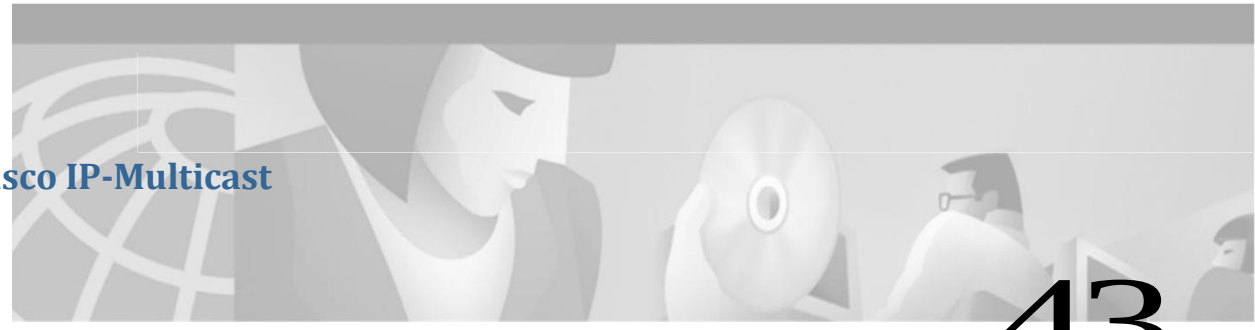

CHAPTER 43

Chapter Goals

- Explain IP multicast addressing.
- Learn the basics of Internet Group Management Protocol (IGMP).
- Explain how multicast in Layer 2 switching works.
- Define multicast distribution trees.
- Learn how multicast forwarding works.
- Explain the basics of protocol-independent multicast (PIM).
- Define multiprotocol BGP.
- Learn how Multicast Source Discovery Protocol (MSDP) works.
- Explain reliable multicast: PGM.

Internet Protocol Multicast

Background

Internet Protocol (IP) multicast is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients and homes. Applications that take advantage of multicast include videoconferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

IP Multicast delivers source traffic to multiple receivers without adding any additional burden on the source or the receivers while using the least network bandwidth of any competing technology. Multicast packets are replicated in the network by Cisco routers enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols resulting in the most efficient delivery of data to multiple receivers possible. All alternatives require the source to send more than one copy of the data. Some even require the source to send an individual copy to each receiver. If there are thousands of receivers, even low-bandwidth applications benefit from using Cisco IP Multicast. High-bandwidth applications, such as MPEG video, may require a large portion of the available network bandwidth for a single stream. In these applications, the only way to send to more than one receiver simultaneously is by using IP Multicast. Figure 43-1 demonstrates how data from one source is delivered to several interested recipients using IP multicast.

Figure 43-1 Multicast Transmission Sends a Single Multicast Packet Addressed to All Intended Recipients

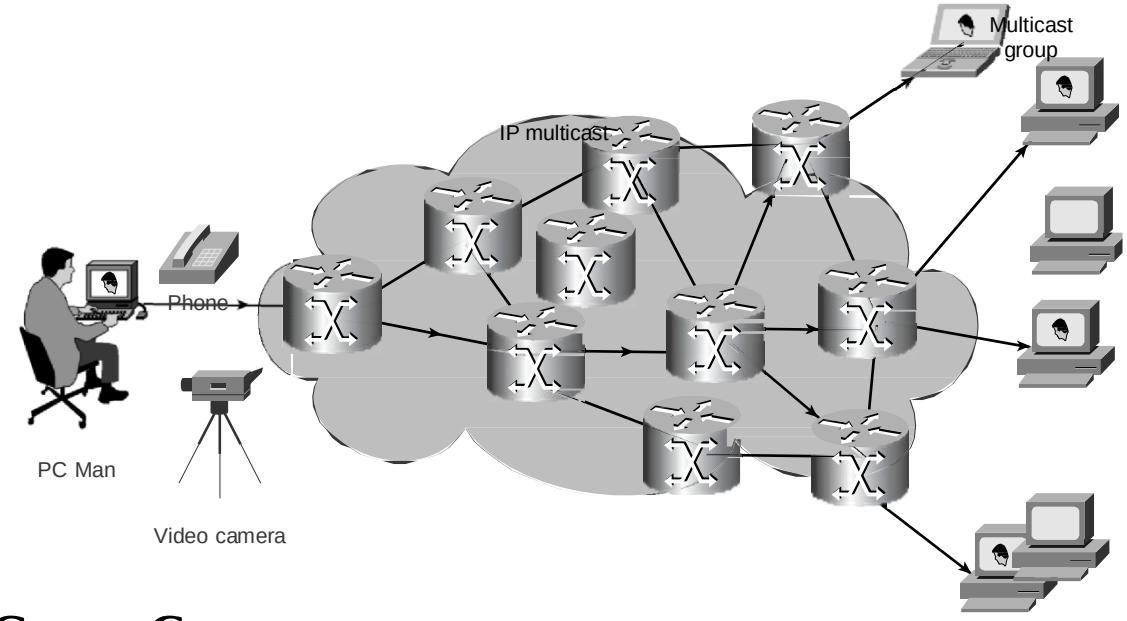

Multicast Group Concept

Multicast is based on the concept of a group. An arbitrary group of receivers expresses an interest in receiving a particular data stream. This group does not have any physical or geographical boundaries—the hosts can be located anywhere on the Internet. Hosts that are interested in receiving data flowing to a particular group must join the group using IGMP. Hosts must be a member of the group to receive the data stream.

IP Multicast Addresses

Multicast addresses specify an arbitrary group of IP hosts that have joined the group and want to receive traffic sent to this group.

IP Class D Addresses

The *Internet Assigned Numbers Authority (IANA)* controls the assignment of IP multicast addresses. It has assigned the old Class D address space to be used for IP multicast. This means that all IP multicast group addresses will fall in the range of 224.0.0.0 to 239.255.255.255.

Note This address range is only for the group address or destination address of IP multicast traffic. The source address for multicast datagrams is always the unicast source address.

Reserved Link Local Addresses

The IANA has reserved addresses in the 224.0.0.0 through 224.0.0.255 to be used by network protocols on a local network segment. Packets with these addresses should never be forwarded by a router; they remain local on a particular LAN segment. They are always transmitted with a time-to-live (TTL) of 1.

Network protocols use these addresses for automatic router discovery and to communicate important routing information. For example, OSPF uses 224.0.0.5 and 224.0.0.6 to exchange link state information. Table 43-1 lists some of the well-known addresses.

Table 43-1 Link Local Addresses

Address	Usage
224.1.1.1	All systems on this subnet
224.1.1.2	All routers on this subnet
224.1.1.5	OSPF routers
224.1.1.6	OSPF designated routers
224.0.0.12	DHCP server/relay agent

Globally Scoped Address

The range of addresses from 224.0.1.0 through 238.255.255.255 are called globally scoped addresses. They can be used to multicast data between organizations and across the Internet.

Some of these addresses have been reserved for use by multicast applications through IANA. For example, 224.0.1.1 has been reserved for Network Time Protocol (NTP).

More information about reserved multicast addresses can be found at http://www.isi.edu/in-notes/iana/assignments/multicast-addresses.

Limited Scope Addresses

The range of addresses from 239.0.0.0 through 239.255.255.255 contains limited scope addresses or administratively scoped addresses. These are defined by RFC 2365 to be constrained to a local group or organization. Routers are typically configured with filters to prevent multicast traffic in this address range from flowing outside an autonomous system (AS) or any user-defined domain. Within an autonomous system or domain, the limited scope address range can be further subdivided so those local multicast boundaries can be defined. This also allows for address reuse among these smaller domains.

Glop Addressing

RFC 2770 proposes that the 233.0.0.0/8 address range be reserved for statically defined addresses by organizations that already have an AS number reserved. The AS number of the domain is embedded into the second and third octets of the 233.0.0.0/8 range.

For example, the AS 62010 is written in hex as F23A. Separating out the two octets F2 and 3A, we get 242 and 58 in decimal. This would give us a subnet of 233.242.58.0 that would be globally reserved for AS 62010 to use.

Layer 2 Multicast Addresses

Normally, network interface cards (NICs) on a LAN segment will receive only packets destined for their burned-in MAC address or the broadcast MAC address. Some means had to be devised so that multiple hosts could receive the same packet and still be capable of differentiating among multicast groups.

Fortunately, the IEEE LAN specifications made provisions for the transmission of broadcast and/or multicast packets. In the 802.3 standard, bit 0 of the first octet is used to indicate a broadcast and/or multicast frame. Figure 43-2 shows the location of the broadcast/multicast bit in an Ethernet frame.

Figure 43-2 IEEE 802.3 MAC Address Format

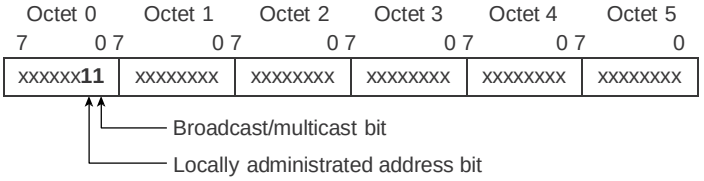

This bit indicates that the frame is destined for an arbitrary group of hosts or all hosts on the network (in the case of the broadcast address, 0xFFFF.FFFF.FFFF).

IP multicast makes use of this capability to transmit IP packets to a group of hosts on a LAN segment.

Ethernet MAC Address Mapping

The IANA owns a block of Ethernet MAC addresses that start with 01:00:5E in hexadecimal. Half of this block is allocated for multicast addresses. This creates the range of available Ethernet MAC addresses to be 0100.5e00.0000 through 0100.5e7f.ffff.

This allocation allows for 23 bits in the Ethernet address to correspond to the IP multicast group address. The mapping places the lower 23 bits of the IP multicast group address into these available 23 bits in the Ethernet address (shown in Figure 43-3).

Figure 43-3 Mapping of IP Multicast to Ethernet/FDDI MAC Address

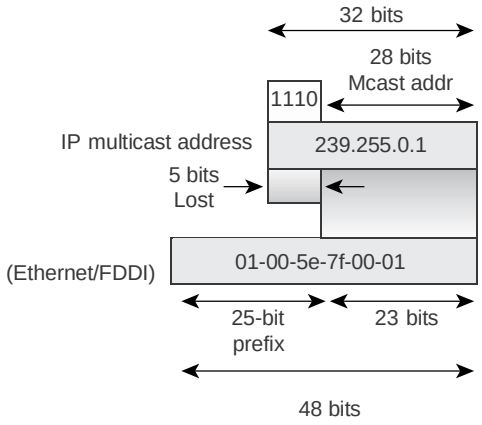

Because the upper 5 bits of the IP multicast address are dropped in this mapping, the resulting address is not unique. In fact, 32 different multicast group IDs all map to the same Ethernet address (see Figure 43-4).

Figure 43-4 MAC Address Ambiguities

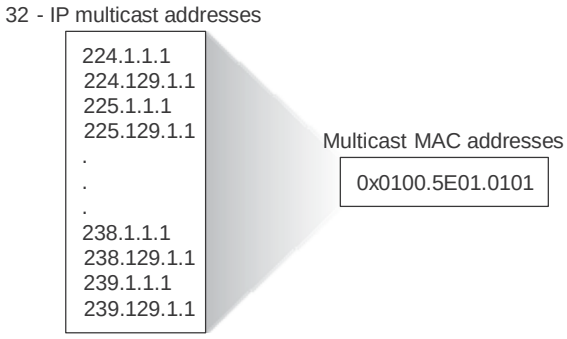

Internet Group Management Protocol

IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

IGMP Version 1

RFC 1112 defines the specification for IGMP Version 1. A diagram of the packet format is found in Figure 43-5.

Figure 43-5 IGMP Version 1 Packet Format

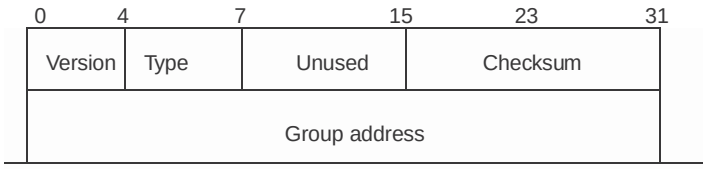

In Version 1, there are just two different types of IGMP messages:

- Membership query
- Membership report

Hosts send out IGMP membership reports corresponding to a particular multicast group to indicate that they are interested in joining that group. The router periodically sends out an IGMP membership query to verify that at least one host on the subnet is still interested in receiving traffic directed to that group. When there is no reply to three consecutive IGMP membership queries, the router times out the group and stops forwarding traffic directed toward that group.

IGMP Version 2

RFC 2236 defines the specification for IGMP Version 2.

A diagram of the packet format follows in Figure 43-6.

Figure 43-6 IGMPv2 Message Format

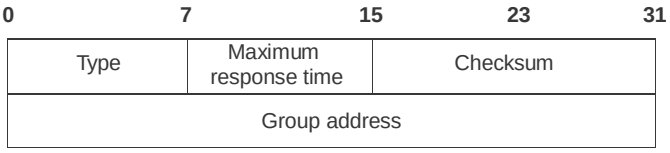

In Version 2, there are four types of IGMP messages:

- Membership query
- Version 1 membership report
- Version 2 membership report
- Leave group

IGMP Version 2 works basically the same as Version 1. The main difference is that there is a leave group message. The hosts now can actively communicate to the local multicast router their intention to leave the group. The router then sends out a group-specific query and determines whether there are any remaining hosts interested in receiving the traffic. If there are no replies, the router times out the group and stops forwarding the traffic. This can greatly reduce the leave latency compared to IGMP Version 1. Unwanted and unnecessary traffic can be stopped much sooner.

Multicast in the Layer 2 Switching Environment

The default behavior for a Layer 2 switch is to forward all multicast traffic to every port that belongs to the destination LAN on the switch. This would defeat the purpose of the switch, which is to limit traffic to the ports that need to receive the data.

Two methods exist by which to deal with multicast in a Layer 2 switching environment efficiently—Cisco Group Management Protocol (CGMP) and IGMP snooping.

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 both on the multicast routers and on the Layer 2 switches. The net result is that with CGMP, IP multicast traffic is delivered only to those Catalyst switch ports that are interested in the traffic. All other ports that have not explicitly requested the traffic will not receive it.

The basic concept of CGMP is shown in Figure 43-7. When a host joins a multicast group (part A), it multicasts an unsolicited IGMP membership report message to the target group (224.1.2.3, in this example). The IGMP report is passed through the switch to the router for the normal IGMP processing. The router (which must have CGMP enabled on this interface) receives this IGMP report and processes it as it normally would, but in addition it creates a CGMP join message and sends it to the switch.

The switch receives this CGMP join message and then adds the port to its content addressable memory (CAM) table for that multicast group. Subsequent traffic directed to this multicast group will be forwarded out the port for that host. The router port is also added to the entry for the multicast group. Multicast routers must listen to all multicast traffic for every group because the IGMP control messages are also sent as multicast traffic. With CGMP, the switch must listen only to CGMP join and CGMP leave messages from the router. The rest of the multicast traffic is forwarded using its CAM table exactly the way the switch was designed.

Figure 43-7 Basic CGMP Operation

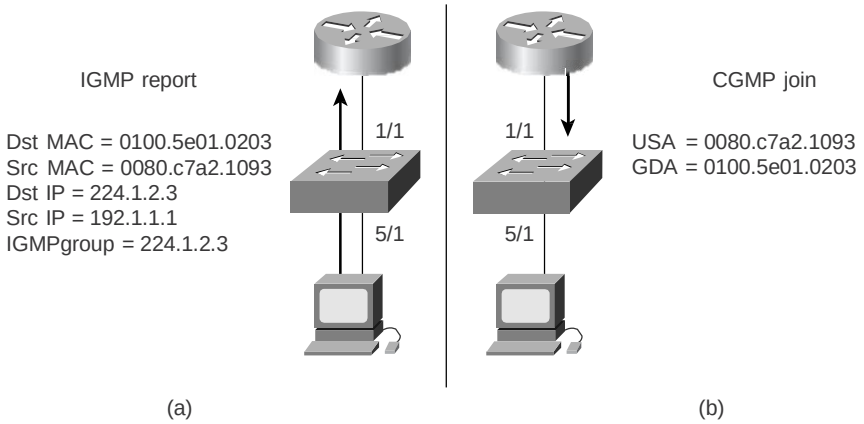

IGMP Snooping

IGMP snooping requires the LAN switch to examine, or snoop, some Layer 3 information in the IGMP packets sent between the hosts and the router. When the switch hears the IGMP host report from a host for a particular multicast group, the switch adds the host's port number to the associated multicast table entry. When the switch hears the IGMP leave group message from a host, it removes the host's port from the table entry.

Because IGMP control messages are transmitted as multicast packets, they are indistinguishable from multicast data at Layer 2. A switch running IGMP snooping examine every multicast data packet to check whether it contains any pertinent IGMP must control information. If IGMP snooping has been implemented on a low-end switch with a slow CPU, this could have a severe performance impact when

data is transmitted at high rates. The solution is to implement IGMP snooping on high-end switches with special ASICs that can perform the IGMP checks in hardware. CGMP is ideal for low-end switches without special hardware.

Multicast Distribution Trees

Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network to deliver traffic to all receivers. 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* whose root is the source of the multicast tree and whose branches form a spanning 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).

Figure 43-8 shows an example of an SPT for group 224.1.1.1 rooted at the source, Host A, and connecting two receivers, hosts B and C.

Figure 43-8 Host A Shortest Path Tree

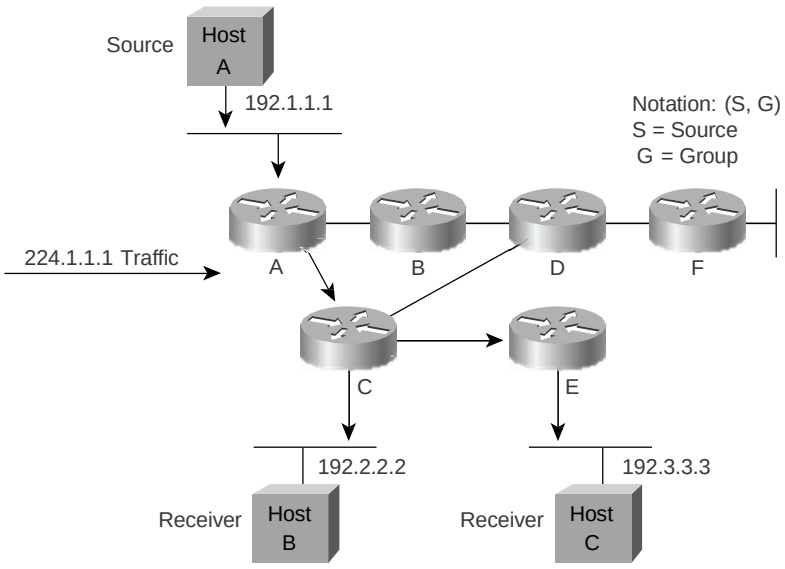

The special notation of (S,G), pronounced "S comma G," enumerates an SPT in which S is the IP address of the source and G is the multicast group address. Using this notation, the SPT for the example in Figure 43-7 would be (192.1.1.1, 224.1.1.1).

The (S,G) notation implies that a separate SPT exists for each individual source sending to each group, which is correct. For example, if Host B is also sending traffic to group 224.1.1.1 and hosts A and C are receivers, then a separate (S,G) SPT would exist with a notation of (192.2.2.2,224.1.1.1).

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 the *rendezvous point (RP)*.

Figure 43-9 shows a shared tree for the group 224.2.2.2 with the root located at Router D. When using a shared tree, sources must send their traffic to the root, and then the traffic is forwarded down the shared tree to reach all receivers.

Figure 43-9 Shared Distribution Tree

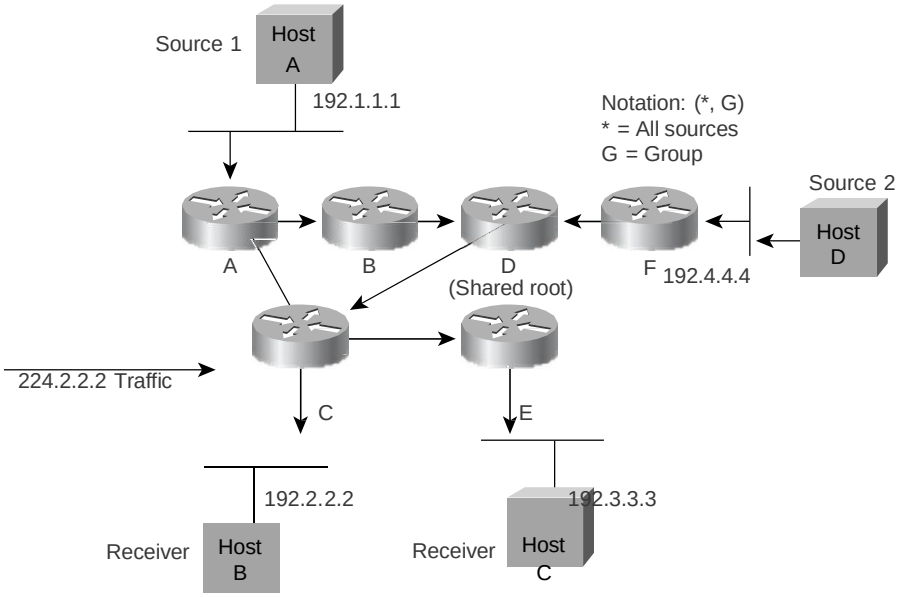

In this example, multicast traffic from the source hosts A and D travels to the root (Router D) and then down the shared tree to the two receivers, hosts B and C. Because all sources in the multicast group use a common shared tree, a wildcard notation written as (*, G), pronounced "star comma G," represents the tree. In this case, * means all sources, and the G represents the multicast group. Therefore, the shared tree shown in Figure 43-8 would be written as (*, 224.2.2.2).

Both SPT and shared trees are loop-free. Messages are replicated only where the tree branches.

Members of multicast groups can join or leave at any time, so the distribution trees must be dynamically updated. When all the active receivers on a particular branch stop requesting the traffic for a particular multicast group, the routers prune that branch from the distribution tree and stop forwarding traffic down that branch. If one receiver on that branch becomes active and requests the multicast traffic, the router dynamically modifies the distribution tree and starts forwarding traffic again.

Shortest path trees have the advantage of creating the optimal path between the source and the receivers. This guarantees the minimum amount of network latency for forwarding multicast traffic. This optimization does come with a price, though: The routers must maintain path information for each source. In a network that has thousands of sources and thousands of groups, this can quickly become a resource issue on the routers. Memory consumption from the size of the multicast routing table is a factor that network designers must take into consideration.

Shared trees have the advantage of requiring the minimum amount of state in each router. This lowers the overall memory requirements for a network that allows only shared trees. The disadvantage of shared trees is that, under certain circumstances, the paths between the source and receivers might not be the optimal paths—which might introduce some latency in packet delivery. Network designers must carefully consider the placement of the RP when implementing an environment with only shared trees.

Multicast Forwarding

In unicast routing, traffic is routed through the network along a single path from the source to the destination host. A unicast router does not really care about the source address—it only cares about the destination address and how to forward the traffic towards that destination. The router scans through its routing table and then forwards a single copy of the unicast packet out the correct interface in the direction of the destination.

In multicast routing, the source is sending traffic to an arbitrary group of hosts represented by a multicast group address. The multicast router must determine which direction is upstream (toward the source) and which direction (or directions) is downstream. If there are multiple downstream paths, the router replicates the packet and forwards the traffic down the appropriate downstream paths—which is not necessarily all paths. This concept of forwarding multicast traffic away from the source, rather than to the receiver, is called *reverse path forwarding*.

Reverse Path Forwarding

Reverse path forwarding (RPF) is a fundamental concept in multicast routing that enables routers to correctly forward multicast traffic down the distribution tree. RPF makes use of the existing unicast routing table to determine the upstream and downstream neighbors. A router forwards a multicast packet only if it is received on the upstream interface. This RPF check helps to guarantee that the distribution tree will be loop-free.

RPF Check

When a multicast packet arrives at a router, the router performs an RPF check on the packet. If the RPF check is successful, the packet is forwarded. Otherwise, it is dropped.

For traffic flowing down a source tree, the RPF check procedure works as follows:

Step 1 Router looks up the source address in the unicast routing table to determine whether it has arrived on the interface that is on the reverse path back to the source.

Step 2 If packet has arrived on the interface leading back to the source, the RPF check is successful and the packet is forwarded.

Step 3 If the RPF check in Step 2 fails, the packet is dropped.

Figure 43-10 shows an example of an unsuccessful RPF check.

Figure 43-10 RPF Check Fails

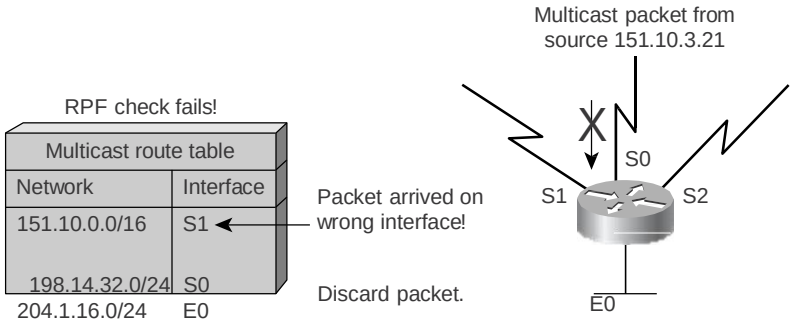

A multicast packet from source 151.10.3.21 is received on interface S0. A check of the unicast route table shows that the interface that this router would use to forward unicast data to 151.10.3.21 is S1. Because the packet has arrived on S0, the packet will be discarded.

Figure 43-11 shows an example of a successful RPF check.

Figure 43-11 RPF Check Succeeds

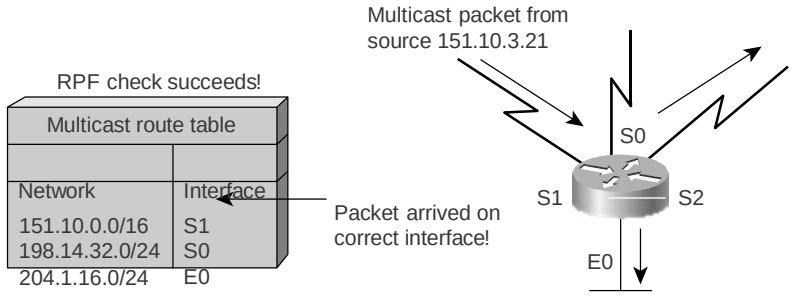

This time the multicast packet has arrived on S1. The router checks the unicast routing table and finds that S1 is the correct interface. The RPF check passes and the packet is forwarded.

Protocol-Independent Multicast

Protocol-independent multicast (PIM) gets its name from the fact that it is IP routing protocol-independent. PIM can leverage whichever unicast routing protocols are used to populate the unicast routing table, including EIGRP, OSPF, BGP, or static routes. PIM uses this unicast routing information to perform the multicast forwarding function, so it is IP protocol-independent. Although PIM is called a multicast routing protocol, it actually uses the unicast routing table to perform the reverse path forwarding (RPF) check function instead of building up a completely independent multicast routing table. PIM does not send and receive multicast routing updates between routers like other routing protocols do.

PIM Dense Mode

PIM Dense Mode (PIM-DM) uses a push model to flood multicast traffic to every corner of the network. This is a brute-force method for delivering data to the receivers, but in certain applications, this might be an efficient mechanism if there are active receivers on every subnet in the network.

PIM-DM initially floods multicast traffic throughout the network. Routers that do not have any downstream neighbors prune back the unwanted traffic. This process repeats every 3 minutes.

The flood and prune mechanism is how the routers accumulate their state information—by receiving the data stream. These data streams contain the source and group information so that downstream routers can build up their multicast forwarding tables. PIM-DM can support only source trees—(S,G) entries. It cannot be used to build a shared distribution tree.

PIM Sparse Mode

PIM Sparse Mode (PIM-SM) uses a pull model to deliver multicast traffic. Only networks that have active receivers that have explicitly requested the data will be forwarded the traffic. PIM-SM is defined in RFC 2362.

PIM-SM uses a shared tree to distribute the information about active sources. Depending on the configuration options, the traffic can remain on the shared tree or switch over to an optimized source distribution tree. The latter is the default behavior for PIM-SM on Cisco routers. The traffic starts to flow down the shared tree, and then routers along the path determine whether there is a better path to the source. If a better, more direct path exists, the designated router (the router closest to the receiver) will send a join message toward the source and then reroute the traffic along this path.

PIM-SM has the concept of an RP, since it uses shared trees—at least initially. The RP must be administratively configured in the network. Sources register with the RP, and then data is forwarded down the shared tree to the receivers. If the shared tree is not an optimal path between the source and the receiver, the routers dynamically create a source tree and stop traffic from flowing down the shared tree. This is the default behavior in IOS. Network administrators can force traffic to stay on the shared tree by using a configuration option (lp pim spt-threshold infinity).

PIM-SM scales well to a network of any size, including those with WAN links. The explicit join mechanism prevents unwanted traffic from flooding the WAN links.

Sparse-Dense Mode

Cisco has implemented an alternative to choosing just dense mode or just sparse mode on a router interface new IP. This was necessitated by a change in the paradigm for forwarding multicast traffic via PIM that became apparent during its development. It turned out that it was more efficient to choose sparse or dense on a per group basis rather than a per router interface basis. Sparse-dense mode facilitates this ability.

Network administrators can also configure sparse-dense mode. This configuration option allows individual groups to be run in either sparse or dense mode, depending on whether RP information is available for that group. If the router learns RP information for a particular group, it will be treated as sparse mode; otherwise, that group will be treated as dense mode.

Multiprotocol Border Gateway Protocol

Multiprotocol Border Gateway Protocol (MBGP) gives a method for providers to distinguish which route prefixes they will use for performing multicast RPF checks. The RPF check is the fundamental mechanism that routers use to determine the paths that multicast forwarding trees will follow and successfully deliver multicast content from sources to receivers.

MBGP is described in RFC 2283, Multiprotocol Extensions for BGP-4. Since MBGP is an extension of BGP, it brings along all the administrative machinery that providers and customers like in their interdomain routing environment. Including all the inter-AS tools to filter and control routing (e.g., route maps). Therefore, by using MBGP, any network utilizing internal or external BGP can apply the multiple policy control knobs familiar in BGP to specify routing (and thereby forwarding) policy for multicast.

Two path attributes, MP_REACH_NLRI and MP_UNREACH_NLRI have been introduced in BGP4+. These new attributes create a simple way to carry two sets of routing information—one for unicast routing and one for multicast routing. The routes associated with multicast routing are used to build the multicast distribution trees.

The main advantage of MBGP is that an internet can support noncongruent unicast and multicast topologies. When the unicast and multicast topologies are congruent, MBGP can support different policies for each. MBGP provides a scalable policy based interdomain routing protocol.

Multicast Source Discovery Protocol

In the PIM Sparse mode model, multicast sources and receivers must register with their local Rendezvous Point (RP). Actually, the closest router to the sources or receivers registers with the RP but the point is that the RP knows about all the sources and receivers for any particular group. RPs in other domains have no way of knowing about sources located in other domains. MSDP is an elegant way to solve this problem. MSDP is a mechanism that connects PIM-SM domains and allows RPs to share information about active sources. When RPs in remote domains know about active sources they can pass on that information to their local receivers and multicast data can be forwarded between the domains. A nice feature of MSDP is that it allows each domain to maintain an independent RP which does not rely on other domains, but it does enable RPs to forward traffic between domains.

The RP in each domain establishes an MSDP peering session using a TCP connection with the RPs in other domains or with border routers leading to the other domains. When the RP learns about a new multicast source within its own domain (through the normal PIM register mechanism), the RP encapsulates the first data packet in a Source Active (SA) message and sends the SA to all MSDP peers. The SA is forwarded by each receiving peer using a modified RPF check, until it reaches every MSDP router in the interconnected networks—theoretically the entire multicast internet. If the receiving MSDP peer is an RP, and the RP has a (*,G) entry for the group in the SA (there is an interested receiver), the RP will create (S,G) state for the source and join to the shortest path tree for the state of the source. The encapsulated data is decapsulated and forwarded down that RP's shared tree. When the packet is received by a receiver's last hop router, the last-hop may also join the shortest path tree to the source. The source's RP periodically sends SAs, which include all sources within that RP's own domain. Figure 43-12 shows how data would flow between a source in domain A to a receiver in domain E.

Figure 43-12 MSDP Example

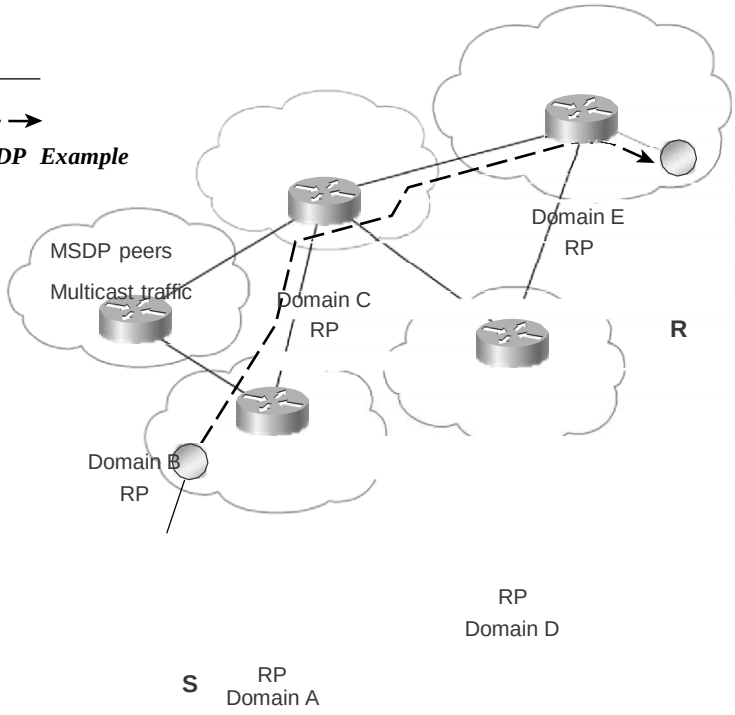

MDSP was developed for peering between Internet Service Providers (ISPs). ISPs did not want to rely on an RP maintained by a competing ISP to service their customers. MSDP allows each ISP to have their own local RP and still forward and receive multicast traffic to the Internet.

Anycast RP-Logical RP

A very useful application of MSDP is called anycast RP. This is a technique for configuring a multicast sparse-mode network to provide for fault tolerance and load sharing within a single multicast domain.

Two or more RPs are configured with the same IP address on loopback interfaces—say, 10.0.0.1, for example (refer to Figure 43-13). The loopback address should be configured as a 32 bit address. All the downstream routers are configured so that they know that their local RP's address is 10.0.0.1. IP routing automatically selects the topologically closest RP for each source and receiver. Because some sources might end up using one RP and some receivers a different RP, there needs to be some way for the RPs to exchange information about active sources. This is done with MSDP. All the RPs are configured to be MSDP peers of each other. Each RP will know about the active sources in the other RP's area. If any of the RPs fail, IP routing will converge and one of the RPs will become the active RP in both areas.

Note The Anycast RP example above uses IP addresses from RFC 1918. These IP addresses are normally blocked at interdomain borders and therefore are not accessible to other ISPs. You must use valid IP addresses if you want the RPs to be reachable from other domains.

Figure 43-13 Anycast RP

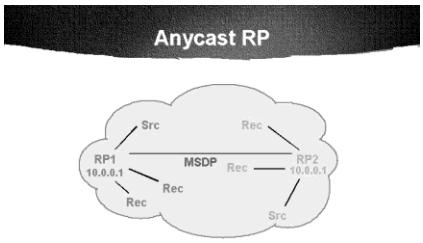

Note The RPs are used only to set up the initial connection between sources and receivers. After the last-hop routers join the shortest path tree, the RP is no longer necessary.

Multicast Address Dynamic Client Allocation Protocol

The *Multicast Address Dynamic Client Allocation Protocol (MADCAP)* is defined in RFC 2730 as a protocol that allows hosts to request a multicast address allocation dynamically from a MADCAP server. The concept is very similar to the way DHCP works today and is built on a client/server model.

Multicast-Scope Zone Announcement Protocol

Multicast-Scope Zone Announcement Protocol (MZAP) is defined in RFC 2776 as a protocol that allows networks to automatically discover administratively scoped zones relative to a particular location.

Reliable Multicast-Pragmatic General Multicast

Pragmatic General Multicast (PGM) is a reliable multicast transport protocol for applications that require ordered, duplicate-free, multicast data delivery from multiple sources to multiple receivers. PGM guarantees that a receiver in a multicast group either receives all data packets from transmissions and retransmissions, or can detect unrecoverable data packet loss.

The PGM Reliable Transport Protocol itself is implemented on the sources and the receivers. The source maintains a transmit window of outgoing data packets and retransmits individual packets when it receives a negative acknowledgment (NAK). The network elements (routers) assist in suppressing an implosion of NAKs (when a failure does occur) and aids in efficient forwarding of the retransmitted data just to the networks that need it.

PGM is intended as a solution for multicast applications with basic reliability requirements. The specification for PGM is network layer-independent. The Cisco implementation of PGM Router Assist supports PGM over IP.

Today, the specification for PGM is an Internet draft that can be found on the IETF web site
(http://www.ietf.org) under the name “PGM Reliable Transport Protocol.”

Review Questions

Q—*What is the range of available IP multicast addresses?*

A—224.0.0.0 to 239.255.255.255.

Q—*What is the purpose of IGMP?*

A—IGMP is used between the hosts and their local multicast router to join and leave multicast groups.

Q—*What is an advantage of IGMPv2 over IGMPv1?*

A—IGMPv2 has a leave group message that can greatly reduce the latency of unwanted traffic on a LAN.

Q—*What is a potential disadvantage of IGMP snooping over CGMP on a low-end Layer 2 switch?*

A—IGMP snooping requires the switch to examine every multicast packet for an IGMP control message. On a low-end switch, this might have a severe performance impact.

Q—*What is an advantage of shortest path (or source) trees compared to shared trees?*

A—Source trees guarantee an optimal path between each source and each receiver, which will minimize network latency.

Q—*What is an advantage of using shared trees?*

A—Shared trees require very little state to be kept in the routers, which requires less memory.

Q—*What information does the router use to do an RPF check?*

A—The unicast routing table.

Q—*Why is protocol-independent multicast called "independent"?*

A—PIM works with any underlying IP unicast routing protocol—RIP, EIGRP, OSPF, BGP or static routes.

Q—*What is the main advantage of MBGP?*

A—Providers can have noncongruent unicast and multicast routing topologies.

Q—*How do RPs learn about sources from other RPs with MSDP*?

A—RPs are configured to be MSDP peers with other RPs. Each RP forwards source active (SA)
messages to each other.

Q—*What is the purpose of the anycast RP?*

A—Load balancing and fault tolerance.

For More Information

Williamson, Beau. *Developing IP Multicast Networks*. Indianapolis: Cisco Press,

2000. Multicast Quick Start Configuration Guide

(http://www.cisco.com/warp/customer/105/48.html)

Appendix B: RFC 4541-Considerations for IGMP

Also found online at http://tools.ietf.org/html/rfc4541

Network Working Group M. Christensen
Request for Comments: 4541 Thrane & Thrane
Category: Informational K. Kimball
Hewlett-Packard
F. Solensky
Calix
May 2006

Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches

Abstract

This memo describes the recommendations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) snooping switches. These are based on best current practices for IGMPv2, with further considerations for IGMPv3- and MLDv2-snooping. Additional areas of relevance, such as link layer topology changes and Ethernet-specific encapsulation issues, are also considered.

1. Introduction

The IEEE bridge standard [BRIDGE] specifies how LAN packets are 'bridged', or as is more commonly used today, switched between LAN segments. The operation of a switch with respect to multicast packets can be summarized as follows. When processing a packet whose destination MAC address is a multicast address, the switch will forward a copy of the packet into each of the remaining network interfaces that are in the forwarding state in accordance with [BRIDGE]. The spanning tree algorithm ensures that the application of this rule at every switch in the network will make the packet accessible to all nodes connected to the network.

This behaviour works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth, particularly when the packet is

intended for only a small number of nodes. Packets will be flooded into network segments where no node has any interest in receiving the packet. While nodes will rarely incur any processing overhead to filter packets addressed to unrequested group addresses, they are unable to transmit new packets onto the shared media for the period of time that the multicast packet is flooded. In general, significant bandwidth can be wasted by flooding.

In recent years, a number of commercial vendors have introduced products described as "IGMP snooping switches" to the market. These devices do not adhere to the conceptual model that provides the strict separation of functionality between different communications layers in the ISO model, and instead utilize information in the upper level protocol headers as factors to be considered in processing at the lower levels. This is analogous to the manner in which a router can act as a firewall by looking into the transport protocol's header before allowing a packet to be forwarded to its destination address.

In the case of IP multicast traffic, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network where no node has expressed interest in receiving packets addressed to the group address. This is in contrast to normal switch behavior where multicast traffic is typically forwarded on all interfaces.

Many switch datasheets state support for IGMP snooping, but no recommendations for this exist today. It is the authors' hope that the information presented in this document will supply this foundation.

The recommendations presented here are based on the following information sources: The IGMP specifications [RFC1112], [RFC2236] and [IGMPv3], vendor-supplied technical documents [CISCO], bug reports [MSOFT], discussions with people involved in the design of IGMP snooping switches, MAGMA mailing list discussions, and on replies by switch vendors to an implementation questionnaire.

Interoperability issues that arise between different versions of IGMP are not the focus of this document. Interested readers are directed to [IGMPv3] for a thorough description of problem areas.

The suggestions in this document are based on IGMP, which applies only to IPv4. For IPv6, Multicast Listener Discovery [MLD] must be used instead. Because MLD is based on IGMP, we do not repeat the entire description and recommendations for MLD snooping switches. Instead, we point out the few cases where there are differences from IGMP.

Note that the IGMP snooping function should apply only to IPv4 multicasts. Other multicast packets, such as IPv6, might be suppressed by IGMP snooping if additional care is not taken in the implementation as mentioned in the recommendations section. It is desired not to restrict the flow of non-IPv4 multicasts other than to the degree which would happen as a result of regular bridging functions. Likewise, MLD snooping switches are discouraged from using topological information learned from IPv6 traffic to alter the forwarding of IPv4 multicast packets.

2. IGMP Snooping Recommendations

The following sections list the recommendations for an IGMP snooping switch. The recommendation is stated and is supplemented by a description of a possible implementation approach. All implementation discussions are examples only and there may well be other ways to achieve the same functionality.

2.1. Forwarding rules

The IGMP snooping functionality is separated into a control section (IGMP forwarding) and a data section (Data forwarding).

2.1.1. IGMP Forwarding Rules

1) A snooping switch should forward IGMP Membership Reports only to those ports where multicast routers are attached.

Alternatively stated: a snooping switch should not forward IGMP Membership Reports to ports on which only hosts are attached. An administrative control may be provided to override this restriction, allowing the report messages to be flooded to other ports.

This is the main IGMP snooping functionality for the control path.

Sending membership reports to other hosts can result, for IGMPv1 and IGMPv2, in unintentionally preventing a host from joining a specific multicast group.

When an IGMPv1 or IGMPv2 host receives a membership report for a group address that it intends to join, the host will suppress its own membership report for the same group. This join or message suppression is a requirement for IGMPv1 and IGMPv2 hosts. However, if a switch does not receive a membership report from the host it will not forward multicast data to it.

This is not a problem in an IGMPv3-only network because there is no suppression of IGMP Membership reports.

The administrative control allows IGMP Membership Report messages to be processed by network monitoring equipment such as packet analyzers or port replicators.

The switch supporting IGMP snooping must maintain a list of multicast routers and the ports on which they are attached. This list can be constructed in any combination of the following ways:

a) This list should be built by the snooping switch sending Multicast Router Solicitation messages as described in IGMP Multicast Router Discovery [MRDISC]. It may also snoop Multicast Router Advertisement messages sent by and to other nodes.

b) The arrival port for IGMP Queries (sent by multicast routers) where the source address is not 0.0.0.0.

The 0.0.0.0 address represents a special case where the switch is proxying IGMP Queries for faster network convergence, but is not itself the Querier. The switch does not use its own IP address (even if it has one), because this would cause the Queries to be seen as coming from a newly elected Querier. The 0.0.0.0 address is used to indicate that the Query packets are NOT from a multicast router.

c) Ports explicitly configured by management to be IGMP-forwarding ports, in addition to or instead of any of the above methods to detect router ports.

2) IGMP networks may also include devices that implement "proxy-reporting", in which reports received from downstream hosts are summarized and used to build internal membership states. Such proxy-reporting devices may use the all-zeros IP Source-Address when forwarding any summarized reports upstream. For this reason, IGMP membership reports received by the snooping switch must not be rejected because the source IP address is set to 0.0.0.0.

3) The switch that supports IGMP snooping must flood all unrecognized IGMP messages to all other ports and must not attempt to make use of any information beyond the end of the network layer header.

In addition, earlier versions of IGMP should interpret IGMP fields as defined for their versions and must not alter these fields when forwarding the message. When generating new messages, a given

IGMP version should set fields to the appropriate values for its

own version. If any fields are reserved or otherwise undefined for a given IGMP version, the fields should be ignored when parsing the message and must be set to zeroes when new messages are generated by implementations of that IGMP version. An exception may occur if the switch is performing a spoofing function, and is aware of the settings for new or reserved fields that would be required to correctly spoof for a different IGMP version.

The reason to worry about these trivialities is that IGMPv3 overloads the old IGMP query message using the same type number (0x11) but with an extended header. Therefore there is a risk that IGMPv3 queries may be interpreted as older version queries by, for example, IGMPv2 snooping switches. This has already been reported [IETF56] and is discussed in section 2.2.

4) An IGMP snooping switch should be aware of link layer topology changes caused by Spanning Tree operation. When a port is enabled or disabled by Spanning Tree, a General Query may be sent on all active non-router ports in order to reduce network convergence time. Non-Querier switches should be aware of whether the Querier is in IGMPv3 mode. If so, the switch should not spoof any General Queries unless it is able to send an IGMPv3 Query that adheres to the most recent information sent by the true Querier. In no case should a switch introduce a spoofed IGMPv2 Query into an IGMPv3 network, as this may create excessive network disruption.

If the switch is not the Querier, it should use the 'all-zeros' IP Source Address in these proxy queries (even though some hosts may elect to not process queries with a 0.0.0.0 IP Source Address). When such proxy queries are received, they must not be included in the Querier election process.

5) An IGMP snooping switch must not make use of information in IGMP packets where the IP or IGMP headers have checksum or integrity errors. The switch should not flood such packets but if it does, it should also take some note of the event (i.e., increment a counter). These errors and their processing are further discussed in [IGMPv3], [MLD] and [MLDv2].

6) The snooping switch must not rely exclusively on the appearance of IGMP Group Leave announcements to determine when entries should be removed from the forwarding table. It should implement a membership timeout mechanism such as the router-side functionality of the IGMP protocol as described in the IGMP and MLD specifications (See Normative Reference section for IGMPv1-3 and MLDv1-2) on all its non-router ports. This timeout value should be configurable.

2.1.2. Data Forwarding Rules

1) Packets with a destination IP address outside 224.0.0.X which are not IGMP should be forwarded according to group-based port membership tables and must also be forwarded on router ports.

   This is the main IGMP snooping functionality for the data path. One approach that an implementation could take would be to maintain separate membership and multicast router tables in software and then "merge" these tables into a forwarding cache.

2) Packets with a destination IP (DIP) address in the 224.0.0.X range which are not IGMP must be forwarded on all ports.

   This recommendation is based on the fact that many host systems do not send Join IP multicast addresses in this range before sending or listening to IP multicast packets. Furthermore, since the 224.0.0.X address range is defined as link-local (not to be routed), it seems unnecessary to keep the state for each address in this range. Additionally, some routers operate in the 224.0.0.X address range without issuing IGMP Joins, and these applications would break if the switch were to prune them due to not having seen a Join Group message from the router.

3) An unregistered packet is defined as an IPv4 multicast packet with a destination address which does not match any of the groups announced in earlier IGMP Membership Reports.

   If a switch receives an unregistered packet, it must forward that packet on all ports to which an IGMP router is attached. A switch may default to forwarding unregistered packets on all ports. Switches that do not forward unregistered packets to all ports must include a configuration option to force the flooding of unregistered packets on specified ports.

   In an environment where IGMPv3 hosts are mixed with snooping switches that do not yet support IGMPv3, the switch's failure to flood unregistered streams could prevent v3 hosts from receiving their traffic. Alternatively, in environments where the snooping switch supports all of the IGMP versions that are present, flooding unregistered streams may cause IGMP hosts to be overwhelmed by multicast traffic, even to the point of not receiving Queries and failing to issue new membership reports for their own groups.

   It is encouraged that snooping switches at least recognize and process IGMPv3 Join Reports, even if this processing is limited to the behavior for IGMPv2 Joins, i.e., is done without considering

any additional "include source" or "exclude source" filtering. When IGMPv3 Joins are not recognized, a snooping switch may incorrectly prune off the unregistered data streams for the groups (as noted above); alternatively, it may fail to add in forwarding to any new IGMPv3 hosts if the group has previously been joined as IGMPv2 (because the data stream is seen as already having been registered).

4) All non-IPv4 multicast packets should continue to be flooded out to all remaining ports in the forwarding state as per normal IEEE bridging operations.

This recommendation is a result of the fact that groups made up of IPv4 hosts and IPv6 hosts are completely separate and distinct groups. As a result, information gleaned from the topology between members of an IPv4 group would not be applicable when forming the topology between members of an IPv6 group.

5) IGMP snooping switches may maintain forwarding tables based on either MAC addresses or IP addresses. If a switch supports both types of forwarding tables then the default behavior should be to use IP addresses. IP address based forwarding is preferred because the mapping between IP multicast addresses and link-layer multicast addresses is ambiguous. In the case of Ethernet, there is a multiplicity of 1 Ethernet address to 32 IP addresses [RFC1112].

6) Switches which rely on information in the IP header should verify that the IP header checksum is correct. If the checksum fails, the information in the packet must not be incorporated into the forwarding table. Further, the packet should be discarded.

7) When IGMPv3 "include source" and "exclude source" membership reports are received on shared segments, the switch needs to forward the superset of all received membership reports on to the shared segment. Forwarding of traffic from a particular source S to a group G must happen if at least one host on the shared segment reports an IGMPv3 membership of the type INCLUDE(G, Slist1) or EXCLUDE(G, Slist2), where S is an element of Slist1 and not an element of Slist2.

The practical implementation of the (G,S1,S2,...) based data forwarding tables are not within the scope of this document. However, one possibility is to maintain two (G,S) forwarding lists: one for the INCLUDE filter where a match of a specific (G,S) is required before forwarding will happen, and one for the EXCLUDE filter where a match of a specific (G,S) will result in no forwarding.

2.2. GMP Snooping-Related Problems

A special problem arises in networks consisting of IGMPv3 routers as well as IGMPv2 and IGMPv3 hosts interconnected by an IGMPv2 snooping switch as recently reported [IETF56]. The router will continue to maintain IGMPv3 even in the presence of IGMPv2 hosts, and thus the network will not converge on IGMPv2. But it is likely that the IGMPv2 snooping switch will not recognize or process the IGMPv3 membership reports. Groups for these unrecognized reports will then either be flooded (with all of the problems that may create for hosts in a network with a heavy multicast load) or pruned by the snooping switch.

Therefore, it is recommended that in such a network, the multicast router be configured to use IGMPv2. If this is not possible, and if the snooping switch cannot recognize and process the IGMPv3 membership reports, it is instead recommended that the switch's IGMP snooping functionality be disabled, as there is no clear solution to this problem.

3. IPv6 Considerations

In order to avoid confusion, the previous discussions have been based on the IGMP protocol which only applies to IPv4 multicast. In the case of IPv6, most of the above discussions are still valid with a few exceptions that we will describe here.

The control and data forwarding rules in the IGMP section can, with a few considerations, also be applied to MLD. This means that the basic functionality of intercepting MLD packets, and building membership lists and multicast router lists, is the same as for IGMP.

In IPv6, the data forwarding rules are more straight forward because MLD is mandated for addresses with scope 2 (link-scope) or greater. The only exception is the address FF02::1 which is the all hosts link-scope address for which MLD messages are never sent. Packets with the all hosts link-scope address should be forwarded on all ports.

MLD messages are also not sent regarding groups with addresses in the range FF00::/15 (which encompasses both the reserved FF00::/16 and node-local FF01::/16 IPv6 address spaces). These addresses should never appear in packets on the link.

Equivalent to the IPv4 behaviors regarding the null IP Source address, MLD membership reports must not be rejected by an MLD snooping switch because of an unspecified IP source address (::).

Additionally, if a non-Querier switch spoofs any General Queries (as

addressed in Section 2.1 above, for Spanning Tree topology changes), the switch should use the null IP source address (::) when sending said queries. When such proxy queries are received, they must not be included in the Querier election process.

The three major differences between IPv4 and IPv6 in relation to multicast are:

-The IPv6 protocol for multicast group maintenance is called Multicast Listener Discovery [MLDv2]. MLDv2 uses ICMPv6 message types instead of IGMP message types.

-The RFCs [IPV6-ETHER] and [IPV6-FDDI] describe how 32 of the 128 bit DIP addresses are used to form the 48 bit DMAC addresses for multicast groups, while [IPV6-TOKEN] describes the mapping for token ring DMAC addresses by using three low-order bits. The specification [IPV6-1394] makes use of a 6 bit channel number.

-Multicast router discovery is accomplished using the Multicast Router Discovery Protocol (MRDISC) defined in [MRDISC].

The IPv6 packet header does not include a checksum field. Nevertheless, the switch should detect other packet integrity issues such as address version and payload length consistencies if possible. When the snooping switch detects such an error, it must not include information from the corresponding packet in the MLD forwarding table. The forwarding code should instead drop the packet and take further reasonable actions as advocated above.

The fact that MLDv2 is using ICMPv6 adds new requirements to a snooping switch because ICMPv6 has multiple uses aside from MLD. This means that it is no longer sufficient to detect that the next-header field of the IP header is ICMPv6 in order to identify packets relevant for MLD snooping. A software-based implementation which treats all ICMPv6 packets as candidates for MLD snooping could easily fill its receive queue and bog down the CPU with irrelevant packets. This would prevent the snooping functionality from performing its intended purpose and the non-MLD packets destined for other hosts could be lost.

A solution is either to require that the snooping switch looks further into the packets, or to be able to detect a multicast DMAC address in conjunction with ICMPv6. The first solution is desirable when a configuration option allows the administrator to specify which ICMPv6 message types should trigger a CPU redirect and which should not. The reason is that a hardcoding of message types is inflexible for the introduction of new message types. The second solution introduces the risk that new protocols that use ICMPv6 and multicast

DMAC addresses could be incorrectly identified as MLD. It is suggested that solution one is preferred when the configuration option is provided. If this is not the case, then the implementor should seriously consider making it available since Neighbor Discovery messages would be among those that fall into this false positive case and are vital for the operational integrity of IPv6 networks.

The mapping from IP multicast addresses to multicast DMAC addresses introduces a potentially enormous overlap. The structure of an IPv6 multicast address is shown in the figure below. As a result, there are 2 ** (112 - 32), or more than 1.2e24 unique DIP addresses which map into a single DMAC address in Ethernet and FDDI. This should be compared to 2**5 in the case of IPv4.

Initial allocation of IPv6 multicast addresses, as described in [RFC3307], however, cover only the lower 32 bits of group ID. While this reduces the problem of address ambiguity to group IDs with different flag and scope values for now, it should be noted that the allocation policy may change in the future. Because of the potential overlap it is recommended that IPv6 address based forwarding is preferred to MAC address based forwarding.

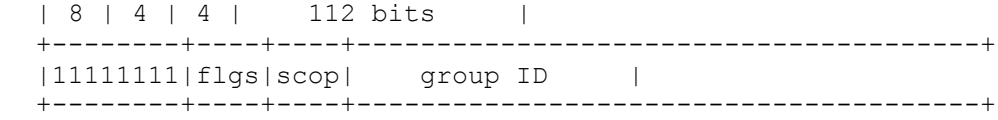

4. IGMP Questionnaire

As part of this work, the following questions were asked on the MAGMA discussion list and were sent to known switch vendors implementing IGMP snooping. The individual contributions have been anonymized upon request and do not necessarily apply to all of the vendors' products.

The questions were:

Q1 Do your switches perform IGMP Join aggregation? In other words, are IGMP joins intercepted, absorbed by the hardware/software so that only one Join is forwarded to the querier?

Q2 Is multicast forwarding based on MAC addresses? Would datagrams addressed to multicast IP addresses 224.1.2.3 and 239.129.2.3 be forwarded on the same ports-groups?

Q3 Is it possible to forward multicast datagrams based on IP addresses (not routed)? In other words, could 224.1.2.3 and 239.129.2.3 be forwarded on different port-groups with unaltered TTL?

Q4 Are multicast datagrams within the range 224.0.0.1 to 224.0.0.255 forwarded on all ports whether or not IGMP Joins have been sent?

Q5 Are multicast frames within the MAC address range 01:00:5E:00:00:01 to 01:00:5E:00:00:FF forwarded on all ports whether or not IGMP joins have been sent?

Q6 Does your switch support forwarding to ports on which IP multicast routers are attached in addition to the ports where IGMP Joins have been received?

Q7 Is your IGMP snooping functionality fully implemented in hardware?

Q8 Is your IGMP snooping functionality partly software implemented?

Q9 Can topology changes (for example spanning tree configuration changes) be detected by the IGMP snooping functionality so that for example new queries can be sent or tables can be updated to ensure robustness?

The answers were:

	Switch Vendor					
	1	2	3	4	5	6
Q1 Join aggregation	x	x	x		x	x
Q2 Layer-2 forwarding	x	x	x	x	(1)	
Q3 Layer-3 forwarding	(1)		(1)		(1)	x
Q4 224.0.0.X aware	(1)	x	(1)	(2)	x	x
Q5 01:00:5e:00:00:XX aware	x	x	x	(2)	x	x
Q6 Mcast router list	x	x	x	x	x	x
Q7 Hardware implemented						
Q8 Software assisted	x	x	x	x	x	x
Q9 Topology change aware	x	x	x	x		(2)

x Means that the answer was Yes.

(1) In some products (typically high-end) Yes; in others No.

(2) Not at the time that the questionnaire was received but expected in the near future.

5. References

5.1. Normative References

[BRIDGE] IEEE Std. 802.1D-2004 IEEE Standard for Local and metropolitan area networks, Media Access Control (MAC) Bridges

[IGMPv3] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[IPV6-1394] Fujisawa, K. and A. Onoe, "Transmission of IPv6 Packets over IEEE 1394 Networks", RFC 3146, October 2001.

[IPV6-ETHER] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.

[IPV6-FDDI] Crawford, M., "Transmission of IPv6 Packets over FDDI Networks", RFC 2467, December 1998.

[IPV6-TOKEN] Crawford, M., Narten, T., and S. Thomas, "Transmission of IPv6 Packets over Token Ring Networks", RFC 2470, December 1998.

[MLD] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.

[MLDv2] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[MRDISC] Haberman, B. and J. Martin, "Multicast Router Discovery", RFC 4286, December 2005.

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.

[RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.

[RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.

5.2. nformative References

[CISCO] Cisco Tech Notes, "Multicast In a Campus Network: CGMP and IGMP snooping", http://www.cisco.com/warp/public/473/22.html

[IETF56] Briefing by Dave Thaler, Microsoft, presented to the MAGMA WG at the 56'th IETF meeting in San Francisco, http://www.ietf.org/proceedings/03mar/index.html

[MSOFT] Microsoft support article Q223136, "Some LAN Switches with IGMP Snooping Stop Forwarding Multicast Packets on RRAS Startup", http://support.microsoft.com/support/articles/Q223/1/36.ASP

6. Security Considerations

Under normal network operation, the snooping switch is expected to improve overall network performance by limiting the scope of multicast flooding to a smaller portion of the local network. In the event of forged IGMP messages, the benefits of using a snooping switch might be reduced or eliminated.

Security considerations for IGMPv3 at the network layer of the protocol stack are described in [IGMPv3]. The introduction of IGMP snooping functionality does not alter the handling of multicast packets by the router as it does not make use of link layer information.

There are, however, changes in the way that the IGMP snooping switch handles multicast packets within the local network. In particular:

-A Query message with a forged source address which is less than that of the current Querier could cause snooping switches to forward subsequent Membership reports to the wrong network interface. It is for this reason that IGMP Membership Reports should be sent to all multicast routers as well as the current Querier.

-It is possible for a host on the local network to generate Current-State Report Messages that would cause the switch to incorrectly believe that there is a multicast listener on the same network segment as the originator of the forged message. This will cause unrequested multicast packets to be forwarded into the network segments between the source and the router. If the router

requires that all Multicast Report messages be authenticated as described in section 9.4 of [IGMPv3], it will discard the forged Report message from the host inside the network in the same way

that it would discard one which originates from a remote location. It is worth noting that if the router accepts unauthenticated Report messages by virtue of them having arrived over a network interface associated with the internal network, investigating the affected network segments will quickly narrow the search for the source of the forged messages.

-As noted in [IGMPv3], there is little motivation for an attacker to forge a Membership report message since joining a group is generally an unprivileged operation. The sender of the forged Membership report will be the only recipient of the multicast traffic to that group. This is in contrast to a shared LAN segment (HUB) or network without snooping switches, where all other hosts on the same segment would be unable to transmit when the network segment is flooding the unwanted traffic.

The worst case result for each attack would remove the performance improvements that the snooping functionality would otherwise provide. It would, however, be no worse than that experienced on a network with switches that do not perform multicast snooping.

7. Acknowledgements

We would like to thank Martin Bak, Les Bell, Yiqun Cai, Ben Carter, Paul Congdon, Toerless Eckert, Bill Fenner, Brian Haberman, Edward Hilquist, Hugh Holbrook, Kevin Humphries, Isidor Kouvelas, Pekka Savola, Suzuki Shinsuke, Jaff Thomas, Rolland Vida, and Margaret Wasserman for comments and suggestions on this document.

Furthermore, the following companies are acknowledged for their contributions: 3Com, Alcatel, Cisco Systems, Enterasys Networks, Hewlett-Packard, Vitesse Semiconductor Corporation, Thrane & Thrane. The ordering of these names do not necessarily correspond to the column numbers in the response table.

Authors' Addresses

Morten Jagd Christensen
Thrane & Thrane
Lundtoftegaardsvej 93 D
2800 Lyngby
DENMARK

EMail: mjc@tt.dk

Karen Kimball
Hewlett-
Packard
8000 Foothills Blvd.
Roseville, CA 95747
USA

EMail: karen.kimball@hp.com

Frank Solensky
Calix
43 Nanog Park
Acton, MA 01720
USA

EMail: frank.solensky@calix.com