4kS3-left-BB." This switch is configured for Auto-RP and is the HSRP standby device.

```
ip multicast-routing                                  Enables IP multicast routing globally.
interface Loopback0                                   Creates interface used for RP and Mapping Agent operation.
 description Interface used for RP/MA
 ip address 1.1.1.1 255.255.255.255
 ip pim sparse-dense-mode
!
interface Vlan2
 description Data VLAN 2
 ip address 10.1.2.2 255.255.255.0
 ip pim sparse-dense-mode                             Enables sparse-dense mode on each interface.
 standby 2 ip 10.1.2.1
 standby 2 priority 110 preempt                       Because priority 110 is less than 120 (on the right switch),
!                                                     this device will be the HSRP standby.
!
interface Vlan3                                       VLAN 3 connects to 3524-right-access, which supports only
 description Data VLAN 3                              CGMP.
 ip address 10.1.3.2 255.255.255.0
 ip pim sparse-dense-mode
 standby 3 ip 10.1.3.1
 standby 3 priority 110 preempt
 ip cgmp                                              Enables CGMP Server.
!
interface Vlan10
 description Server Farm VLAN 10 - MOH
 ip address 10.1.10.2 255.255.255.0
 ip pim sparse-dense-mode
 standby 10 ip 10.1.10.1
 standby 10 priority 110 preempt
!
interface Vlan11
 description Server Farm VLAN 11 - IP/TV
 ip address 10.1.11.2 255.255.255.0
 ip pim sparse-dense-mode
 standby 11 ip 10.1.11.1
 standby 11 priority 110 preempt
!
interface Vlan12
 description Voice VLAN 12
 ip address 10.1.12.2 255.255.255.0
 ip pim sparse-dense-mode
 standby 12 ip 10.1.12.1
 standby 12 priority 110 preempt
!
interface Vlan13                                      VLAN 13 connects to 3524-right-access, which supports
 description Voice VLAN 13                            only CGMP.
 ip address 10.1.13.2 255.255.255.0
 ip pim sparse-dense-mode
 standby 13 10.1.13.1
 standby 13 priority 110 preempt
 ip cgmp                                              Enables CGMP Server.
!
ip pim send-rp-announce Loopback0 scope 16 group-list 1   Sends an Auto-RP announcement
!                                                         message to the 224.0.1.39 group.
ip pim send-rp-discovery Loopback0 scope 16               Sends an Auto-RP announcement
!                                                         message to the 224.0.1.40 group.
```

```
access-list 1 permit 239.192.240.0 0.0.3.255
access-list 1 permit 239.192.244.0 0.0.3.255
access-list 1 permit 239.192.248.0 0.0.3.255
access-list 1 permit 239.255.0.0 0.0.3.255
```

The following example shows the configuration of "4kS3-right-BB." This switch is configured for Auto-RP and is the HSRP active device.

```
ip multicast-routing                          Enables IP multicast routing globally.
!
interface Loopback0                           Creates an interface used for RP and Mapping Agent
 description Interface used for RP/MA         operation.
 ip address 1.1.1.2 255.255.255.255
 ip pim sparse-dense-mode
!
interface Vlan2
 description Data VLAN 2
 ip address 10.1.2.3 255.255.255.0
 ip pim sparse-dense-mode                     Enables sparse-dense mode on each interface.
 standby 2 ip 10.1.2.1
 standby 2 priority 120 preempt               Because priority 120 is more than 110 (on the left
!                                             switch), this device will be the HSRP active.
!
interface Vlan3                               VLAN 3 connects to 3524-right-access, which supports
 description Data VLAN 3                      only CGMP.
 ip address 10.1.3.3 255.255.255.0
 ip pim sparse-dense-mode
 standby 3 ip 10.1.3.1
 standby 3 priority 120 preempt
 ip cgmp                                      Enables CGMP Server.
!
interface Vlan10
 description Server Farm VLAN 10 - MOH
 ip address 10.1.10.3 255.255.255.0
 ip pim sparse-dense-mode
 standby 10 ip 10.1.10.1
 standby 10 priority 120 preempt
!
interface Vlan11
 description Server Farm VLAN 11 - IP/TV
 ip address 10.1.11.3 255.255.255.0
 ip pim sparse-dense-mode
 standby 11 ip 10.1.11.1
 standby 11 priority 120 preempt
!
interface Vlan12
 description Voice VLAN 12
 ip address 10.1.12.3 255.255.255.0
 ip pim sparse-dense-mode
 standby 12 ip 10.1.12.1
 standby 12 priority 120 preempt
!
interface Vlan13                              VLAN 13 connects to 3524-right-access, which supports
 description Voice VLAN 13                    only CGMP.
 ip address 10.1.13.3 255.255.255.0
 ip pim sparse-dense-mode
 standby 13 10.1.13.1
 standby 13 priority 120 preempt
 ip cgmp                                      Enables CGMP Server.
!
```

```
ip pim send-rp-announce Loopback0 scope 16 group-list 1
!
ip pim send-rp-discovery Loopback0 scope 16
!
!
access-list 1 permit 239.192.240.0 0.0.3.255
access-list 1 permit 239.192.244.0 0.0.3.255
access-list 1 permit 239.192.248.0 0.0.3.255
access-list 1 permit 239.255.0.0 0.0.3.255
```

Sends an Auto-RP announcement message to the 224.0.1.39 group.
Sends an Auto-RP announcement message to the 224.0.1.40 group.

The following examples show how to verify that the access layer switches have found their attached multicast routers. On 3550-left-access, display the IGMP snooping multicast router information:

```
3550-left-access#show ip igmp snooping mrouter
Vlan    ports
----    -----
2       Gi0/1(dynamic                 Gi0/1 connects to "4kS3-right-BB"
12      Gi0/1(dynamic)
2       Gi0/2(dynamic                 Gi0/2 connects to "4kS3-left-BB"
12      Gi0/2(dynamic)
```

On 3524-right-access, display the CGMP Server and Client operation:

```
3524-right-access#show cgmp
CGMP is running.
CGMP Fast Leave is not running.
CGMP Allow reserved address to join GDA .
Default router timeout is 300 sec.

vLAN     IGMP MAC Address   Interfaces
------  -----------------   -----------

vLAN     IGMP Router        Expire   Interface
------  -----------------  --------  ----------
    3    0010.7bab.983f     281 sec   Gi0/1       Gi0/1 connects to "4kS3-right-BB"
   13    0010.7bab.983f     281 sec   Gi0/1
    3    0010.7bab.983e     281 sec   Gi0/2       Gi0/2 connects to "4kS3-left-BB"
   13    0010.7bab.983e     281 sec   Gi0/2
```

Once Auto-RP multicast traffic is flowing to the switches, check to see that multicast group entries show up in their CAM tables.

```
3550-left-access#show mac-address-table multicast
Vlan    Mac Address       Type       Ports
----    -----------       ----       -----
   2    0100.5e00.0002    IGMP       Gi0/1, Gi0/2
   2    0100.5e00.0127    IGMP       Gi0/1, Gi0/2
   2    0100.5e00.0128    IGMP       Gi0/1, Gi0/2
  12    0100.5e00.0002    IGMP       Gi0/1, Gi0/2
  12    0100.5e00.0127    IGMP       Gi0/1, Gi0/2
  12    0100.5e00.0128    IGMP       Gi0/1, Gi0/2
```

Because the access-layer switches are Layer 2 switches, they will display Layer 2 Multicast address information instead of Layer 3 IP multicast addresses. The two key group addresses listed are:

- 0100.5e00.0127—RP announcement address of 224.0.1.39
- 0100.5e00.0128—Mapping agent discovery address of 224.0.1.40

Note For recommendations and configurations for securing the network from unauthorized sources and RPs, see Chapter 8, "Security, Timers, and Traffic Engineering in IP Multicast Networks.".

IP Multicast Medium Campus Design

This section provides a sample design for IP multicast in a medium campus network.

Note Unlike the small campus design, which used collapsed distribution and core layer, the medium campus network has a distinct distribution and core layer. The IP addressing, VLAN numbering, and HSRP numbering conventions are the same as those used in the small campus design.

For IP multicast in a medium campus (shown in Figure 2-6):

- Use IGMP snooping or CGMP enabled access switches.
- Enable PIM-SM on the distribution layer switches with their RPs defined as the Anycast RP address of the core switches.
- Reduce multicast state (S,G) from the leaf routers by keeping traffic on the shared tree. (Optional)
- Place the RPs in the core layer switches running PIM-SM and Anycast RP.

Figure 2-6 Medium Campus Design Reference Diagram

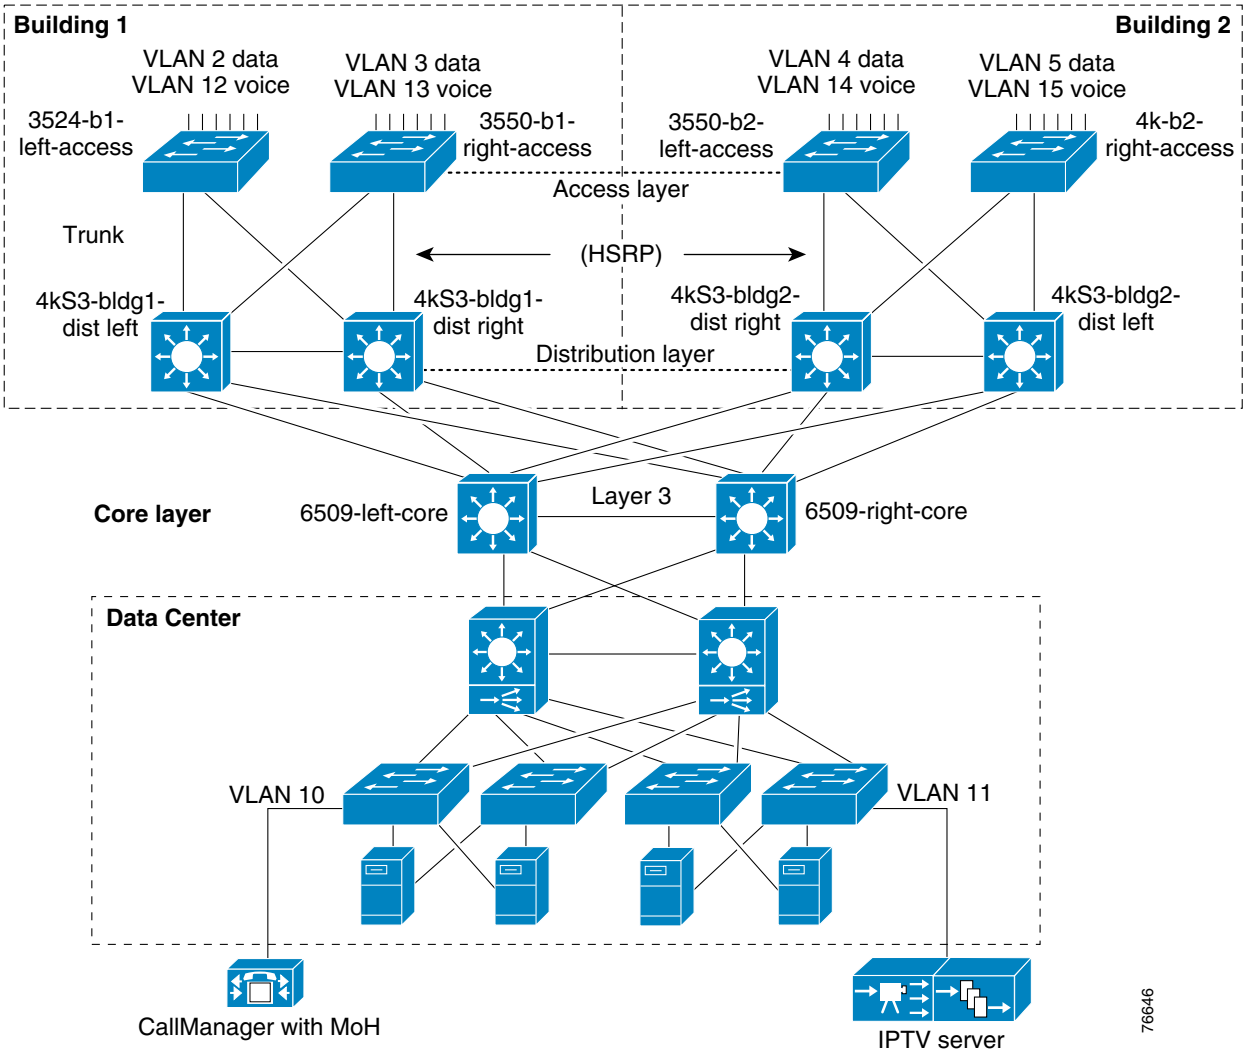

In this design:

- The access layer uses switches that support both IGMP snooping (Catalyst 3550 and Catalyst 4006 with Supervisor III) and CGMP Client (Catalyst 3524-PWR).
- The RPs are located on the two core layer switches (Catalyst 6509-left-core and Catalyst 6509-right-core).
- PIM-SM is configured on all distribution switches and core switches.
- Anycast RP is configured for fast recovery of IP multicast traffic.
- PIM-SM and MSDP are enabled on both core switches.
- Each distribution switch points to the Anycast RP address as its RP.
- MSDP is used to synchronize SA states between both core switches.

Core-Layer Switch Configuration

The following example shows the configuration of "6509-left-core." This switch is configured for Anycast RP and is the HSRP standby device.

```
ip multicast-routing
!
interface Loopback0
 description MSDP local peer address
 ip address 10.6.1.1 255.255.255.255
!
interface Loopback1
 description Anycast RP address
 ip address 10.6.2.1 255.255.255.255          Loopback 1 has same address on both the right and left
!                                              routers.
!
interface Vlanxx
 description VLANs for other L3 links
 ip address 10.0.0.x 255.255.255.252
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1                     Identifies the address of the RP.
!
ip msdp peer 10.6.1.2 connect-source Loopback0 Enables MSDP and identifies the address of the
!                                              MSDP peer.
!
ip msdp cache-sa-state                         Creates SA state (S,G). The SA cache entry is created
!                                              when either MSDP peer has an active source.
!
ip msdp originator-id Loopback0                Sets RP address in SA messages to be Loopback 0.
```

The following example shows the configuration of "6509-right-core." This switch is configured for Anycast RP and is the HSRP active device.

```
ip multicast-routing
!
interface Loopback0
 description MSDP local peer address
 ip address 10.6.1.2 255.255.255.255
!
interface Loopback1
 description Anycast RP address
 ip address 10.6.2.1 255.255.255.255          Loopback 1 has same address on both the right and left
!                                              routers.
!
interface Vlanxx
 description VLANs for other L3 links
 ip address 10.0.0.x 255.255.255.252
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1                     Identifies the address of the RP.
!
ip msdp peer 10.6.1.1 connect-source Loopback0 Enables MSDP and identifies the address of the
!                                              MSDP peer. The connect-source identifies the
!                                              primary interface used to source the TCP connection
!                                              between peers.
!
```

```
ip msdp cache-sa-state                    Creates SA state (S,G). The SA cache entry is created
!                                         when either MSDP peer has an active source.
!
ip msdp originator-id Loopback0           Sets RP address in SA messages to be Loopback 0.
```

Distribution-Layer Switch Configuration

The distribution switches for both building 1 and 2 run PIM-SM on each Layer 3 interface and point to the Anycast RP of 10.6.2.1 (Loopback 1 on both RPs).

The following example shows the multicast configuration for each distribution switch.

```
ip multicast-routing                      Enables IP multicast routing globally.
!
interface VLAN/GigabitEthernet
 ip pim sparse-mode                       Enables PIM-SM on the interface.
 ip cgmp                                  Enables CGMP.
!
interface Vlanxx
 description VLANs for other L3 links to Core
 ip address 10.0.0.x 255.255.255.252
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1                Points to the Loopback 1 address of both core switches.
!
ip pim spt-threshold infinity             Reduces multicast state (S,G) from the leaf routers by
!                                         keeping traffic on the shared tree. (Optional)
!
!To simplify this example, the HSRP configuration has been omitted.
```

Note If the Layer 3 interface connects to a CGMP-enabled switch (3524-PWR), CGMP Server operation must be enabled. If the interface connects to an IGMP snooping enabled router, IP CGMP does not need to be enabled, but PIM must be enabled.

The following example shows how to verify that each distribution switch has established a PIM neighbor relationship.

```
4kS3-bldg1-dist-left#show ip pim neighbor
PIM Neighbor Table
Neighbor Address  Interface               Uptime    Expires   Ver  Mode
10.1.2.3          Vlan2                   2d14h     00:01:27  v2    (DR)
10.1.12.3         Vlan12                  2d14h     00:01:37  v2    (DR)
10.1.3.3          Vlan3                   1w1d      00:01:16  v2    (DR)
10.1.13.3         Vlan13                  1w1d      00:00:16  v2    (DR)
10.0.0.2          GigabitEthernet0/1      00:50:59  00:00:14  v2    (DR)
10.0.0.6          GigabitEthernet0/2      00:50:59  00:00:14  v2    (DR)
```

Note For information about configuring the Data Center portion of the network for IP multicast, see Chapter 4, "IP Multicast in a Data Center."

IP Multicast Large Campus Design

This section provides a sample design for IP multicast in a large campus network. Multicast design in a large enterprise network can be difficult if the design is too granular. The optimal design provides fast-failover and a simplistic approach to traffic control. Although there are a large number of possible combinations in deploying multicast in a large campus and even more combinations for each type of multicast application, the sample design in this chapter focuses on keeping things simple and the traffic reliable.

Note The large campus network also has a distinct distribution and core layer. The IP addressing, VLAN numbering, and HSRP numbering conventions are the same as those used in the small campus design.

For IP multicast in a large campus (shown in Figure 2-7):

- Use IGMP snooping enabled access-layer switches.
- Enable PIM-SM on the distribution-layer switches that have their RPs defined as the Anycast RP address of the two RPs located on 6k-botleft-core and 6k-botright-core.
- Reduce multicast state (S,G) from the leaf routers by keeping traffic on the shared tree. (Optional)
- Place the RPs in the core-layer switches. Run PIM-SM and Anycast RP. A minimum of two RPs must be used to deploy Anycast RP.

Figure 2-7 Large Campus Design Reference Diagram

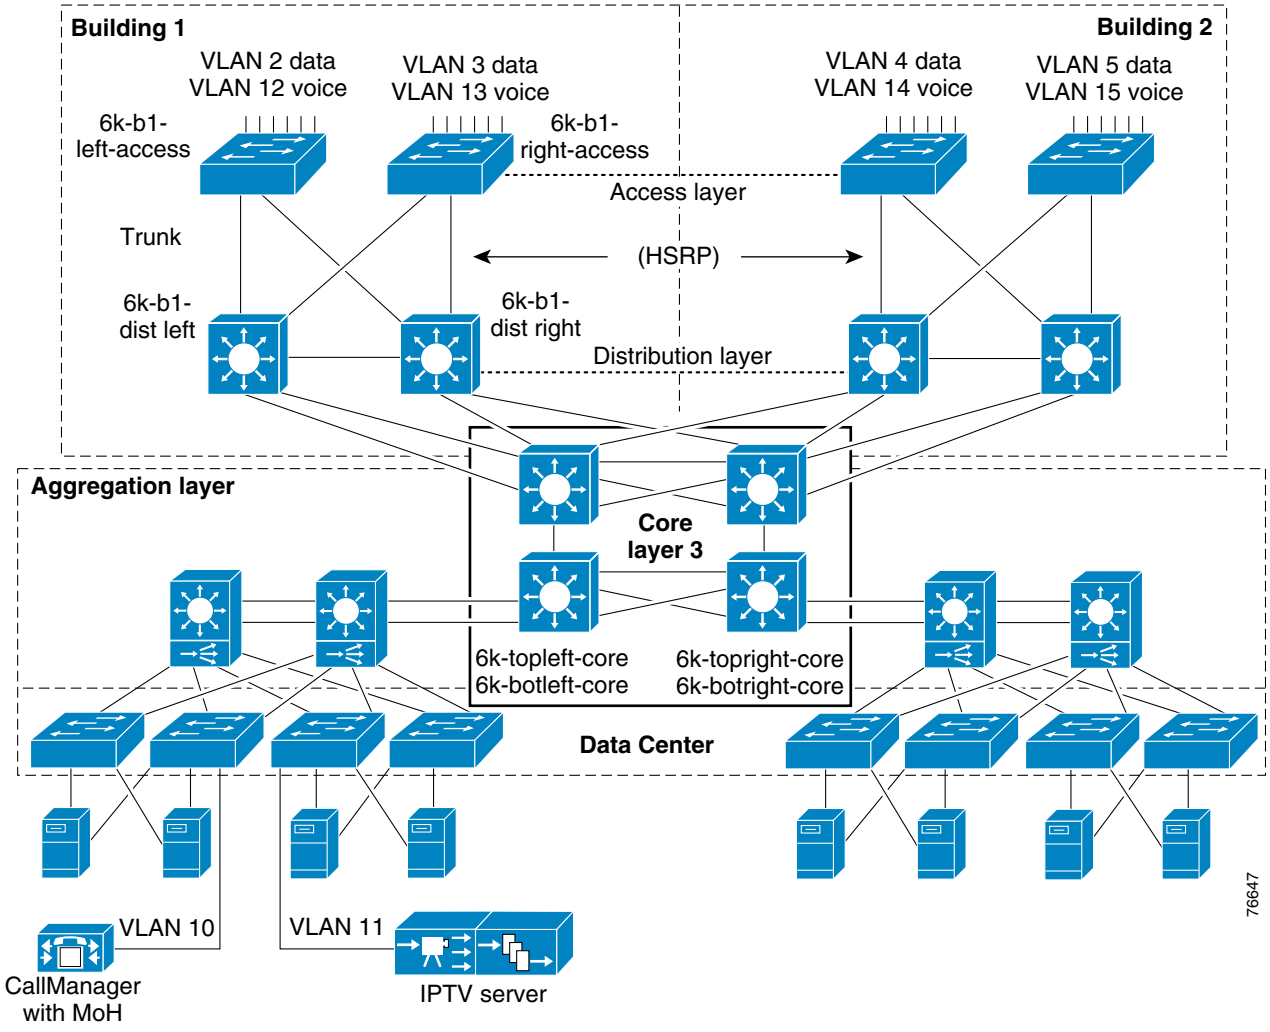

Looking at this design layer-by-layer:

- The access layer uses switches that support IGMP snooping and good port density to serve the end stations. The Catalyst 6500 with Supervisor II and 4006 with Supervisor III are used in this design.
- The distribution layer uses switches that have good multicast forwarding performance, the capability to have large multicast routing tables, and the ability to house multiple Gigabit links. The Catalyst 6500 with Supervisor II is used in both the distribution layer and the server farm aggregation layer.
- The server farm uses switches that support IGMP snooping and advanced features that are useful for security, QoS, and management. The Catalyst 6500 with Supervisor II and 4006 with Supervisor III are used in the server farm.
- The core layer uses switches that support the same features that the distribution-layer switches support. In addition, the core-layer switches must support a dense population of Gigabit ports for connectivity to the distribution-layer switches and other core-layer switches. The Catalyst 6500 with Supervisor II is used in the core layer.

Each client and server access-layer switch is dual-connected to, and served by, a pair of distribution-layer routers running HSRP. For multicast, one of the two routers is the DR with the other being the IGMP querier. The IP Unicast routing protocol is configured such that the trunk from the access-layer switch to the DR is always preferred over that of the non-DR. This forces the unicast and multicast paths to be the same. The IGMP querier is responsible for sending IGMP queries. Both the DR and the non-DR receive the subsequent reports from clients, but only the DR will act on them. If the DR fails, the non-DR will take its role. If the IGMP querier fails, the DR will take over its role. The distribution routers have dual connections to the core.

Keep the RP placement simple. With Anycast RP, it is recommended that the RPs be placed in the center of the network. Placing RP operations on the core-layer switches is a good idea because the core is the central point in the network servicing the distribution-layer switches in each building, the aggregation-layer switches in the server farms, and the WAN and Internet blocks.

The applications used in this sample design (MoH and IP/TV) have few sources to many receivers. The sources are located in the server farm. So, a complex distribution of RPs throughout the network is not required.

Note Due to the number of devices in a large campus design, this section presents configurations for a sampling of the devices. The remaining devices should have similar configurations with the exception of the unique IP addresses, VLANs, HSRP, and other specifics.

In this design:

- The access layer switches have IGMP snooping enabled.
- The RPs are located on the two core-layer switches.
- PIM-SM is configured on all distribution-layer switches and core-layer switches.
- Anycast RP is configured for fast recovery of IP multicast traffic.
- PIM-SM and MSDP are enabled on all core-layer switches.
- Each distribution-layer switch points to the Anycast RP address as their RP.
- MSDP is used to synchronize SA state between the core switches.

Note Only two RPs are required to run Anycast RP. In most situations, two RPs will sufficiently provide redundancy for multicast. The following sections show the RP configuration for “6k-botleft-core” and “6k-botright-core”.

Core-Layer Switch Configuration

The following example shows the multicast configuration for “6k-botleft-core.”

```
ip multicast-routing
!
interface Loopback0
description MSDP local peer address
ip address 10.6.1.1 255.255.255.255
!
interface Loopback1
description Anycast RP address
ip address 10.6.2.1 255.255.255.255
```

```
ip pim sparse-mode
!
interface Vlanxx
description VLANs for other L3 links
ip address 10.0.0.x 255.255.255.252
ip pim sparse-mode
!
ip pim rp-address 10.6.2.1
ip msdp peer 10.6.1.2 connect-source Loopback0          Identifies the TCP peer and source interface.
!
ip msdp cache-sa-state                                  Creates SA state (S,G).
ip msdp originator-id Loopback0                         Sets RP address in SA messages to Loopback 0.
```

The following example shows the multicast configuration for "6k-botright-core."

```
ip multicast-routing
!
interface Loopback0
description MSDP local peer address
ip address 10.6.1.2 255.255.255.255
!
interface Loopback1
description Anycast RP address
ip address 10.6.2.1 255.255.255.255
ip pim sparse-mode
!
interface Vlanxx
description VLANs for other L3 links
ip address 10.0.0.x 255.255.255.252
ip pim sparse-mode
!
ip pim rp-address 10.6.2.1
ip msdp peer 10.6.1.1 connect-source Loopback0
!
ip msdp cache-sa-state
ip msdp originator-id Loopback0
```

Figure 2-8 provides a logical view of the MSDP configuration of the core switches.

Figure 2-8 Logical View of MSDP Configuration

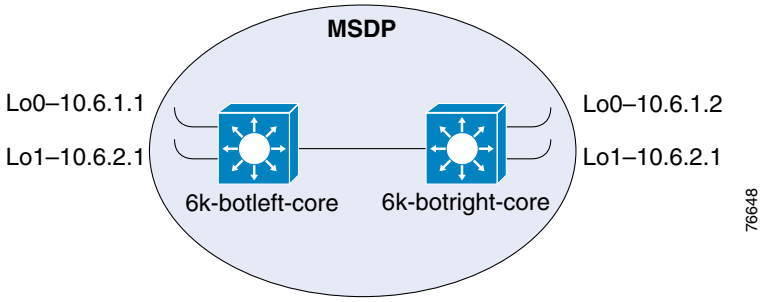

Distribution-Layer Switch Configuration

The following example shows the multicast configuration for "6k-b1-dist-left."

```
ip multicast-routing
!
interface Vlanxx
 description VLANs for Access Layer VLANs
 ip address 10.1.x.x 255.255.255.0
 ip pim sparse-mode
!
interface Vlanxx
 description VLANs for other L3 links to Core
 ip address 10.0.0.x 255.255.255.252
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1          Identifies the RP addresses for this device pointing to the Anycast RP
!                                   address in the core layer.
ip pim spt-threshold infinity       Reduces multicast state (S,G) from the leaf routers by keeping traffic
!                                   on the shared tree. (Optional)
```

The following example shows the multicast configuration for "6k-b1-dist-right."

```
ip multicast-routing
!
interface Vlanxx
 description VLANs for Access Layer VLANs
 ip address 10.1.x.x 255.255.255.0
 ip pim sparse-mode
!
interface Vlanxx
 description VLANs for other L3 links to Core
 ip address 10.0.0.x 255.255.255.252
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1          Identifies the RP addresses for this device pointing to the Anycast RP
!                                   address in the core layer.
ip pim spt-threshold infinity       Reduces multicast state (S,G) from the leaf routers by keeping traffic
!                                   on the shared tree. (Optional)
```

Summary

In summary, when using IP multicast with MoH or IP/TV in the campus follow these recommendations.

- Use PIM-SM (recommended)
- Ensure that the DR is configured to be the HSRP Active router (recommended)
- Use Anycast RP (recommended)
- Keep RP placement simple (recommended)

CHAPTER 3

IP Multicast in a Wireless LAN

This chapter describes the configurations needed to control IP Multicast traffic over a Wireless LAN (WLAN) and includes the following sections:

- Multicast WLAN Deployment Recommendations
- IP Multicast WLAN Configuration
- Other Considerations
- Summary

Tip For information about wireless theory, deployment, and configuration, please see the *Cisco AVVID Wireless LAN Design SRND*.

Note This chapter uses MoH and IP/TV in the examples. It does not, however, provide detailed configurations and designs for MoH and IP/TV. A basic MoH and IP/TV implementation is covered in Chapter 7, "Multicast Music-on-Hold and IP/TV Configurations."

Also, other types of IP multicast implementations, such as IP multicast for financial deployments, are not covered.

Multicast WLAN Deployment Recommendations

By default, IP multicast traffic is permitted to stream across a WLAN. However, because WLANs use shared bandwidth, certain measures should be taken to prevent saturation of the available bandwidth. If IP multicast traffic is not required on the wireless network, it is recommended that a boundary be configured to block the multicast traffic. The best place to control IP Multicast traffic is on the routers and switches that connect to the wireless Access Points (AP) and bridges. If a Layer 3 device is not available for use in deploying the configurations described in this chapter, then see the *Cisco AVVID Network Infrastructure Wireless LAN Design SRND* for recommendations for using AP and bridge MAC and IP filters to block traffic.

Note Filters on the AP and bridge do not provide the flexibility needed for true multicast control.

If IP Multicast is to be deployed and streamed across the wireless network, then the following recommendations should be implemented:

- Prevent unwanted multicast traffic from being sent out on the air interface.
  - Place the WLAN in its own subnet.
  - Control which multicast groups are allowed by implementing multicast boundaries on the egress Layer 3 interface connecting to the VLAN or interface to the AP or bridge.
- To gain the highest AP/bridge performance for multicast traffic and data traffic, configure the APs and bridges to run at the highest possible fixed data rate. This removes the requirement for multicast to clock out at a slower rate, which can impact the range of the AP/bridge and must be taken into account in the site survey.
- If multicast reliability is a problem (seen as dropped packets), ignore the preceding recommendation and use a slower data rate (base rate) for multicast. This gives the multicast a better signal-to-noise ratio and can reduce the number of dropped packets.
- Test the multicast application for suitability in the WLAN environment. Determine the application and user performance effects when packet loss is higher than that seen on wired networks.

IP Multicast WLAN Configuration

The **ip multicast boundary** command configures an administratively-scoped boundary on an interface for multicast group addresses found in the range defined by an access list. No multicast packets are allowed to flow across the boundary from either direction, except those packets explicitly allowed by the access list.

Controlling IP Multicast in a WLAN with Access Points

Figure 3-1 shows the topology for a WLAN using an AP. The IP multicast source is the IP/TV server (10.5.10.22). There are two multicast streams being sourced from the IP/TV server.

- 239.255.0.1 is a high-rate (1.4 Mbps) video stream.
- 239.192.248.1 is a low-rate (100 Kbps) video stream.

The low-rate stream is allowed and the high-rate stream is disallowed on the WLAN link. A multicast boundary is used to control multicast forwarding and IGMP packets.

Note IP multicast configuration for the campus is covered in Chapter 2, “IP Multicast in a Campus Network.”

Figure 3-1 Testbed for Wireless LAN using an Access Point

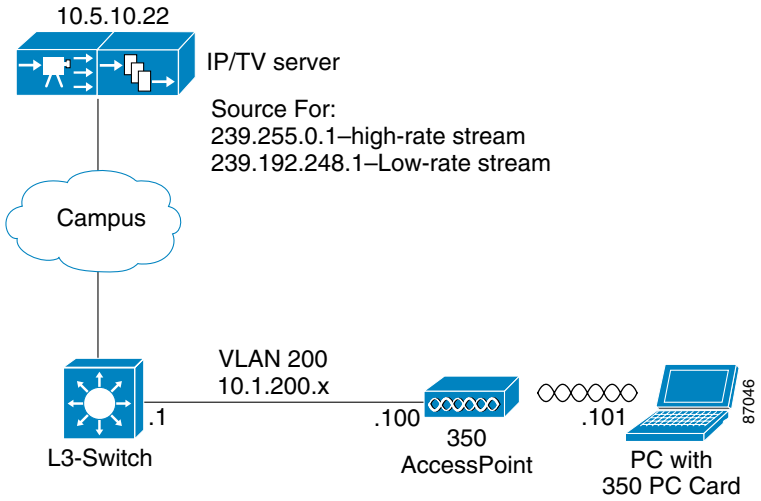

In this configuration:

- L3-SWITCH connects to the campus network and the Cisco Aironet 350 Access Point (10.1.200.100).
- The VLAN 200 interface on L3-SWITCH has the IP address of 10.1.200.1 and is the interface that provides the boundary for IP multicast.
- The laptop computer (10.1.200.101) has a Cisco Aironet 350 PC Card and is running the IP/TV Viewer software.

Below is the configuration is for L3-SWITCH.

```
interface Vlan200
 description WLAN VLAN
 ip address 10.1.200.1 255.255.255.0
 ip pim sparse-mode                          Enables PIM on the interface.
 ip multicast boundary IPMC-WLAN             Boundary refers to named ACL "IPMC-WLAN" and controls
!                                            multicast forwarding AND IGMP packets.
ip access-list standard IPMC-WLAN
 permit 239.192.248.1                        Permits low-rate stream (239.192.248.1).
```

See Verification and Testing for steps to verify this configuration.

Controlling IP Multicast in a P2P WLAN using Bridges

The same boundary that was deployed in the AP scenario is used with the bridge scenario. Figure 3-2 shows the topology for a WLAN using a bridge for a Point-to-Point (P2P) connection. The IP/TV server (10.5.10.22) is sourcing the same groups as in the previous example:

- 239.255.0.1 is a high-rate (1.4 Mbps) video stream.
- 239.192.248.1 is a low-rate (100 Kbps) video stream.

The low-rate stream is allowed and the high-rate stream is disallowed on the P2P wireless link. To control what multicast traffic passes over the P2P link, only the **ip multicast boundary** configuration on ROUTER is needed. Because the multicast boundary prevents hosts from joining unwanted groups, the network never knows to forward unwanted traffic over the P2P link.

Figure 3-2 Testbed for Point-to-Point Wireless Network using Bridges

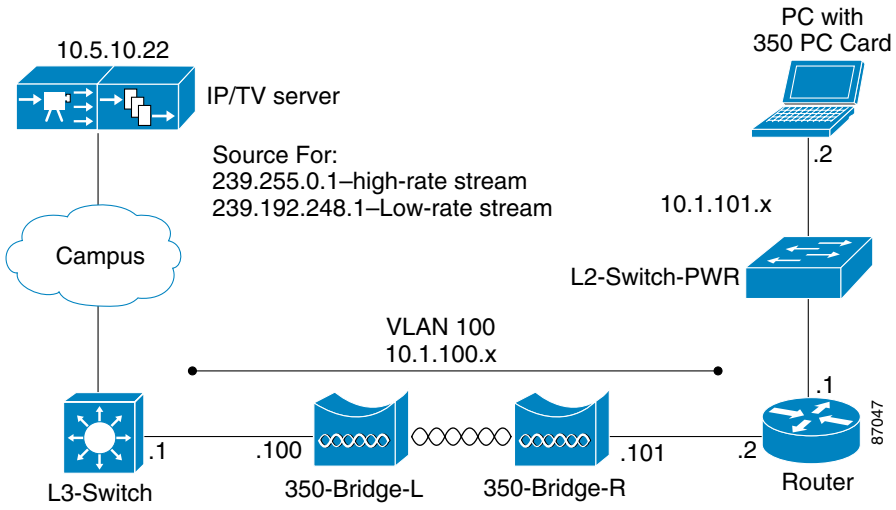

In this configuration:

- L3-SWITCH (VLAN 100-10.1.100.1) connects to the campus network and the P2P wireless network.
- The P2P wireless link is made possible by two Cisco Aironet 350 Bridges, 350-Bridge-L (10.1.100.100) and 350-Bridge-R (10.1.100.101).
- ROUTER (10.1.100.2) connects to the P2P wireless network and the remote site network (10.1.101.1) via L2-SWITCH-PWR.
- The laptop computer (10.1.101.2) is running the IP/TV Viewer software.

If the remote side of the P2P link has a Layer 2 switch and no Layer 3 switch or router, then a boundary can be placed on the VLAN 100 interface of L3-SWITCH2. Also, in a Point-to-Multipoint (P2MP) deployment, a mix of both may be needed. Both configurations are shown here for reference.

Following is the configuration for L3-SWITCH.

```
interface Vlan100
 description VLAN for P2P Bridge
 ip address 10.1.100.1 255.255.255.0
 ip pim sparse-mode                          Enables PIM on the interface.
 ip multicast boundary IPMC-BRIDGE           Boundary refers to named ACL "IPMC-BRIDGE."
!
ip access-list standard IPMC-BRIDGE
 permit 239.192.248.1                        Permits low-rate stream (239.192.248.1).
```

To prevent unwanted IGMP messaging and multicast traffic from traversing the P2P wireless link on the receiver side (remote LAN - 10.1.101.x), an **ip multicast boundary** is configured on the Fast Ethernet 0/1 interface of ROUTER.

Following is the configuration for ROUTER.

```
interface FastEthernet 0/1
 description Local LAN in Remote Site
 ip address 10.1.101.1 255.255.255.0
 ip pim sparse-mode                          Enables PIM on the interface.
 ip multicast boundary IPMC-BRIDGE           Boundary refers to named ACL "IPMC-BRIDGE."

ip access-list standard IPMC-BRIDGE
 permit 239.192.248.1                        Permits low-rate stream (239.192.248.1).
```

Verification and Testing

To ensure that the multicast boundary feature is working properly, use the following commands:

- **show ip mroute active**
- **debug ip igmp 239.192.248.1**
- **debug ip igmp 239.255.0.1**

The following examples show how to test the controls that have been implemented. In these examples, the PC joins the 239.192.248.1 (low-rate IP/TV stream); the wireless clients and P2P WLAN should be able to access the stream. Then the PC attempts to join the disallowed group 239.255.0.1 (high-rate IP/TV stream); no IGMP joins and no multicast traffic for that group should be allowed.

A deny and log clause have been added to the access lists to show basic permit and deny messages.

```
ip access-list standard acl
 permit 239.192.248.1 log
 deny   any log
```

Test 1: WLAN with AP

To test the WLAN with AP deployment, do the following:

Step 1 Ensure that there are no active multicast streams.

```
L3-SWITCH#show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps
L3-SWITCH#
```

Step 2 On the PC, open the IP/TV viewer and request the program associated with the 239.192.248.1 group.

The following ACL console message should appear on L3-SWITCH showing that the traffic was permitted (number of packets may vary):

```
5w1d: %SEC-6-IPACCESSLOGS: list IPMC-WLAN permitted 239.192.248.1 7 packets
```

Step 3 Issue the **show ip mroute active** command on L3-SWITCH to see the active multicast stream. The result is an active multicast stream for group 239.192.248.1 being sourced from the IP/TV server 10.5.10.22 with a rate of 100 kbps.

```
L3-SWITCH#show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps
```

```
Group: 239.192.248.1, (low-rate)
   Source: 10.5.10.22 (IPTV)
     Rate: 23 pps/100 kbps(1sec), 94 kbps(last 30 secs), 115 kbps(life avg)
```

The IP/TV viewer should be displaying the video content on the screen.

Step 4 On the PC, stop the IP/TV program that is running for the 239.192.248.1 group.

Step 5 Enable the debug for the high-rate stream:

```
L3-SWITCH#debug ip igmp 239.255.0.1
```

Step 6 On the PC, open the IP/TV viewer and request the program associated with the 239.255.0.1 group.

The following ACL console message should appear on L3-SWITCH showing that the traffic was denied (number of packets may vary):

```
5w1d: %SEC-6-IPACCESSLOGS: list IPMC-WLAN denied 239.255.0.1 1 packet
```

Step 7 The following debug entry should appear for the discarded IGMP join attempt:

```
5w1d: IGMP(0): Discard report at boundary (Vlan200) for 239.255.0.1
```

There should be no multicast state active for the 239.255.0.1.

Test 2: WLAN with P2P Bridges

To test the WLAN with P2P bridges deployment, do the following:

Step 1 Ensure that there are no active multicast streams.

```
L3-SWITCH#show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps
L3-SWITCH#

ROUTER#show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps
ROUTER#
```

Step 2 On the PC, open the IP/TV viewer and request the program associated with the 239.192.248.1 group. As an alternative to using the IP/TV viewer, run the **ip igmp join-group 239.192.248.1** command on the Fast Ethernet 0/1 interface on ROUTER.

The following ACL console message should appear on ROUTER showing that the traffic was permitted (number of packets may vary):

```
1w1d: %SEC-6-IPACCESSLOGS: list IPMC-BRIDGE permitted 239.192.248.1 5 packets
```

Step 3 Run the **show ip mroute active** command on L3-SWITCH and ROUTER.

```
ROUTER#show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps

Group: 239.192.248.1, (low-rate)
   Source: 10.5.10.22 (IPTV)
     Rate: 23 pps/100 kbps(1sec), 91 kbps(last 20 secs), 123 kbps(life avg)
```

The IP/TV viewer should be displaying the video content on the screen.

Step 4 On the PC, stop the IP/TV program that is running for the 239.192.248.1 group.

Step 5 Enable the debug for the high-rate stream.

```
ROUTER#debug ip igmp 239.255.0.1
```

Step 6 On the PC, open the IP/TV viewer and request the program associated with the 239.255.0.1 group.

The following ACL console message should appear on ROUTER showing that the traffic was denied (number of packets may vary):

```
1w1d: %SEC-6-IPACCESSLOGS: list IPMC-BRIDGE denied 239.255.0.1 1 packet
```

The following debug entry should appear for the discarded IGMP join attempt:

```
1w1d: IGMP(0): Discard report at boundary (FastEthernet0/1) for 239.255.0.1
```

There should be no multicast state active for the 239.255.0.1

Other Considerations

The following additional considerations apply to deploying IP multicast in a WLAN environment:

- The WLAN LAN extension via EAP and WLAN static WEP solutions can support multicast traffic on the WLAN; the WLAN LAN extension via IPSec solution cannot.
- The WLAN has an 11 Mbps available bit rate that must be shared by all clients of an AP. If the AP is configured to operate at multiple bit-rates, multicasts and broadcasts are sent at the lowest rate to ensure that all clients receive them. This reduces the available throughput of the network because traffic must queue behind traffic that is being clocked out at a slower rate.
- WLAN clients can roam from one AP to another seamlessly within the same subnet. If roaming multicast is to be supported, Cisco Group Management Protocol (CGMP) and/or Internet Group Management Protocol (IGMP) snooping must be turned off for the port to which the AP is connected because a multicast user roaming from one AP to another is roaming from one switch port to another. The new switch port might not have this stream setup and it has no reliable way of determining the required multicast stream.
- Multicast and broadcast from the AP are sent without requiring link-layer acknowledgement. Every unicast packet is acknowledged and retransmitted if unacknowledged. The purpose of the acknowledgement is to overcome the inherent unreliable nature of wireless links. Broadcasts and multicasts are unacknowledged due to the difficulty in managing and scaling the acknowledgements. This means that a network that is seen as operating well for unicast applications, can experience degraded performance in multicast applications.
- Enterprise customers who are using WLAN in laptops would normally use (Constant Awake Mode) CAM as the Power-Save Mode. If delay-sensitive multicast traffic is being sent over the WLAN, customers should ensure that only the CAM configuration is used on their WLAN clients. Based on the 802.11 standard, if the client is in power-save mode, then the AP will buffer broadcast and multicast traffic until the next beacon period that contains a delivery traffic information map (DTIM) transmission. The default period is 200ms. Enterprises that use WLAN on small handheld devices will most likely need to use the WLAN power-save features (Max or Fast) and should not attempt to run delay-sensitive multicast traffic over the same WLAN.

Summary

In summary, when using IP multicast in the WLAN, follow these recommendations.

- Place the WLAN AP or bridge on a separate VLAN or Layer 3 interface so multicast boundaries can be implemented.
- Use the **ip multicast boundary** command to prevent IGMP joins and multicast forwarding on denied multicast groups.
- In a WLAN using AP, the boundary should be placed on the VLAN or Layer 3 interface connecting to the AP.
- In a WLAN using bridges, the boundary is placed on the VLAN or Layer 3 interface connecting to the remote receiver side. If no Layer 3 capable device is used at the remote site, the boundary is placed on the VLAN or Layer 3 interface connecting to the bridge at the main site. Also, a combination of a boundary at the receiver side and bridge connection at the main site, may be needed in a P2MP deployment.
- Set the highest possible fixed data rate on the APs and bridges to ensure the best possible performance for multicast and data traffic.
- If dropped packets occur and impact the performance of the application, the fixed data rate on the APs and bridges may need to be reduced to ensure a better signal-to-noise ratio, which can reduce dropped packets.

CHAPTER 4

IP Multicast in a Data Center

Because many of the servers residing in the data center may use multicast to communicate, it becomes necessary to enable IP multicasting on the data center switches. Platforms that deliver video services, such as Cisco's IP/TV broadcast servers, rely on multicast to deliver their services to end users. Many back-end applications use multicast for replication purposes.

This chapter provides an overview of the data center architecture and recommendations for implementing IP multicast in the data center. For more information about designing data centers, see the *Designing Enterprise Data Centers* Solution Reference Network Design guide.

Data Center Architecture Overview

As with the campus network, the data center uses a hierarchical architecture composed of layers. The layers of a data center are:

- Aggregation Layer
- Front-End Layer
- Application Layer
- Back-End Layer
- Storage Layer

Aggregation Layer

The aggregation layer consists of network infrastructure components and other devices supporting application, security, and management services. The aggregation layer is analogous to the traditional *distribution layer* in the campus network in its Layer 3 and Layer 2 functionality. Those services that are common to the devices in the front-end layer and the layers behind it should be centrally located for consistency, manageability, and predictability. This makes the aggregation layer the location where services are centralized. The devices in this layer include aggregation switches that connect the server farms to the core layer, content switches, firewalls, intrusion detection systems (IDSs), content engines, and SSL offloading.

For IP multicast:

- The aggregation layer requires multicast routing with PIM Sparse-mode to be configured.
- The Layer 3 switches in the aggregation layer point to the RPs previously defined in the campus core.

- A dedicated VLAN is used to connect multicast sources that are not located in the front-end layer. An example multicast source that will be placed in the dedicated VLAN is a streaming media application, like an IP/TV broadcast server.
- Some sources are part of another server role that may need to be located in the front-end layer. An example of a multicast source that is located in the front-end layer is Multicast Music-on-Hold (MMoH). MMoH is often deployed in a co-resident fashion with Cisco Call Manager.

Front-End Layer

The front-end layer consists of infrastructure, security, and management devices supporting the front-end server farms. The front-end layer is analogous to the traditional *access layer* in the campus network and provides the same functionality. The front-end server farms typically include FTP, Telnet, TN3270, SMTP, Web servers, and other business application servers. In addition, it includes network-based application servers, such as IPTV Broadcast servers, and call managers that are not placed at the aggregation layer due to port density or other design caveats.

The specific features required depend on the server and their functions. For example, if Video streaming over IP is supported, multicast must be enabled, or if Voice over IP is supported, QoS must be enabled. Layer 2 connectivity through VLANs is required between servers and service devices, such as content switches, and between servers that belong to the same server farm or subnet and are located in the same or different access switches. This is known as *server-to-server* communication, which could also span multiple tiers. Other features may include the use of host IDS if the servers are critical and need constant monitoring. In general, the infrastructure components such as the Layer 2 switches provide intelligent network services that enable front-end servers to provide their functions.

Cisco Catalyst switches support IGMP snooping or CGMP at Layer 2. Because IGMP snooping or CGMP (platform dependent) is enabled on Layer 2 switches by default, no multicast configuration is required at the front-end layer.

Application Layer

The application layer consists of the infrastructure, security, and management devices that support the application servers. Applications servers run a portion of the software used by business applications and provide the communication logic between front-end and the back-end, which is typically referred to as the middleware or business logic. Application servers translate user requests to commands the back-end database systems understand. Increasing the security at this layer is focused on controlling the protocols used between the front-end servers and the application servers.

The features required at this layer are almost identical to those needed in the front-end layer. Like the front-end layer, the application layer infrastructure must support intelligent network services as a direct result of the functions provided by the application services. However, the application layer requires additional security.

Additional security is based on how much protection is needed by the application servers as they have direct access to the database systems. Depending on the security policies, firewalls between web and application servers, IDS, and host IDSs are used. By default firewalls do not permit nor do they support multicast forwarding. Careful consideration must be given to the deployment of firewall services when multicast traffic is to be permitted across a secure boundary. It is not uncommon to deploy GRE tunneling, multicast helper, or Policy Based Routing (PBR) to support multicast across secured boundaries. These methods are difficult to deploy and troubleshoot. They also require Layer 3 intelligence in an area of the network that needs only Layer 2 features.

The best way to accommodate IP Multicast in the data center is to place multicast sources on a separate VLAN that is located in the aggregation layer. Placing the multicast sources on a dedicated VLAN will bypass the difficulties of using firewalls and multicast, as well as other issues like multicast on router-less subnets.

Back-End Layer

The back-end layer consists of the infrastructure, security, and management devices supporting the interaction with the database systems that hold the business data. The back-end layer features are almost identical to those needed at the application layer, yet the security considerations are more stringent and aimed at protecting the data, critical or not.

The back-end layer is primarily for the relational database systems that provide the mechanisms to access the Enterprise information, which makes them highly critical. The hardware supporting the relational database systems range from medium sized servers to mainframes, some with locally attached disks and others with separate storage.

Multicast requirements are not common at the back-end layer. The synchronization of database contents requires an acknowledgement-based method of synchronization and multicast does not provide a mechanism to support an acknowledgement-based service, unless a vendor has written a proprietary extension.

Storage Layer

The storage layer consists of the storage infrastructure such as Fibre-Channel switches or iSCSI routers that connect server to storage devices, as well as the storage devices to where the data resides. At the data center, storage devices, such as tape and disk subsystems, need a high-speed connection to provide block level access to information from the database servers. This implies the disk subsystems are used to consolidate information from dedicated local disks to a centralized repository supporting the database systems. Multicast is not typically used at this layer.

Data Center Logical Topology

The logical architecture, presented in Figure 4-1, shows the relationship between the data center layers and the traditional campus layers. From a logical viewpoint, the front-end servers are separated from the applications servers, which are separated from the database servers. The traffic flow is then client-to-front-end servers, front-end servers-to-application servers, and finally application servers-to-database servers. The logical separation just implies each layer is a distinct functional area. Regardless of the logical separation, the Layer 2 switches providing connectivity within each layer can be though of as comparable to access switches and the aggregation switches equivalent to distribution switches.

Figure 4-1 Data Center Logical Topology

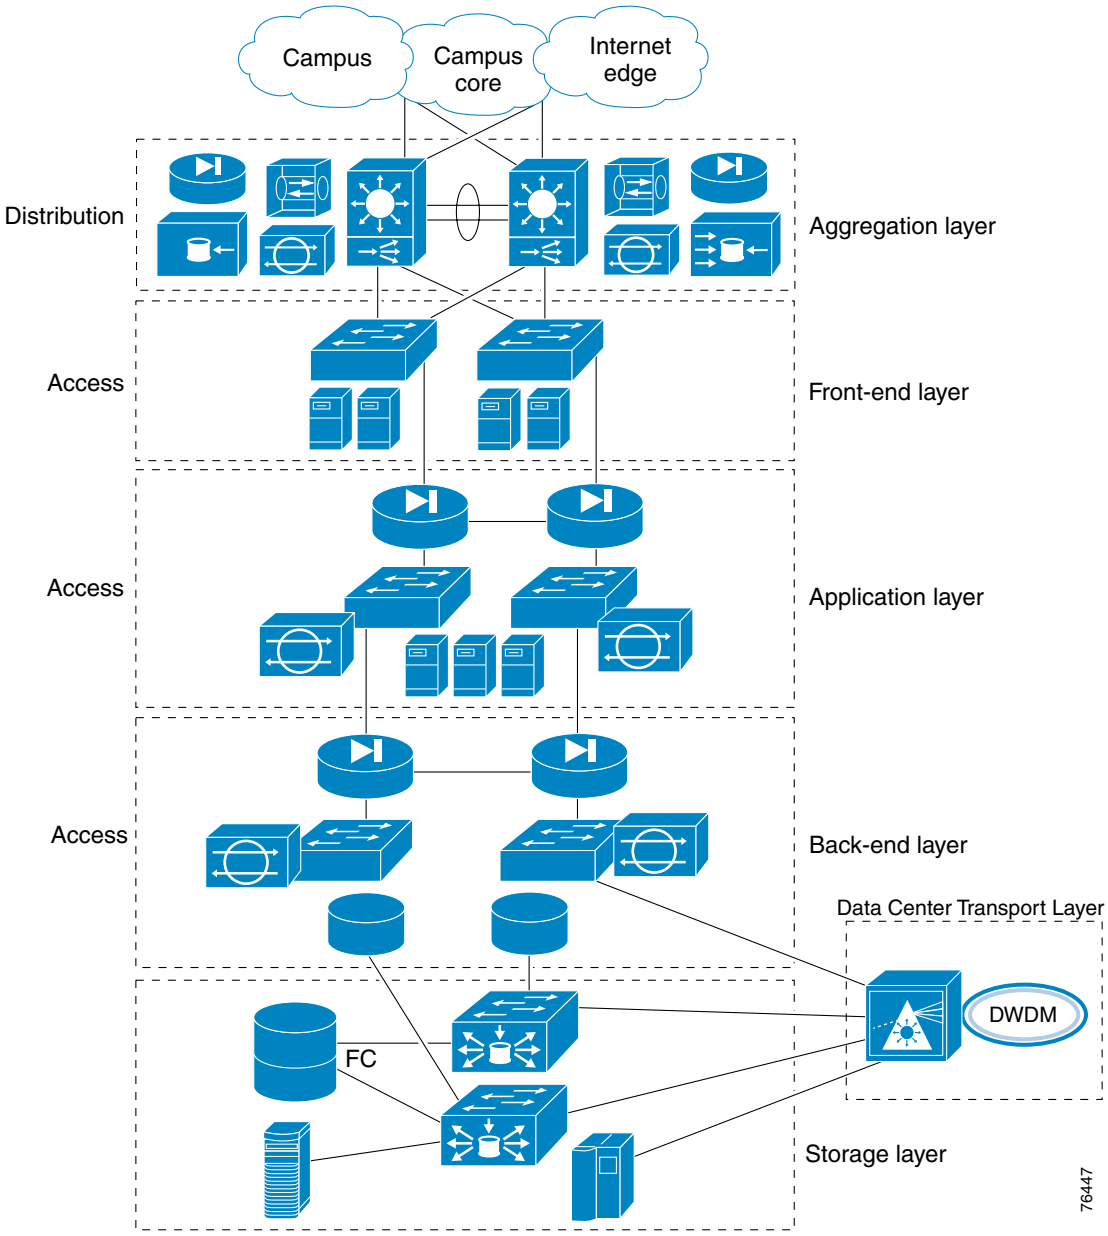

Multicast Data Center Deployment Recommendations

This section discusses the recommended and optional configurations for IP multicast data center deployment. The recommended guidelines are summarized below:

- Place multicast sources on a separate VLAN that is attached to the aggregation layer. Sources that require direct access to back-end resources (such as the application, database, or storage layers) are placed in the front-end layer.
- Use administratively scoped addresses to differentiate multicast applications by type and bandwidth.

- Deploy PIM Sparse-mode on the aggregation-layer switches and define the RPs as the previously configured campus core-layer switches.
- Select Catalyst switches that have IGMP snooping and use CGMP in low-end switches that do not support IGMP snooping.
- Use recommended commands to ensure that the correct RPs and sources are used.
- Use "show" commands to ensure proper operation of the multicast configurations and enable SNMP traps to log multicast events.
- If a firewall module (PIX) or CSM is used in the aggregation layer, then a VLAN that is terminated on the Layer 3 switch (aggregation layer) must be created for IP multicast traffic to bypass the firewall and CSM. Currently, the firewall and CSM do not support IP multicast forwarding. Once PIM is supported on the various services modules, then this step will not be necessary.

IP Multicast Data Center Configuration

The applications (MoH and IP/TV) used in this sample design (shown in Figure 4-2) have few sources to many receivers. The sources are located in the server farm. So, a complex distribution of RPs throughout the network is not required.

Figure 4-2 Data Center Design Reference Diagram

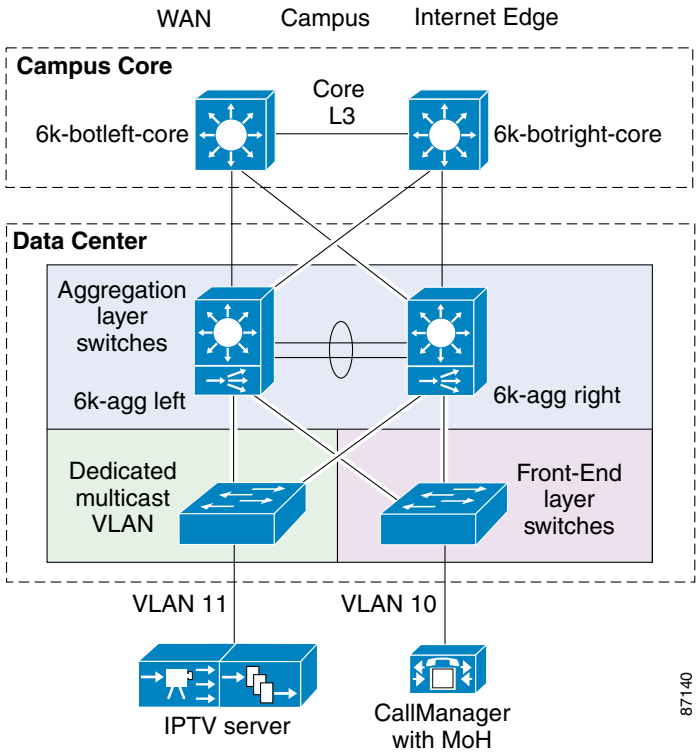

Core-Layer Switch Configuration

The following example shows excerpts of the configuration of "6k-botleft-core" that pertain to the aggregation layer.

```
interface Vlan50
 description To agg (6k-agg-left) for front-end/dedicated IPmc VLAN
 ip pim sparse-mode

!
```

The following example shows excerpts of the configuration of "6k-botright-core" that pertain to the aggregation layer.

```
interface Vlan60
 description To agg (6k-agg-right) for front-end/dedicated IPmc VLAN
 ip pim sparse-mode
!
```

Server Farm Aggregation Switch Configuration

The following example shows the multicast configuration for "6k-agg-left."

```
ip multicast-routing
!
interface Vlan50
 description To Campus Core (6k-botleft-core)
 ip pim sparse-mode

!
interface Vlan10
 description To front-end layer - MMoH server
 ip pim sparse-mode
!
interface Vlan11
 description To dedicated VLAN - IP/TV Server
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1          Identifies the RP addresses for this device pointing to the Anycast RP
!                                   address in the core layer.
```

The following example shows the multicast configuration for "6k-agg-right."

```
ip multicast-routing
!
interface Vlan60
 description To Campus Core (6k-botright-core)
 ip pim sparse-mode
!
```

```
interface Vlan10
 description To front-end layer - MMoH Server
 ip pim sparse-mode

!
interface Vlan11
 description To dedicated VLAN - IP/TV Server
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1
!
```

Identifies the RP addresses for this device pointing to the Anycast RP address in the core layer.

CHAPTER 5

IP Multicast in a WAN

This chapter discusses the basic layout needed to use IP multicast in a WAN and includes the following sections:

- Multicast WAN Deployment Recommendations
- IP Multicast WAN Configuration
- Summary

Note This chapter uses MoH and IP/TV in the examples. It does not, however, provide detailed configurations and designs for MoH and IP/TV. A basic MoH and IP/TV implementation is covered in Chapter 7, "Multicast Music-on-Hold and IP/TV Configurations."

Also, other types of IP multicast implementations, such as IP multicast for financial deployments, are not covered.

Multicast WAN Deployment Recommendations

This chapter discusses the recommended and optional configurations for IP multicast WAN deployment. The recommended guidelines are summarized below:

- Use IP Multicast to scale streaming applications, such as Music-on-Hold and IP/TV.
- Use administratively scoped addresses to differentiate multicast applications by type and bandwidth.
- Use Anycast RP.
- Select Catalyst switches that have IGMP snooping and use CGMP in low-end switches that do not support IGMP snooping.
- Use the recommended commands to ensure that the correct RPs and sources are used.
- Use IP Multicast boundaries to control where certain multicast streams go.
- Use "show" commands to ensure proper operation of the multicast configurations and enable SNMP traps to log multicast events.

IP Multicast WAN Configuration

The WAN topology, shown in Figure 5-1, is based on a dual-PVC Frame Relay model. The PVCs are not load-balanced. Preference is given to the PVC connected to "wanaggleft" by configuring the delay variable to influence the preference given by the IP Unicast routing protocol, EIGRP.

RP placement will occur on Layer 3 switches in the campus (in a server farm or core). RP resource requirements are discussed in Chapter 1, "IP Multicast Overview."

Key IP addresses are listed in the diagram for reference. The "Corporate Network" cloud represents the campus as a whole. The example design uses the two core switches as the RPs. IP multicast configuration for the campus is covered in Chapter 2, "IP Multicast in a Campus Network." In this chapter, a brief configuration for the campus RPs is shown for reference.

Figure 5-1 Multicast WAN Topology

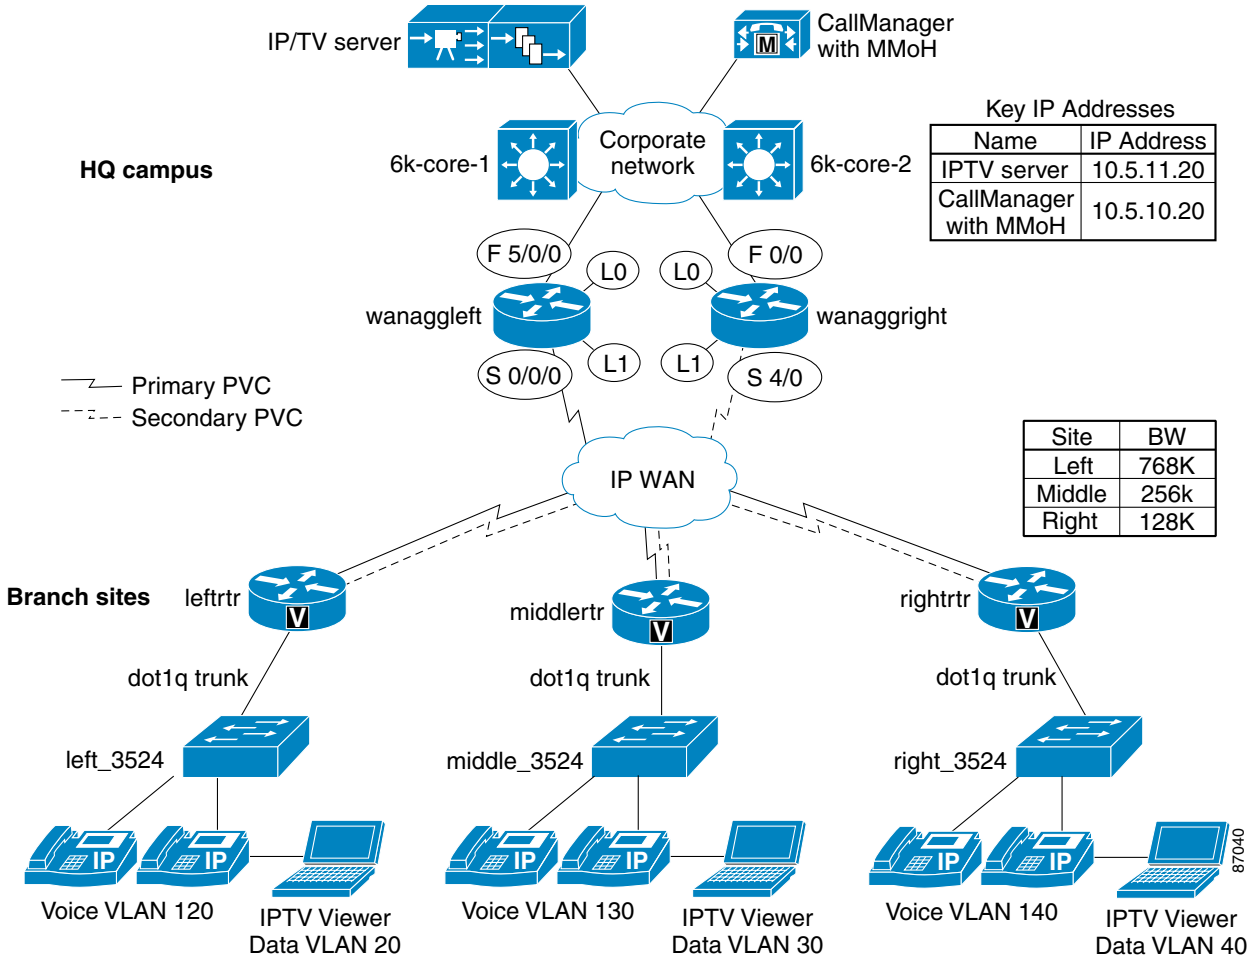

This section provides information and sample configurations for deploying the following IP multicast elements in a WAN:

- Anycast RP
- IGMP Snooping and CGMP

Anycast RP

This section provides sample Anycast RP configurations for the routers shown in Figure 5-1.

Branch

The IP multicast configuration steps for the branch routers are as follows:

- Enable IP multicast routing.
- Enable IP PIM Sparse Mode on each interface that will receive and forward multicast traffic.

Identify the RP for the router. With Anycast RP, the address defined is the Loopback address that is duplicated on each participating Anycast RP router. (Loopback 1 - 10.6.2.1).

Following is the basic multicast configuration for "leftrtr." This configuration is duplicated on each branch router.

```
ip multicast-routing                                    Enable IP multicast routing globally.
!
interface FastEthernet0/0.20                            PIM-SM on the interface.
 description dot1q trunk interface for DATA VLAN
 ip pim sparse-mode
!
interface FastEthernet0/0.120
 description dot1q trunk interface for VOICE VLAN
 ip pim sparse-mode
!
interface Serial0/0.1 point-to-point
 description To wanaggright
 ip pim sparse-mode
!
interface Serial0/0.2 point-to-point
 description To wanaggleft
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1                              Loopback 1 address of both WAN aggregation
!                                                       routers.
```

WAN Aggregation

The IP multicast configuration steps for the WAN aggregation routers are as follows:

- Enable IP multicast routing.
- Enable IP PIM Sparse Mode on each interface that will receive and forward multicast traffic.

Identify the RP for the router. With Anycast RP, the address defined is the Loopback address that is duplicated on each participating Anycast RP router. (Loopback 1 - 10.6.2.1).

Note Multicast Distributed Fast Switching (MDFS) is **not** supported on any type of virtual interface, such as virtual-template or multilink interfaces. Enabling **ip mroute-cache distributed** on unsupported interfaces can result in process switching of the multicast packets or dropped packets.

Following is the basic multicast configuration for "wanaggleft."

```
ip multicast-routing distributed                    Distributed keyword used on Cisco 7500 routers.
!


interface FastEthernet5/0/0
 description To HQ Switch
 ip pim sparse-mode
 ip mroute-cache distributed                        Enable distributed multicast fast-switching.
!
interface Serial0/0/0
 description WANAGGLEFT to Branch Routers
 ip mroute-cache distributed
!
interface Serial0/0/0.1 point-to-point
 description To LEFTRTR DLCI 121
 ip pim sparse-mode
!
interface Serial0/0/0.2 point-to-point
 description To MIDDLERTR DLCI 131
 ip pim sparse-mode
!
interface Serial0/0/0.3 point-to-point
 description To RIGHTRTR DLCI 141
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1
```

Following is the basic multicast configuration for "wanaggright."

```
ip multicast-routing
!
interface FastEthernet0/0
 description To HQ Switch
 ip pim sparse-mode
!
interface Serial4/0.1 point-to-point
 description To LEFTRTR DLCI 120
 ip pim sparse-mode
!
interface Serial4/0.2 point-to-point
 description To MIDDLERTR DLCI 130
 ip pim sparse-mode
!
interface Serial4/0.3 point-to-point
 description To RIGHTRTR DLCI 140
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1
```

MSDP Filters

There are recommended MSDP filters that should be applied if the network is connected to native IP Multicast on the Internet. The MSDP filters are used to reduce the excessive amount of (S, G) state that is passed between Internet MSDP peers.

Tip For more information about MSDP filters, see:
ftp://ftpeng.cisco.com/ipmulticast/config-notes/msdp-sa-filter.txt

IGMP Snooping and CGMP

The topology shown in Figure 5-1 uses CGMP in the branch office switches. CGMP is a better option than IGMP snooping for low-end switches without special hardware.

To take advantage of CGMP, CGMP support must be enabled on the locally attached branch router. The branch office switches in this design scenario are Catalyst 3524-PWR XLs. The Catalyst 3524 supports only CGMP, which is on by default.

Following is the CGMP configuration for "leftrtr."

```
ip multicast-routing
!
interface FastEthernet0/0.20
ip pim sparse-mode
ip cgmp                                    Enable CGMP on the router interface. PIM must be
!                                          enabled prior to CGMP being enabled.
interface FastEthernet0/0.120
ip pim sparse-mode
ip cgmp
```

CGMP operation between the switch and router can be verified using the **show cgmp**, as shown below:

```
left_3524#show cgmp
CGMP is running.
CGMP Fast Leave is not running.
CGMP Allow reserved address to join GDA .
Default router timeout is 300 sec.

vLAN     IGMP MAC Address   Interfaces
------  -----------------   -----------


vLAN     IGMP Router        Expire   Interface
------  -----------------  --------  ----------
   20    0004.c16d.49a0     281 sec   Fa0/24          MAC address and port associated with
  120    0004.c16d.49a0     281 sec   Fa0/24          attached router.
```

Summary

In summary, when using IP multicast with MoH or IP/TV in the WAN follow these recommendations.

- Use PIM-SM (recommended)
- Anycast RP with RPs located in the campus (recommended)
  - Fast convergence, redundancy, load balancing

CHAPTER 6

IP Multicast in a Site-to-Site VPN

This chapter discusses the basic layout needed to use IP multicast in a Virtual Private Network (VPN) and includes the following sections:

- Site-to-Site VPN Overview
- VPN Deployment Model
- Multicast VPN Deployment Recommendations
- Multicast Site-to-Site VPN Deployment
- Summary

Site-to-Site VPN Overview

The following section is an overview of Site-to-Site VPNs. The following topics are discussed:

- IPSec Deployment with GRE
- Managing IPSec and GRE Overhead
- Redundant VPN Head-end Design

IPSec Deployment with GRE

Generic routing encapsulation (GRE) is often deployed with IPSec for several reasons, including:

- IPSec supports unicast IP only. If network layer protocols other than IP are to be supported, an IP encapsulation method must be chosen so that those protocols can be transported in IP packets.
- Multicast is not supported with IPSec. Because IPSec was created to be a security protocol between two and only two devices, a service such as multicast is problematic. An IPSec peer encrypts a packet so that only one other IPSec peer can successfully perform the de-encryption. Multicast is not compatible with this mode of operation.
- IPSec tunnels are not logical tunnel interfaces for routing purposes. A GRE tunnel, on the other hand, is a logical router interface for purposes of forwarding IP (or any other network protocol) traffic. A GRE interface may appear as a next hop interface in a routing table. If a routing protocol using unicast as a peer communication method (such as BGP) were to be run over an IPSec tunnel alone, the router would learn about the available routes from the interface that the IPSec was configured over. This would be problematic if the IPSec peer is not directly connected to that physical interface.

Managing IPSec and GRE Overhead

The use of IPSec and GRE causes some packet expansion. This is a concern when working with MTU-sized packets with either of these protocols. When both protocols are used together, the packet expansion can be 90 bytes or more depending upon the IPSec options and transforms used. This has an effect in two places.

- Networks with a high percentage of small packets. A voice-over-IP network is a good example. Because of the higher percentage of packet expansion that is caused with small packets, additional link bandwidth may need to be configured in such a network.
- MTU-sized packets. Because of the packet expansion that takes place during the encapsulation/encryption process, these packets will be larger than the MTU of many media types commonly used in networks today. These packets will be fragmented upon forwarding if the MTU of the packet has not been set to a lower value prior to the encryption. A work around for this exists with path MTU discovery, which will dynamically discover the smaller MTU of an encapsulated packet.

Figure 6-1 shows how the IP packet is expanded over the MTU size when GRE and IPSec are added.

Figure 6-1 IPSec/GRE Packet Expansion

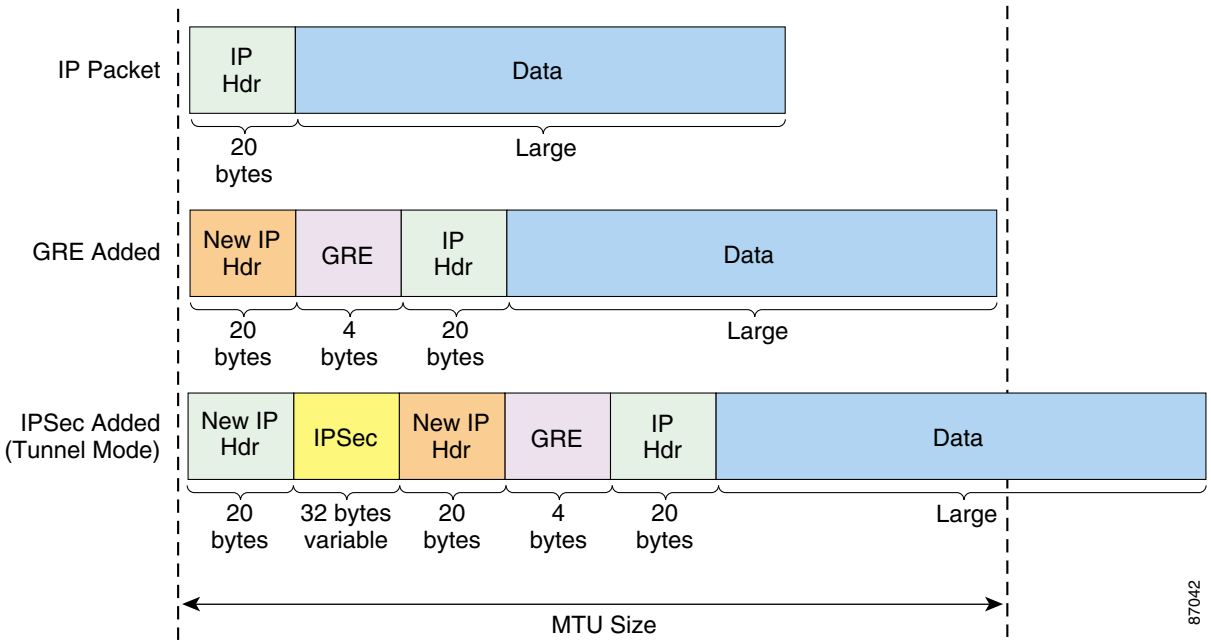

Redundant VPN Head-end Design

Because fail-safe operation is a mandatory feature in any enterprise network, redundancy should be built into head-end designs. From each branch location, a minimum of two tunnels are configured back to different head end devices. When sizing the head end installation, the failure of a single head end device should be taken into consideration. When adding a intelligent services like IP multicast, adding additional head-end routers and spreading the load of the VPN terminations across more devices will allow for the head-end routers to “share” CPU time, thus making the solution more scalable.

Tip Additional information for tunnel allocation and sizing devices can be found in the *Cisco AVVID Network Infrastructure Enterprise Site-to-Site VPN Design SRND*.

Figure 6-2 Tunnel Aggregation for Resiliency

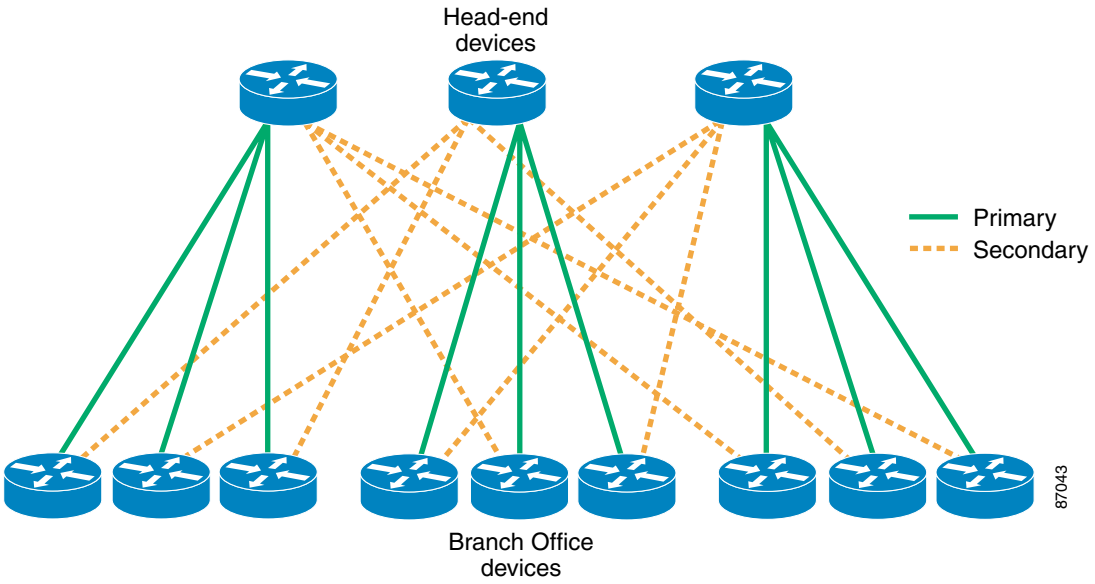

To plan for proper tunnel aggregation and load distribution in the case of a head-end device failure, the following process should be used:

- Start with the number of total branch devices to be aggregated at the head-end.
- Divide this number by the number of head-end devices.
- Multiply the result by 2 for primary and secondary tunnels. This is the total tunnels per head-end device.
- Allocate the primary tunnels to each head-end device in the arrangement shown in Figure 2 above (in green).
- For each group, allocate the secondary tunnels in a round-robin fashion to all head-end devices except the one serving as a primary for that group. This arrangement is also shown in Figure 2 above (in yellow).
- Check to see that each head-end device is now allocated the same number of total tunnels per head-end device.

VPN Deployment Model

This section describes the Site-to-Site VPN configuration that is used as the foundation for delivering IP multicast. The following topics are discussed:

- IKE Configuration
- IPSec Transform and Protocol Configuration
- Access List Configuration for Encryption
- Crypto Map Configuration
- Applying Crypto Maps
- Static Route Configuration

Figure 6-3 provides a reference for the configuration of both the VPN and the IP multicast deployment.

Figure 6-3 Example VPN Deployment

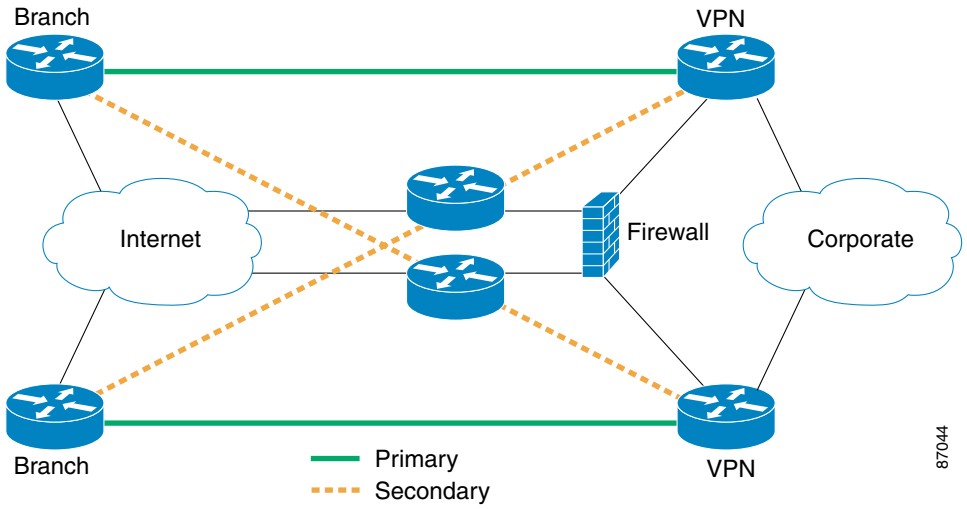

IKE Configuration

There must be at least one matching Internet Key Exchange (IKE) policy between two potential IPSec peers. The example configuration below shows a policy using pre-shared keys with 3DES as the encryption transform. There is a default IKE policy that contains the default values for the transform, hash method, Diffie-Helman group, authentication and lifetime parameters. This is the lowest priority IKE policy.

When using pre-shared keys, Cisco recommends that wildcard keys not be used. Instead, the example shows two keys configured for two separate IPSec peers. The keys should be carefully chosen; “cisco” is used only as an example. The use of alpha-numeric and punctuation characters as keys is recommended.

The IKE configurations shown below are all the same for each device, with the exception of the unique IP address used for each router.

Head-End

Following is the IKE configuration for VPN-HE-1.

```
interface FastEthernet0/1
 description to ISP for VPN
 ip address 131.108.1.1 255.255.255.252
!
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto isakmp key cisco address 131.108.101.1
crypto isakmp key cisco address 131.108.102.1
```

Following is the IKE configuration for VPN-HE-1.

```
interface FastEthernet0/1
 description to ISP for VPN
 ip address 131.108.1.5 255.255.255.252
!
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto isakmp key cisco address 131.108.102.1
crypto isakmp key cisco address 131.108.101.1
```

Branch

Following is the IKE configuration for VPN-Branch-1.

```
interface Serial0/0
 description to ISP for VPN
 ip address 131.108.101.1 255.255.255.252
!
crypto isakmp policy 1
 encr 3des
 authentication pre-share group 2
crypto isakmp key cisco address 131.108.1.1
crypto isakmp key cisco address 131.108.1.5
```

Following is the IKE configuration for VPN-Branch-2.

```
interface Serial0
 description to ISP for VPN
 ip address 131.108.102.1 255.255.255.252
!
crypto isakmp policy 1
 encr 3des
 authentication pre-share group 2
```

```
crypto isakmp key cisco address 131.108.1.1
crypto isakmp key cisco address 131.108.1.5
```

Tip These defaults and more information can be found at: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/fipsencr/srfike.htm#xtocid17729

IPSec Transform and Protocol Configuration

Transform is the list of operations performed on a dataflow to provide data authentication, data confidentiality, and data compression. For example, one transform is the ESP protocol with the HMAC-MD5 authentication algorithm, another transform is the AH protocol with the 56-bit DES encryption algorithm, and yet another is the ESP protocol with the HMAC-SHA authentication algorithm.

The transform set must match on the two IPSec peers. The transform set names are locally significant only. However, the encryption transform, hash method, and the particular protocols used (ESP or AH) must match. You may also configure data compression here but it is not recommended on peers with high-speed links. There can be multiple transform sets for use between different peers. The example below shows the exact same transform set for the head-end and branch routers.

Head-End

Following is the transform configuration for the head-end routers.

```
crypto ipsec transform-set strong
esp-3des esp-sha-hmac
```

Branch

Following is the transform configuration for the branch routers.

```
crypto ipsec transform-set strong
esp-3des esp-sha-hmac
```

Tip More information can be found at: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/fipsencr/srfipsec.htm#xtocid105784

Access List Configuration for Encryption

The access list entries that define the traffic to be encrypted must be mirror images of each other on the IPSec peers. If the access list entries include ranges of ports, then a mirror image of those same ranges must be included in the remote peer access lists. The addresses specified in the access lists are independent of the addresses used by the IPSec peers.

In the examples below, GRE entries have been specified for both source and destination addresses. The source address is the local router's side of the ISP connection and the destination address is the ISP connection on the other router's interface. All traffic encapsulated in the GRE packets will be protected.

The examples below show the same flow for the ACL. The name of the ACL and the unique addresses are the only differentiating factors.

Head-End

Following is the ACL configuration for VPN-HE-1.

```
ip access-list extended toBranch1
 permit gre host 131.108.1.1 host 131.108.101.1
ip access-list extended toBranch2
 permit gre host 131.108.1.1 host 131.108.102.1
```

Following is the ACL configuration for VPN-HE-2.

```
ip access-list extended toBranch2
 permit gre host 131.108.1.5 host 131.108.102.1
ip access-list extended toBranch1
 permit gre host 131.108.1.5 host 131.108.101.1
```

Branch

Following is the ACL configuration for VPN-Branch-1.

```
ip access-list extended toHE-1
 permit gre host 131.108.101.1 host 131.108.1.1
ip access-list extended toHE-2
 permit gre host 131.108.101.1 host 131.108.1.5
```

Following is the ACL configuration for VPN-Branch-2.

```
ip access-list extended toHE-2
 permit gre host 131.108.102.1 host 131.108.1.5
ip access-list extended toHE-1
 permit gre host 131.108.102.1 host 131.108.1.1
```

Crypto Map Configuration

The crypto map entry ties together the IPSec peers, the transform set used, and the access list used to define the traffic to be encrypted. The crypto map entries are evaluated sequentially.

In the example below, the crypto map name “static-map” and crypto map numbers (for example, “1” and “2”) are locally significant only. The first statement sets the IP address used by this peer to identify itself to other IPSec peers in this crypto map. This address must match the set peer statement in the remote IPSec peers' crypto map entries. This address must also match the address used with any preshared keys the remote peers might have configured. The IPSec mode defaults to tunnel mode.

Head-End

Following is the crypto map configuration for VPN-HE-1.

```
interface FastEthernet0/1
 description to ISP for VPN
 ip address 131.108.1.1 255.255.255.252
!
crypto map static-map local-address
FastEthernet0/1
!
crypto map static-map 1 ipsec-isakmp
 set peer 131.108.101.1
 set transform-set strong
 match address toBranch1

crypto map static-map 2 ipsec-isakmp
 set peer 131.108.102.1
 set transform-set strong
 match address toBranch2
```

Following is the crypto map configuration for VPN-HE-2.

```
interface FastEthernet0/1
 description to ISP for VPN
 ip address 131.108.1.5 255.255.255.252
!
crypto map static-map local-address
FastEthernet0/1
!
crypto map static-map 1 ipsec-isakmp
 set peer 131.108.102.1
 set transform-set strong
 match address toBranch2
crypto map static-map 2 ipsec-isakmp
 set peer 131.108.101.1
 set transform-set strong
 match address toBranch1
```

Branch

Following is the crypto map configuration for VPN-Branch-1.

```
interface Serial0/0
 description to ISP for VPN
 ip address 131.108.101.1 255.255.255.252
!
crypto map static-map local-address Serial0/0
!
crypto map static-map 1 ipsec-isakmp
 set peer 131.108.1.1
 set transform-set strong
 match address toHE-1
crypto map static-map 2 ipsec-isakmp
 set peer 131.108.1.5
 set transform-set strong
 match address toHE-2
```

Following is the crypto map configuration for VPN-Branch-2.

```
interface Serial0
 description to ISP for VPN
 ip address 131.108.102.1 255.255.255.252
!
crypto map static-map local-address Serial0
!
crypto map static-map 1 ipsec-isakmp
 set peer 131.108.1.5
 set transform-set strong
 match address toHE-2
crypto map static-map 2 ipsec-isakmp
 set peer 131.108.1.1
 set transform-set strong
 match address toHE-1
```

Tip A more complete description can be found at: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/fipsencr/srfipsec.htm#xtocid105785

Applying Crypto Maps

The crypto maps must be applied to both the physical interface and the logical interfaces, such as the GRE tunnel interfaces.In the examples below, the delay statements applied to the tunnel interfaces have been weighted so that traffic will prefer the tunnel that has been designated the primary tunnel.

Head-End

Following is the configuration that applies the crypto maps for VPN-HE-1.

```
interface FastEthernet0/1
 description to ISP for VPN
 ip address 131.108.1.1 255.255.255.252
 crypto map static-map
!
interface Tunnel0
 description Primary Tunnel to Branch1
 ip address 10.200.1.1 255.255.255.252
 tunnel source FastEthernet0/1
 tunnel destination 131.108.101.1
 crypto map static-map
!
interface Tunnel1
 description Secondary Tunnel to Branch2
 ip address 10.200.1.13 255.255.255.252
 delay 6000
 tunnel source FastEthernet0/1
 tunnel destination 131.108.102.1
 crypto map static-map
```

Following is the configuration that applies the crypto maps for VPN-HE-2.

```
interface FastEthernet0/1
 description to ISP for VPN
 ip address 131.108.1.5 255.255.255.252
 crypto map static-map
!
interface Tunnel0
 description Primary Tunnel to Branch2
 ip address 10.200.1.9 255.255.255.252
 tunnel source FastEthernet0/1
 tunnel destination 131.108.102.1
 crypto map static-map
!
interface Tunnel1
 description Secondary Tunnel to Branch1
 ip address 10.200.1.5 255.255.255.252
 delay 6000
 tunnel source FastEthernet0/1
 tunnel destination 131.108.101.1
 crypto map static-map
```

Note EIGRP is the routing protocol used in this design. See the *Cisco AVVID Network Infrastructure Enterprise Site-to-Site VPN Design SRND* for alternative ways to influence cost in other IGP routing protocols.

Branch

Following is the configuration that applies the crypto maps for VPN-Branch-1.

```
interface Serial0/0
 description to ISP for VPN
 ip address 131.108.101.1 255.255.255.252
 crypto map static-map
!
interface Tunnel0
 description Primary Tunnel to HE1
 ip address 10.200.1.2 255.255.255.252
 tunnel source Serial0/0
 tunnel destination 131.108.1.1
 crypto map static-map
!
interface Tunnel1
 description Secondary Tunnel to HE2
 ip address 10.200.1.6 255.255.255.252
 delay 60000
 tunnel source Serial0/0
 tunnel destination 131.108.1.5
 crypto map static-map
```

Following is the configuration that applies the crypto maps for VPN-Branch-2.

```
interface Serial0
 description to ISP for VPN
 ip address 131.108.102.1 255.255.255.252
 crypto map static-map
!
interface Tunnel0
 description Primary Tunnel to HE2
 ip address 10.200.1.10 255.255.255.252
 tunnel source Serial0
 tunnel destination 131.108.1.5
 crypto map static-map
!
interface Tunnel1
 description Secondary Tunnel to HE1
 ip address 10.200.1.14 255.255.255.252
 delay 60000
 tunnel source Serial0
 tunnel destination 131.108.1.1
 crypto map static-map
```

Static Route Configuration

At a minimum, each router must have a static host route for each for the other end-point. The head-end peers (primary and secondary) will point to ISP as the gateway to reach the branches. And the branches will point back to the ISP as the gateway to reach the head-end routers.

Following is the configuration of static routes on VPN-Branch-1.

```
! Route anything going to VPN-HE-1 and VPN-HE-2 to the ISP
ip route 131.108.1.1 255.255.255.255 131.108.101.2
ip route 131.108.1.5 255.255.255.255 131.108.101.2
```

Tip See the *Cisco AVVID Network Infrastructure Enterprise Site-to-Site VPN Design SRND* for more on static and dynamic routing with VPN.

Multicast VPN Deployment Recommendations

This chapter discusses the recommended and optional configurations for IP multicast WAN deployment. The recommended guidelines are summarized below:

- Use IP Multicast to scale streaming applications, such as Music-on-Hold and IP/TV.
- Use administratively scoped addresses to differentiate multicast applications by type and bandwidth.
- Use Anycast RP.
- Select Catalyst switches that have IGMP snooping and use CGMP in low-end switches that do not support IGMP snooping.
- Use the recommended commands to ensure that the correct RPs and sources are used.
- Use IP Multicast boundaries to control where certain multicast streams go.
- Use "show" commands to ensure proper operation of the multicast configurations and enable SNMP traps to log multicast events.

Multicast Site-to-Site VPN Deployment

The VPN topology, shown in Figure 6-4, is based on a dual-VPN head-end model.

There are a few different design scenarios that can be used with IP Multicast when using a site-to-site VPN deployment. The design in this chapter uses RPs that are located in the campus network. IP multicast configuration for the campus is covered in Chapter 2, "IP Multicast in a Campus Network." In this chapter, a brief configuration for the campus RPs is shown for reference.

For IP multicast forwarding, the branch sites use the appropriate VPN head-end router to which they have a primary tunnel connection. If there is a failure in the VPN head-end router, the branch will use the secondary tunnel configuration to receive IP multicast traffic and will receive an update stating which device to use as the RP for the specified groups.

Figure 6-4 Multicast VPN Topology

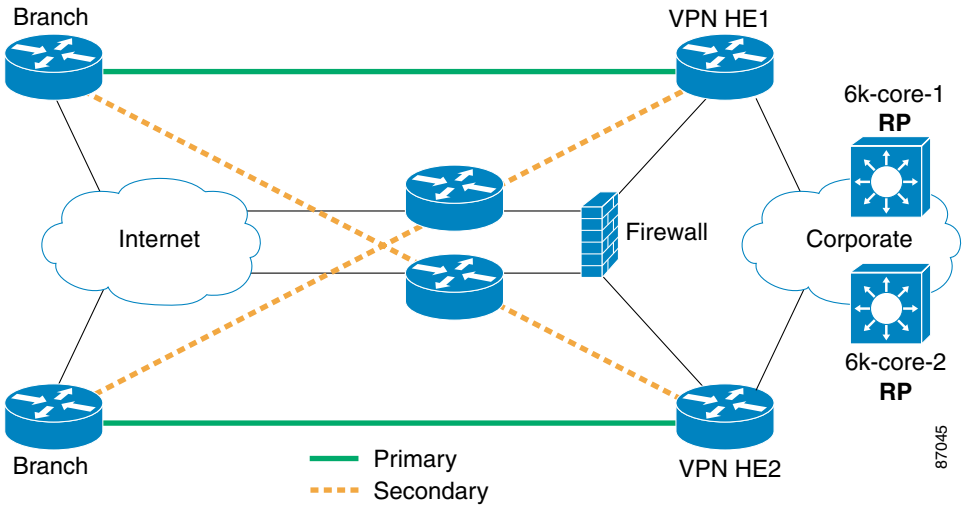

This section provides information and sample configurations for deploying Anycast RP in a WAN.

Branch and Head-End

The IP multicast configuration steps for the branch and head-end routers are as follows:

- Enable IP multicast routing.
- Enable IP PIM Sparse Mode on each interface that will receive and forward multicast traffic.

Identify the RP for the router. With Anycast RP, the address defined is the Loopback address that is duplicated on each participating Anycast RP router. (Loopback 1 - 10.6.2.1).

Branch

Following is the basic multicast configuration for Branch1.

```
ip multicast-routing
!
interface Tunnel0
 description Primary Tunnel to HE1
 ip pim sparse-mode
!
interface Tunnel1
 description Secondary Tunnel to HE2
 ip pim sparse-mode
!
interface FastEthernet0/0.1
 description DATA VLAN 4
 encapsulation dot1Q 4
 ip address 40.1.4.1 255.255.255.0
 ip pim sparse-mode
!
```

```
interface FastEthernet0/0.2
 description VOICE VLAN 40
 encapsulation dot1Q 40
 ip address 40.1.40.1 255.255.255.0
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1                    Loopback 1 address of both RPs.
```

Following is the basic multicast configuration for Branch2.

```
ip multicast-routing
!
interface Tunnel0
 description Primary Tunnel to HE2
 ip pim sparse-mode
!
interface Tunnel1
 description Secondary Tunnel to HE1
 ip pim sparse-mode
!
interface FastEthernet0/0.1
 description DATA VLAN 5
 encapsulation dot1Q 5
 ip address 50.1.5.1 255.255.255.0
 ip pim sparse-mode
!
interface FastEthernet0/0.2
 description VOICE VLAN 50
 encapsulation dot1Q 50
 ip address 50.1.50.1 255.255.255.0
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1                    Loopback 1 address of both RPs.
```

Head-End

Following is the basic multicast configuration for VPN-HE-1.

```
ip multicast-routing
!
interface Tunnel0
 description Primary Tunnel to Branch1
 ip pim sparse-mode
!
interface Tunnel1
 description Secondary Tunnel to Branch2
 ip pim sparse-mode
!
interface FastEthernet0/0
 description to Corporate Network
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1                    Loopback 1 address of both RPs.
```

Following is the basic multicast configuration for VPN-HE-2.

```
ip multicast-routing
!
interface Tunnel0
 description Primary Tunnel to Branch2
 ip pim sparse-mode
!
interface Tunnel1
 description Secondary Tunnel to Branch1
 ip pim sparse-mode
!
interface FastEthernet0/0
 description to Corporate Network
 ip pim sparse-mode
!
ip pim rp-address 10.6.2.1                    Loopback 1 address of both RPs.
```

Note Ensure that the IP subnets for the GRE tunnels are advertised in the IGP routing protocol. If not, RPF failures will occur.

Summary

In summary, when using IP multicast with MoH or IP/TV in the WAN follow these recommendations.

- Use PIM-SM (recommended)
- Anycast RP (recommended)
  - Fast convergence, redundancy, load balancing

Multicast Music-on-Hold and IP/TV Configurations

This chapter discusses the basics of MMoH and IP/TV implementation. Detailed configuration and design for both MoH and IP/TV are not covered; only the basic layout is needed to enable multicast operation.

Multicast Music-on-Hold

Call Manager version 3.1 and higher allows for the configuration of MoH services. Unicast and Multicast MoH streams can be generated by two source types.

- A file source uses actual files located on the MoH server to “play” or stream audio music out to the network.
- A fixed source uses a soundcard located in the MoH server. The soundcard can connect a wide-range of devices, such as CD and cassette players.

There are two types of hold that can be used with MoH.

- Network hold is used when features, such as transfer, conference, and call-park, are used.
- User hold is used when the IP phone user selects the “hold” softkey.

The following are guidelines for implementing MoH with IP multicast:

- Use an IP multicast address group range from the administratively scoped address range.
- Enable multicast on the MoH server and enter the starting address range selected from the scoped range.
- Increment the multicast streams on IP addresses.
- Use the G.711 CODEC on the LAN and the G.729 CODEC for streams crossing the WAN.
- Change the default “Max Hops” or TTL to a value representative of the routed network.
- Enable multicast for each of the selected audio sources.
- Enable the user and network hold sources for each IP phone.
- Over provision the Low-Latency Queue (LLQ) by one or two streams to account for the additional traffic brought on by MoH.

It is important to understand the basic operation of how the MoH server provides music to the IP phone. Figure 7-1 illustrates a MoH setup.

Figure 7-1 Multicast Music-on-Hold Operation

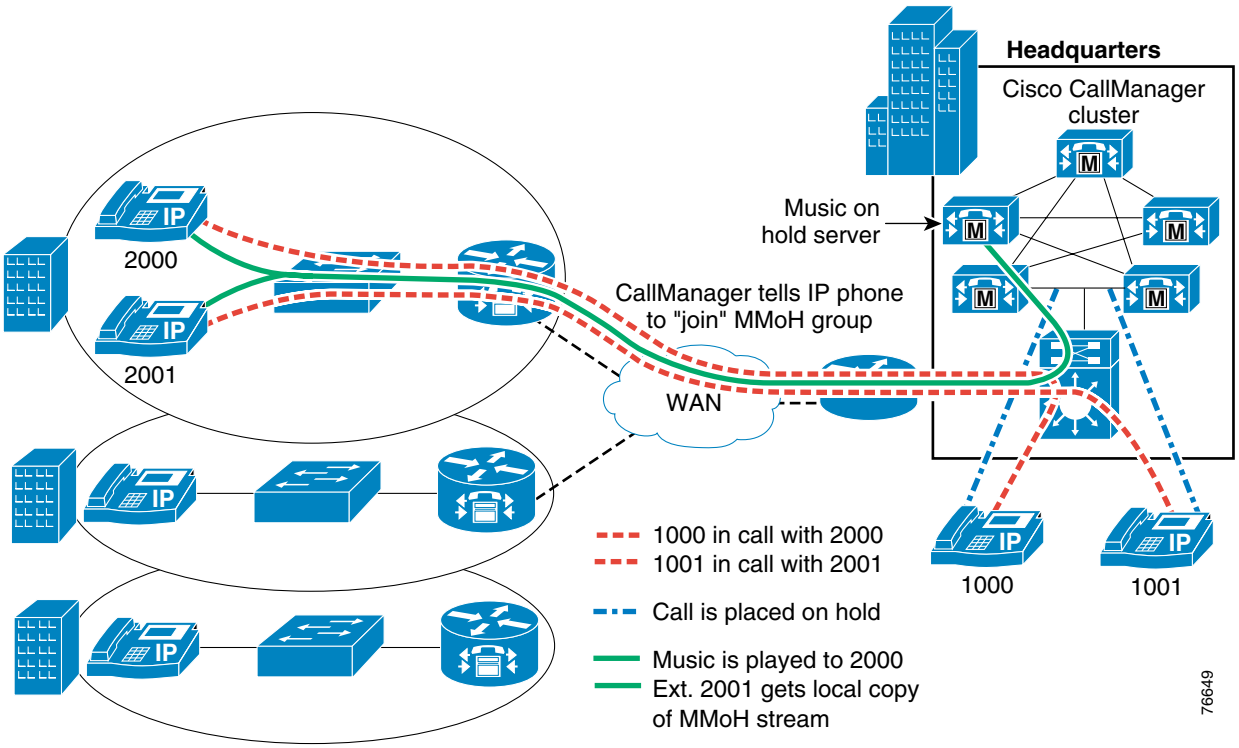

Step 1 Extension 1000 is in a call with Extension 2000 (branch office).

Step 2 Extension 1000 places Extension 2000 on hold. Call Manager signals the IP phone (Extension 2000) via the Skinny protocol to perform an IGMP join for the Multicast MoH group.

Step 3 The IP phone (Extension 2000) signals to the network that it wishes to receive the multicast stream associated with the MoH group number. PIM is used to forward the stream from the first-hop router for the MoH server, through the WAN, and on to the IP phone at the branch.

Step 4 Extension 1001 is on a call with Extension 2001 (same branch office as Extension 2000).

Step 5 Extension 1001 places Extension 2001 on hold. Call Manager signals the IP phone (Extension 2001) via the Skinny protocol to perform an IGMP join for the same Multicast MoH group.

Note It is important to make sure that the audio source, region, location and Media Resource Group List (MRGL) are configured properly for the branch office. If not, the wrong MoH audio CODEC may be used and a separate audio stream will be sent to the branch. This defeats the purpose of saving bandwidth with multicast.

Step 6 The IP phone (Extension 2001) signals the network that it wishes to receive the multicast stream associated with the MoH group number. The local switch and router know, via CGMP/PIM or IGMP snooping, that an existing stream is already being forwarded to Extension 2000. So a copy of the existing stream is forwarded from the local branch office switch.

Step 7 When the holding phone resumes the call, the held phone sends an IGMP leave to signal that it no longer wants the multicast traffic.

Increment Multicast on IP Address

If the MoH server is configured with "Increment Multicast on IP Address," the starting address is for the G.711ulaw CODEC and the addresses are incremented by 1 for each additional CODEC.

As a result, for every file that is added as a multicast audio source, four addresses are consumed (one for every CODEC). The current maximum for audio sources on a single MoH server is 51 (50 file sources and 1 fixed source). Resulting in a maximum requirement of 204 multicast addresses (51 files each with 4 CODECs).

Even so, it is better to use "Increment Multicast on IP Address" and not "Increment Multicast on Port Number" because IP multicast routers understand (S,G) notations and not port numbers.

For example, let's assume that "Increment Multicast on Port Number" is configured for audio sources and G.729 is to be used for the branch office phones. Then, when the branch office IP phone issues an IGMPv2 Join for the group associated with the audio source, all four streams are sent to the branch. Because the four streams have the same group address (different port numbers are ignored), the router at the branch pulls all four streams even though only one stream is needed.

The four streams together use a bandwidth total of about 480 Kbps.

- G.711ulaw =77Kbps
- G.711alaw= 77Kbps
- G.729 = 26Kbps
- Wideband = almost 300Kbps

The result could be a terrible load on the network, given that there could be up to 51 sources streaming 480Kbps each.

Note The recommendation is to use the G.711ulaw or Wideband CODEC in the campus network. G.729 consumes less bandwidth than the G.711 or Wideband CODEC and is recommended for use in the WAN. Using G.729 will conserve bandwidth, but the audio quality is poor.

Table 7-1 lists the maximum number of MoH sessions for each model of server. The maximum sessions refers to unicast, multicast, or unicast and multicast sessions. This is the recommended maximum sessions a platform can support, irrespective of the transport mechanism.

Table 7-1 Maximum Number of MoH Sessions.

Server Platform	CODECs Supported	MoH Sessions
MCS 7815 MCS 7825 (800/1133) MCS 7830 SPE-310 Compaq DL310 Compaq DL320 IBM xSeries 330	G.711 (A-law and µ -law) G.729aWideband Audio	20 MoH sessions (co-resident) 50 MoH sessions (Stand-alone MoH Server)
MCS 7835 (1000/1266) Compaq DL380 IBM xSeries 340	G.711 (A-law and µ -law) G.729aWideband Audio	20 MoH sessions (co-resident) 250 MoH sessions (Stand-alone MoH Server)

Multicast MoH Configuration

This section provides an example for configuring multicast MoH, including:

- Configuring the MoH Server for Multicast
- Configuring the MoH Audio Source
- Configuring the IP Phones
- Changing the Default CODEC
- Verifying the Configuration
- QoS for Music-on-Hold

Configuring the MoH Server for Multicast

To configure the MoH server for Multicast, do the following:

Step 1 From the Call Manager Admin main page, select **Service >Music On Hold > Configure Music On Hold Servers**. Figure 7-2 shows the MoH Server Configuration screen.

Figure 7-2 Call Manager—MoH Server Configuration

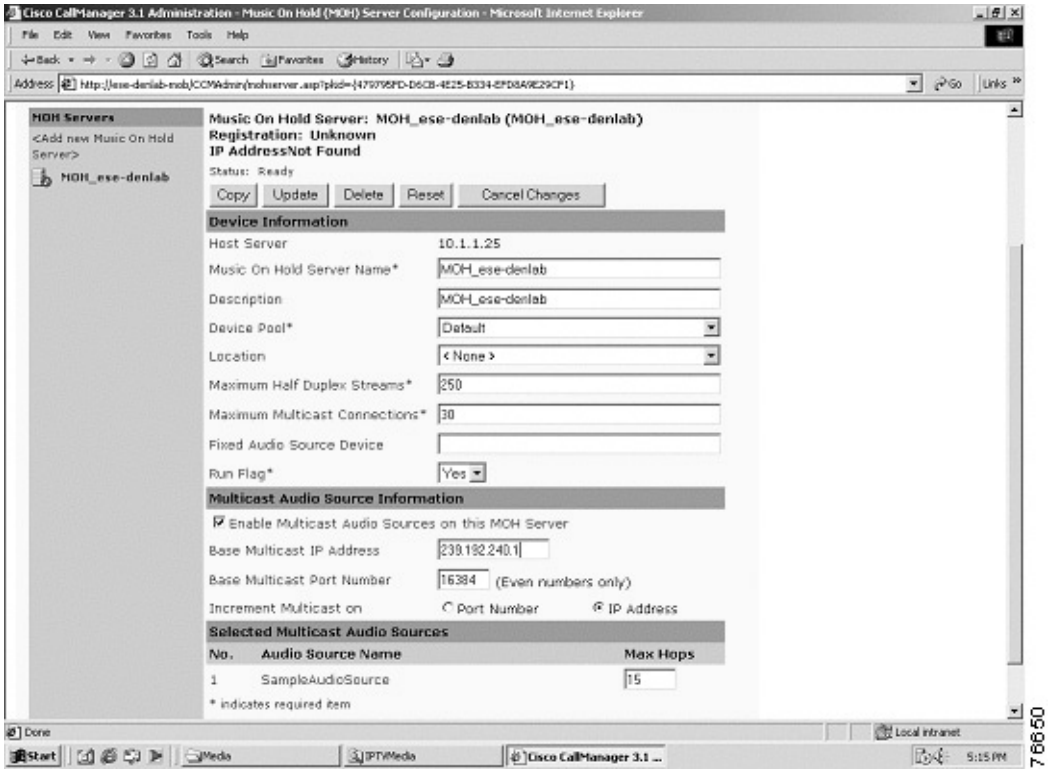

Step 2 Select the desired server from the list on the left.

Step 3 Under the "Multicast Audio Source Information" section, select **Enable Multicast Audio Sources on this MoH Server**.

Step 4 Change the Base Multicast IP Address to the range identified from the administratively scoped range.

Step 5 Ensure that **Increment Multicast on IP Address** is selected. The default depends on version of Call Manager.

Step 6 Change the default max hops (TTL) to reflect the routed network.

Step 7 Click on **Update** to save the changes.

Note A warning message will be displayed if the MoH Server has not been associated with a Media Resource Group (MRG) and Media Resource Group List (MRGL). The MoH Server should be associated with the appropriate MRG and MRGL before the source is assigned to the IP phone. MRG and MRGL configurations are not covered in this document as this involves the configurations of other Media Resources such as conference bridges.

Configuring the MoH Audio Source

To configure the MoH audio source for Multicast, do the following:

Step 1 From the Call Manager Admin main page, select **Service >Music On Hold**. Figure 7-3 shows the MoH Audio Source Configuration screen.

Figure 7-3 Call Manager—MoH Audio Source Configuration

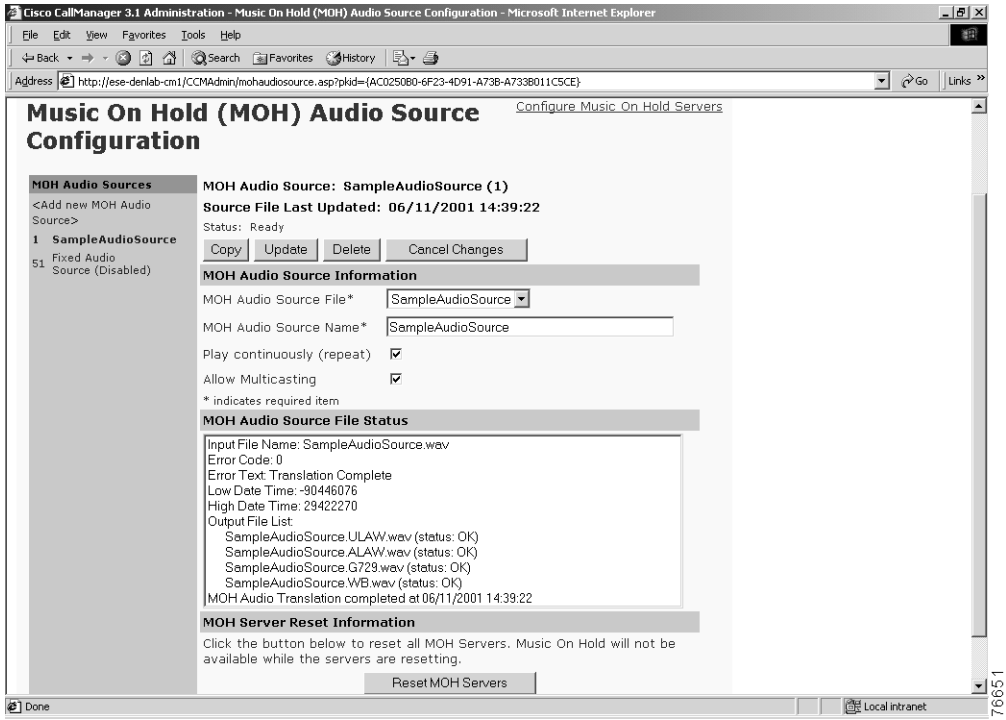

Step 2 Select the desired audio source from the list on the left.

Step 3 Select **Allow Multicasting**.

Step 4 Click on **Update** to save the changes.

Step 5 A message is displayed indicating that the MoH Servers must be reset. To reset the servers, click on **Reset MoH Servers**.

Configuring the IP Phones

To configure the IP Phones for Multicast MoH, do the following:

Step 1 From the CallManager Admin main page, select **Device>Phone>Find** and enter the search criteria. Figure 7-4 shows the Phone Configuration screen.

Figure 7-4 Call Manager—Phone Configuration

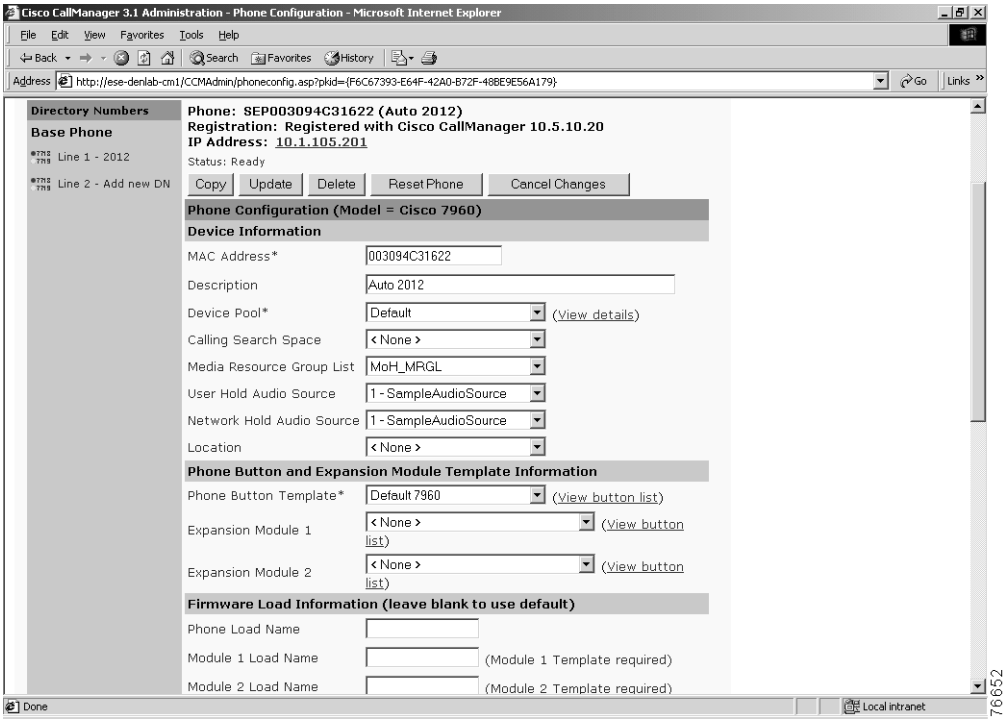

Step 2 Select the desired phone from the list.

Step 3 For the "User Hold Audio Source," select the configured audio source (see Configuring the MoH Audio Source) from the drop-down list.

Step 4 For the "Network Hold Audio Source," select the audio source (see Configuring the MoH Audio Source) from the drop-down list. This may be a different audio source than the one used for the user hold.

Step 5 If there is configured a MRGL for this device or audio source, select that also.

Note A different user and network hold source can be configured for each Directory Number (DN) listed for this device.

Step 6 Click on **Update** to save the changes.

Changing the Default CODEC

To change the default CODEC used for streaming MoH, do the following:

Step 1 From the Call Manager Admin main page, select **Service >System Parameters >Cisco IP Voice Media Streaming App>DefaultMOHCodec**.

Step 2 Change the defaults as desired,

Step 3 Click on **Update** to save the changes.

Verifying the Configuration

After placing a call on hold that has been configured with a multicast-enabled audio source, run the show **ip mroute active** command to see the Multicast MoH stream at the branch office. For example:

```
show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps

Group: 239.192.240.3, (?)                    Multicast MoH - G.729 address
   Source: 10.5.10.20 (CM-MoH)
     Rate: 50 pps/24 kbps(1sec), 23 kbps(last 53 secs), 22 kbps(life avg)
```

QoS for Music-on-Hold

Although this document does not go into detail on the QoS features and configurations for IP Telephony and Music-on-Hold, there is one recommendation for MoH that should be mentioned here:

- Over provision the Low-Latency Queue (LLQ) by one or two streams to account for the additional traffic brought on by MoH.

Multicast MoH is classified as RTP bearer traffic (DSCP=EF). MoH streams will inherently take advantage of QoS deployed on the network for IP Telephony.

Tip For information about configuring QoS for IP Telephony, see the *Cisco AVVID Network Infrastructure Enterprise Quality of Service Design* guide.

IP/TV Server

The Cisco IP/TV solution is used to serve media streams for live broadcast, stored broadcast, and on-demand content. The server used in this design is streaming out to multicast groups using pre-recorded files. IP/TV servers provide streaming of Video, Audio, and Slidecast (presentations) media.

The following are guidelines for implementing IP/TV with IP multicast:

- Use multicast addresses found in the administratively scoped address range.
- Associate multicast group ranges for each type of stream.

- Use IP multicast boundaries to restrict the forwarding of certain streams to selected areas of the network.
- Configure QoS rules for the IP/TV server at the ingress point in the campus.
- Use the recommended DSCP of CS4.

Multicast IP/TV Configuration

This section provides an example configuration for multicast IP/TV. Table 7-2 lists the stream, name, bandwidth consumption, multicast address, and scope assignment used in this example for the IP/TV server and multicast MoH traffic.

Table 7-2 IP/TV Stream Sample Layout

Stream	Name	Bandwidth	Multicast Groups /22	Scope
High-Rate	1.5Mbps-Local	1.5Mbps	239.255.0.1 - Video 239.255.0.2 - Audio 239.255.0.3 - Slidecast	Site-Local
Medium	256k-Local-Branch -768k	256k	239.192.248.1 - Video 239.192.248.2 - Audio 239.192.248.3 - Slidecast	Organization-Local
Low	100k-Local-Branch -768-256k	100k	239.192.244.1 - Video 239.192.244.2 - Audio 239.192.244.3 - Slidecast	Organization-Local

Following is an example IP/TV configuration for a WAN aggregation router (Frame-relay). In this example:

- The high-rate stream is configured to stream in the campus network and stop at the WAN aggregation router. If satellite feeds or high-speed WAN links exist, the High-Rate stream can pass to those areas of the network. However, the design discussed in this section prevents this to illustrate how IP multicast boundaries and filters can be used to keep certain traffic from passing over the WAN.
- The medium stream is configured to stream in the campus and links 768kbps and higher.
- The low stream is configured to stream in the campus and links 256kbps and higher.
- Multicast MoH is configured to stream to all locations.
- The group range for Multicast MoH is added for full configuration reference (239.192.240.0 /22).

```
interface Serial4/0.1 point-to-point
 description To leftrtr - 768k
 ip multicast boundary medium-low-moh

ip access-list standard medium-low-moh
 remark Deny High-rate (.255)
 deny   239.255.0.0 0.0.255.255
 permit any
```

```
interface Serial4/0.2 point-to-point
 description To middlertr - 256k
 ip multicast boundary low-moh

ip access-list standard low-moh
 remark Deny High-rate (.255), Medium (.248)
 deny   239.255.0.0 0.0.255.255
 deny   239.192.248.0 0.0.3.255
 permit any

interface Serial4/0.3 point-to-point
 description To rightrtr - 128k
 ip multicast boundary moh

ip access-list standard moh
 remark Deny High-rate (.255), Medium (.248), Low (.244)
 deny   239.255.0.0 0.0.255.255
 deny   239.192.248.0 0.0.3.255
 deny   239.192.244.0 0.0.3.255
 permit any
```

Note For more information about the boundaries, see the "Traffic Engineering" section on page 8-4.

To verify the configuration, first use the IP/TV Viewer client to request the stream at the branch office. Then use the **show ip mroute active** command. The following example shows the low stream (100k).

```
show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps

Group: 239.192.244.2, (?)                    Low-rate stream - audio group address
   Source: 10.5.11.20 (IPTV)
     Rate: 6 pps/12 kbps(1sec), 16 kbps(last 17 secs), 20 kbps(life avg)


Group: 239.192.244.1, (?)                    Low-rate stream - video group address
   Source: 10.5.11.20 (IPTV)
     Rate: 20 pps/90 kbps(1sec), 94 kbps(last 23 secs), 115 kbps(life avg)
```

QoS for IP/TV Server

Although this document does provide full details on the QoS features and configurations for IP/TV, a basic recommendation and example configuration is given for reference.

The recommended QoS guidelines for IP/TV are:

- Classify streaming media with a DSCP of CS4 (which maps, by default, to a COS of 4).
- Configure the ingress port on the switch that connects to the IP/TV server to "tag" the streaming media with the recommended value of CoS=4. The IP/TV server will require a 801.2Q trunk connection.

The following example illustrates the recommended configuration for IP/TV streaming media. On the server-farm switch, configure the port connected to the IP/TV server to classify all traffic from the source as CoS=4 and to override any CoS value previously set.

```
interface GigabitEthernet0/1
 description To IP/TV Server
 mls qos cos 4                          Assign a CoS of 4 to any traffic arriving on this port
 mls qos cos override                   Override previously classified CoS value with the
.                                       one set above
```

Because a COS of 4 maps, by default, to a DSCP of 32 (CS4), the **mls qos map** command is optional for the IP/TV server-farm switch. Also, on the server-farm switch there is no need to remap any of the CoS values for Voice (CoS 5 by defalut maps to DSCP 46) or Voice-Control (CoS 3 by default maps to DSCP 26). All of these should be left at defaults for this example.

To verify that the default mappings are in place, issue the following command:

```
show mls qos maps cos-dscp

Cos-dscp map:
cos:  0  1  2  3   4   5   6   7
--------------------------------
dscp: 0  8  16 24  32  40  48 56
```

Once this mapping is complete, the IP/TV traffic can be scheduled and provisioned according to the CoS/DSCP values identified above.

Summary

In summary, when using IP multicast with MoH or IP/TV follow these recommendations.

Multicast Music-on-Hold

- Increment on IP Address (recommended)
- Change max Hops to reflect specific routing topology (recommended)
- Allow G.729 Multicast Audio Source to Branch offices (recommended)
- Restrict G.711 and Wideband to local Campus (recommended)
- Provision LLQ appropriately to account for Multicast MoH streams in QoS Policy (recommended)

IP/TV Server

- Use administrative scoping for different streams sizes (recommended)
- Boundary streams by using the scoped group address (recommended)
- Classify IP/TV streaming media at the ingress port in the server-farm switch
  - Use QoS markings of CoS=4 and DSCP=CS4 (recommended)

CHAPTER 8

Security, Timers, and Traffic Engineering in IP Multicast Networks

This chapter provides recommendations for security measures, timer adjustments, and traffic engineering for an IP multicast network.

Security

With IP multicast, it is important to protect the traffic from Denial-of-Service (DoS) attacks or stream hijacking by rogue sources or rogue RPs.

- DoS attacks affect the availability and efficiency of a network. If a rogue application or device can generate enough traffic and target that traffic at a source, then CPU and memory resources can be severely impacted.
- Stream hijacking allows any host on the network to become an active source for any legitimate multicast group. It is easy to download a free multicast-enabled chat application from the Internet and change the IP Multicast group address assignment to be the same as that used by legitimately configured multicast applications. If the network devices are not secured from "accepting" unauthorized sources, the rogue source can impact the IP Multicast streams. For the most part, receivers are ignorant of the details associated with which source is really responsible for which group.

Use the following commands on IP Multicast-enabled routers to guard against rogue sources and rogue RPs:

- **ip pim accept-register**
- **rp-announce-filter**
- **ip pim rp-address**
- **ip igmp access-group**

Rogue Sources

A source is any host that is capable of sending IP Multicast traffic. Rogue sources are unauthorized sources that send IP Multicast traffic.

Sources send group traffic to the first-hop router. The first-hop router sends a Register message to the RP with information about the active source. To protect the router from unauthorized Register messages, use the **ip pim accept-register** command. This command, which can be used only on candidate RPs, configures the RP to accept Register messages only from a specific source. If a Register message is denied, a Register-Stop is sent back to the originator of the Register.

- If the **list** *acl* attribute is used, extended access lists can be configured to determine which pairs (source and group) are permitted or denied when seen in a Register message.
- If the **route-map** *map* attribute is used, typical route-map operations can be applied on the router for the source address that appears in a Register message.

Note The keywords list and route-map cannot be used together.

The following example illustrates a configuration that permits a registration from the sources listed (10.5.10.20 MoH server and 10.5.11.20 IP/TV Server).

```
ip pim accept-register list 101

access-list 101 permit ip host 10.5.10.20 any
access-list 101 permit ip host 10.5.11.20 any
```

Note For more information about the addresses used, see Table 6-1.

Additionally, a list can be configured that indicates which groups are permitted from the sources at time of registration. The following example illustrates a configuration that permits the MoH group address from the MoH server and the three IP/TV groups from the IP/TV server.

```
access-list 101 permit ip host 10.5.10.20 239.192.240.0 0.0.3.255
access-list 101 permit ip host 10.5.11.20 239.192.244.0 0.0.3.255
access-list 101 permit ip host 10.5.11.20 239.192.248.0 0.0.3.255
access-list 101 permit ip host 10.5.11.20 239.255.0.0 0.0.255.255
```

If an unauthorized source comes online and the first-hop router attempts to register this new source with the RP, the registration will be rejected. The following example shows the debug output for a failed registration attempt by router 10.0.0.37 for source 10.5.12.1 and group 239.194.1.1.

```
1d03h: PIM: Received v2 Register on Vlan59 from 10.0.0.37
1d03h:      (Data-header) for 10.5.12.1, group 239.194.1.1
1d03h: PIM Register for 10.5.12.1, group 239.194.1.1 rejected
1d03h: PIM: Send v2 Register-Stop to 10.0.0.37 for 10.5.12.1, group 239.194.1.1
```

Note The streams sent by rogue sources would flow on the local subnet where the source resides. In addition to not blocking the source on the local subnet, there are other topological cases where the "accept-register" mechanism fails to block rogue sources.

Rogue RPs

A rogue RP is any router that, by mistake or maliciously, acts as an RP for a group. To guard against maliciously configured routers acting as a candidate RP, use the following commands:

- The **ip pim rp-announce-filter** command is used on Mapping Agents to filter Auto-RP announcement messages coming from the RP. This command can be used only when Auto-RP is deployed.

  In the following example, the router is configured to accept RP announcements from RPs in access list 11 for group ranges described in access list 12.

```
ip pim rp-announce-filter rp-list 11 group-list 12

access-list 11 permit 10.6.2.1                    IP address of permitted RP
access-list 12 permit 239.192.240.0 0.0.3.255     Permit MoH
access-list 12 permit 239.192.244.0 0.0.3.255     Permit low stream
access-list 12 permit 239.192.248.0 0.0.3.255     Permit medium stream
access-list 12 permit 239.255.0.0 0.0.255.255     Permit high stream
access-list 12 deny 239.0.0.0 0.255.255.255       Deny remaining administratively scoped range
access-list 12 permit 224.0.0.0 15.255.255.255    Permit link local/reserved address
```

- The **ip pim rp-address** command configures the PIM RP address for a particular group or group range. Without an associated group-acl, the default group range is 224.0.0.0/4. The RP address is used by first-hop routers to send Register messages on behalf of source multicast hosts. The RP address is also used by routers on behalf of multicast hosts that want to become members of a group. These routers send Join and Prune messages to the RP. Although the command is not used specifically for security purposes, it does help to ensure that a non-RP router uses the authorized RPs for the network.

  In Chapter 2, "IP Multicast in a Campus Network," a filter was used to control which groups an RP was responsible for. If a source becomes active for a group that is **not** in the ACL for the RP group-list, then there will be no active RP for the newly defined group. This will cause the group to fall into dense-mode. As an extra layer of precaution against configurations mistakes or acts of DoS, an RP should be defined that covers all unused multicast groups. This will ensure that undefined groups have an RP on the network and they will not fall into dense-mode nor will the groups be forwarded. This method, commonly referred to as a "Garbage-can RP" can also be used to "log" attempts by rogue sources and groups to register on the network.

- The **ip igmp access-group** command is applied to an interface to restrict the group ranges to which devices are permitted to become members. The interface will discard Join messages for illegal groups. The following example shows that the members of VLAN 10 are only allowed to join the group *w.x.y.z*.

```
interface Vlan 10
ip igmp access-group 1
!
access-list 1 permit w.x.y.z
```

Note The use of IGMP-based ACLs can become a management issue if widely deployed. If there is a need to restrict which groups that the clients can join, try to restrict the groups on the RP, Mapping Agents, or PIM-enable, first-hop routers. If IGMP-based ACLs have been deployed and a group is added or deleted, the ACLs will have to be reconfigured on each VLAN or physical interface to which the clients are attached.

The streams sent by rogue sources would flow on the local subnet where the source resides. Also, depending on the topological layout of the network, the accept-register feature may not block all sources.

Adjusting Timers for IP Multicast

Two timers can be adjusted to facilitate faster failover of multicast streams. The timers control the:

- Query Interval
- Announce Interval

Query Interval

The **ip pim query-interval** command configures the frequency of PIM Router-Query messages. Router Query messages are used to elect a PIM DR. The default value is 30 seconds. If the default value is changed, the recommended interval is 1 second.

To verify the interval for each interface, issue the **show ip pim interface** command, as shown below.

```
3550-svrL-dist#show ip pim interface

Address          Interface                Version/Mode    Nbr    Query     DR
                                                          Count  Intvl
10.5.10.1        Vlan10                   v2/Sparse        0     1       10.5.10.1
10.0.0.37        GigabitEthernet0/1       v2/Sparse        1     1       10.0.0.38
10.0.0.41        GigabitEthernet0/2       v2/Sparse        1     1       10.0.0.42
```

Announce Interval

The **ip pim send-rp-announce** command has an interval option. Adjusting the interval allows for faster RP failover when using Auto-RP. The default interval is 60 seconds and the holdtime is 3 times the interval. So the default failover time is 3 minutes. The lower the interval, the faster the failover time. Decreasing the interval will increase Auto-RP traffic but not enough to cause any kind of a performance impact. If the default is to be changed, use the recommended values of 3 to 5 seconds.

Traffic Engineering

Traffic engineering adds a great deal of control to IP multicast deployment and operation. It also adds complexity. One option for controlling IP multicast is through the use of scoped boundaries.

The **ip multicast boundary** command configures an administratively scoped boundary on an interface and permits or denies multicast group addresses found in the access-list. No multicast packets will be allowed to flow across the boundary from either direction. This allows reuse of the same multicast group address in different administrative domains.

If the RPF interface for a multicast route has a multicast boundary configured for that group, its outgoing interfaces will not be populated with multicast forwarding state. Join messages received on other interfaces will be ignored as long as the boundary remains on the RPF interface. If the RPF interface changes and the boundary no longer applies to the new RPF interface, join latency will be introduced because of the delay in populating outgoing interfaces.

The boundary controls traffic based on the permit/deny configuration of the ACL associated with the boundary command.

The following example shows a boundary on VLAN2 connected to an access layer. The boundary shown permits the group range 239.255.0.0 and denies all other streams on the VLAN2 interface.

```
interface VLAN2
 description To Access VLAN 2
 ip multicast boundary moh
!
ip access-list standard moh
 remark Permit 239.255, Deny all others
 permit 239.255.0.0 0.0.255.255
 deny any
```

Figure 8-1 shows a high-level view of how scoped boundaries can be used to restrict traffic to certain areas of the network. The "Allowed Scopes" list the permitted streams at each location.

Figure 8-1 Graphical View of Administrative Scoping Used with Boundaries

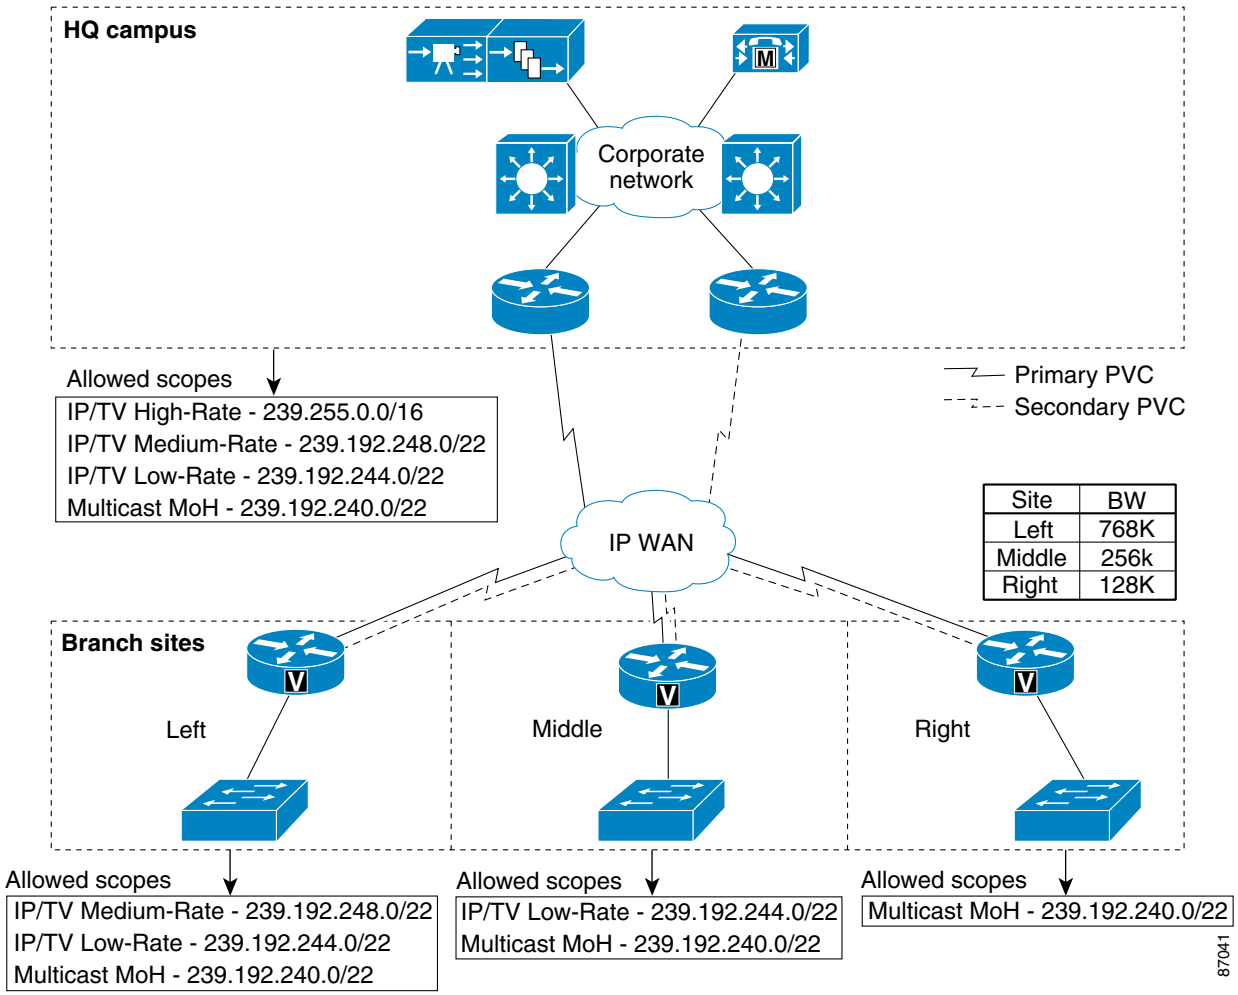

The following is an example configuration for a WAN aggregation router (Frame-relay). The boundary is placed on the serial sub-interfaces that connect to each branch office. Boundary commands could also be configured on the VLAN/interface from the core-layer switches to the WAN aggregation routers.

```
interface Serial4/0.1 point-to-point
 description To leftrtr - 768k
 ip multicast boundary medium-low-moh          Enable a multicast boundary for the 768kbps link

ip access-list standard medium-low-moh
 remark Deny High-rate (.255)
 deny   239.255.0.0 0.0.255.255
 permit any

interface Serial4/0.2 point-to-point
 description To middlertr - 256k
 ip multicast boundary low-moh                 Enable a multicast boundary for the 256kbps link

ip access-list standard low-moh
 remark Deny High-rate (.255), Medium (.248)
 deny   239.255.0.0 0.0.255.255
 deny   239.192.248.0 0.0.3.255
 permit any

interface Serial4/0.3 point-to-point
 description To rightrtr - 128k
 ip multicast boundary moh                     Enable a multicast boundary for the 128kbps link

ip access-list standard moh
 remark Deny High-rate (.255), Medium (.248), Low (.244)
 deny   239.255.0.0 0.0.255.255
 deny   239.192.248.0 0.0.3.255
 deny   239.192.244.0 0.0.3.255
 permit any
```

Managing IP Multicast

To assist in management of the IP multicast, routers can be enabled to send SNMP traps (with IP multicast, MSDP, and PIM information) to the SNMP server. To enable the traps, use the following commands:

- **snmp-server enable traps ipmulticast**
- **snmp-server enable traps msdp**
- **snmp-server enable traps pim**

Tip Details for each of the MIBs for IP multicast can be found at: ftp://ftpeng.cisco.com/ipmulticast/config-notes/mib-info.txt

Layer 2 switches can be enabled to send IGMP traps.

Table 9-1 lists the show commands and tools that can be used to manage and monitor IP multicast configurations and traffic.

Table 9-1 Show Commands and Tools

Layer 2 Switch	Layer 3 Switch	MSDP	Tools
show ip igmp profile	show ip mcache	show ip msdp count	mstat
show ip igmp snooping	show ip mpacket	show ip msdp peer	mtrace
show igmp mode	show ip mroute	show ip msdp sa-cache	mrinfo
show igmp querier information	show ip pim interface	show ip msdp summary	
show igmp statistics	show ip pim neighbor		
show cgmp	show ip pim rp		
show mls multicast	show mls rp ip multicast		
show multicast group	show ip rpf		
show multicast group count			
show multicast protocols status			
show multicast router			

Note The Layer 2 Switch commands are platform dependent.