White Paper

Introduction to IGMP for IPTV Networks

Scott Shoaf
Consulting Engineer

Marc Bernstein
IPTV Solutions Architect

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
408 745 2000 or 888 JUNIPER
www.juniper.net

Part Number: 200188-001 June 2006

Contents

Executive Summary

This document provides a basic overview of the Internet Group Multicast Protocol (IGMP) and how it can be implemented by various elements in a broadband video architecture. This document describes basic IGMP operation.

In an IPTV network, traditional broadcast television channels are delivered via IP multicasting. IGMP is the control mechanism used to control the delivery of multicast traffic to interested and authorized users. IGMP commands tell the upstream equipment to stop sending ("leave") one channel or begin sending ("join") another channel. Depending on the architectural choices, this process occurs in the DSLAM, an aggregation switch, or at an edge router such as Juniper's E-series.

IPv6, Multicast Listener Discovery (MLD, the IPv6 equivalent to IGMP) and multicast routing protocols such as PIM are outside the scope of this document.

Terms and Acronyms

ASM	Any Source Multicast allows a multicast receiver to listen to all traffic sent to a multicast group, regardless of who is sending the information
BSR	Broadband Services Router used for subscriber management and edge routing
IGMP	Internet Group Membership Protocol is a host-router signaling protocol for IPv4 used to support IP multicasting
IPTV	IP Television is the capability to delivery broadcast TV services using an IP network
MLD	Multicast Listener Discovery is a host-router signaling protocol for IPv6
MRC / MRT	Maximum Response Code (Time) is the longest time that an IGMP-aware device will wait for a response
PIM	Protocol Independent Multicast is a multicast routing protocol for delivery multicast in a routed environment
OIF	Outgoing InterFace is used by multicast functions within a router to determine which egress ports will be used for forwarding each multicast group
RG	Routing (or Residential) Gateway is a firewall, NAT, routing device used as CPE termination in the home/office
SSM	Single Source Multicast allows a multicast receiver to listen to only the specific identified sender within a multicast group
STB	Set Top Box is the end host used to receive IPTV video
VoD	Video on Demand is a unicast streaming video offering providing an isolated video session per user with rewind, pause and similar VCR-like capabilities

Baseline IPTV Model

Figure 1 depicts the baseline DSL access network supporting IPTV service. There are up to five network elements involved. At the subscriber site, the television set [or more precisely, a set-top box] initiates channel change requests and responds to status inquiries. The routing gateway (RG) at the subscriber's site and DSLAM aggregate traffic from multiple subscribers and may act on requests from the STB. Some networks include an Ethernet switch to provide an additional layer of aggregation. Finally, the edge router (BSR, such as Juniper's E320) is the gateway into the backbone network.

Figure 1: Basic IPTV Topology

IGMP Overview

Before discussing the options available in a multicast-enabled access network, it is first helpful to understand how IGMP operates. The sections below give a brief overview of IGMPv2 and IGMPv3 when used in an IPTV architecture.

Basic IGMP operation involves two devices:

- IGMP host (or client), which issues messages to join or leave a multicast group. The client also responds to queries from the multicast router. A set-top box is an example of an IGMP host.
- IGMP router, which responds to the join and leave messages to determine if multicast groups should be forwarded out an interface. Periodic queries are used to recover from error conditions and verify requests. The IGMP router will receive multicast groups, either through the use of a multicast protocol such as PIM or static flooding. It is the termination point for IGMP messages, so does not send any IGMP information to its upstream neighbors. The IGMP router is also called a multicast router.

For this discussion, think of the STB as the IGMP host and the BSR as the IGMP router. The capabilities of the intermediate devices are discussed later.

- IGMP provides three basic functions for IP multicast networks:
- JOIN: An IGMP host indicates that it wants to receive information from ("become a member of") a multicast group.
- LEAVE: An IGMP host indicates that it no longer wishes to receive information from a multicast group.

- QUERY: An IGMP router can ask the hosts which groups they are members of. This is done to verify a JOIN/LEAVE request or to look for error conditions. For example, a set-top box may be been unplugged so did not issue a LEAVE command.

In an IPTV network, each broadcast television channel is an IP multicast group. The subscriber changes the channel by LEAVE-ing one group and JOINing a different group.

There are two versions of IGMP used for IPTV, version 2 (IGMPv2) and version 3 (IGMPv3). Both provide these basic functions, although the commands to do these differ.

IGMPv2

IGMPv2 is defined by IETF RFC 2236. It is the earliest IGMP version used to support IPTV. The key advantage of IGMPv2 over the older IGMPv1 is the addition of the Leave Group Message allowing a host to immediately signal its request to leave a multicast group. This is a major requirement in IPTV applications where bandwidth is at a premium and the forwarding of inactive multicast groups across a broadband connection can lead to periods of congestion and poor video quality.

IGMPv2 is designed to support Any Source Multicast (ASM) networks. In an ASM network, the IGMP host specifies the multicast group that it wishes to join, and listens to all traffic in that group *regardless of who is sending the traffic.*

IGMPv2 is backwards compatible with IGMPv1.

IGMPv2 Frames

The packet format for IGMPv2 is shown below. It is a simple packet format containing the Type of packet, a Maximum Response Time (MRT), Checksum, and Group Address.

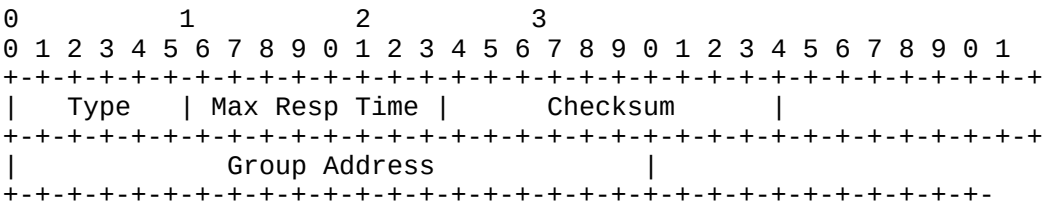

Figure 2: IGMPv2 Packet Format

The Type field indicates what message is being sent, and is discussed below. The Maximum Response Time indicates how long the sender should wait for a response. The Group Address specifies which multicast group this frame refers to.

There are three types of IGMPv2 packets (ignoring IGMPv1 backwards compatibility):

- *Membership Report*, which is used by a host to JOIN a group or to respond to Membership Queries. Membership Reports may be sent by an IGMP host without prompting (unsolicited) or in response to a Membership Query (solicited). For IGMPv2, Membership Reports have a Type of 0x16.[1]
- *Leave Group Message*, which is used by a host to explicitly LEAVE a multicast group. Leave Group Messages have a Type of 0x17.

[1] IGMPv1 Membership Reports have a Type field of 0x12. For backward compatibility, IGMPv2 supports IGMPv1 commands.

- *Membership Query*, which is sent by the multicast router to determine if any hosts are listening to a group. All membership Queries have the Type field set to 0x11. There are two types of Membership Queries.
  - *General Query* asks whether any host is listening to any group
  - *Group-Specific Query* is used to determine whether any host is listening to a specific multicast group.

IGMPv2 makes use of two reserved multicast addresses:

- 224.0.0.1 is used by the IGMP router to send messages to "all multicast hosts" (otherwise known as "all systems").
- 224.0.0.2 is used by the IGMP host to send messages to "all multicast routers".

IGMPv2 Operation

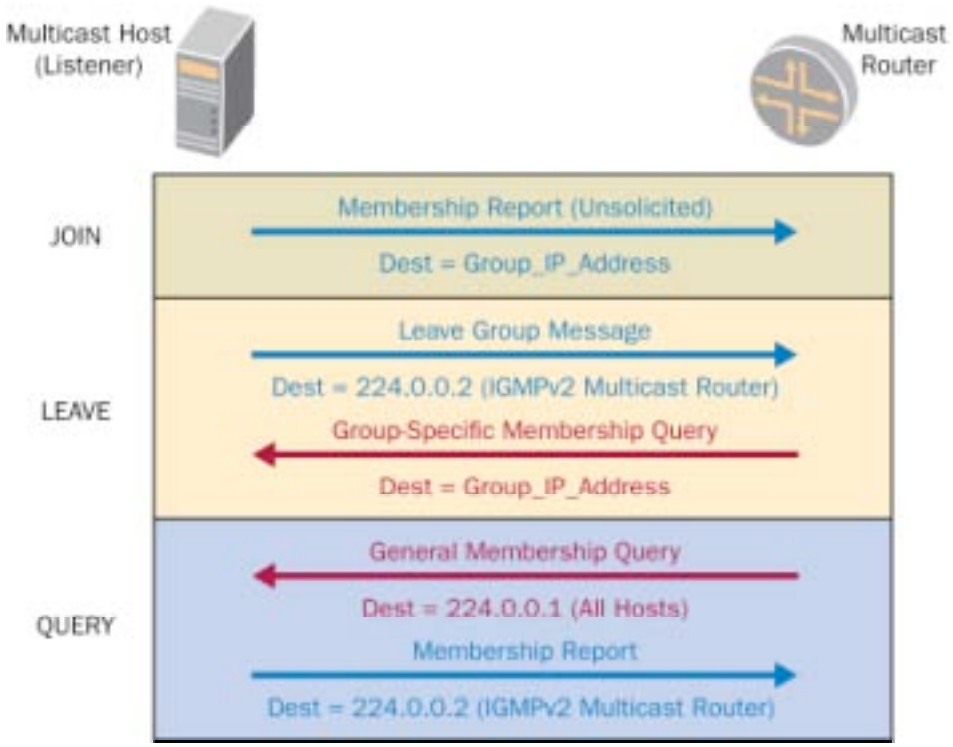

Figure 3: Common Operations Using IGMPv2

Figure 3 depicts basic IGMPv2 operation for each of the three basic functions.

- JOIN: When an IGMPv2 host needs to join a group [such as the STB wishing to view a new channel], it will send an unsolicited Membership Report destined to the group it wishes to join. Upon receiving the membership report, the multicast router will begin forwarding the channel out the appropriate interface (if it is not already being forwarded out that interface).
- LEAVE: When an IGMPv2 host leaves a group [such as the STB is no longer watching a channel], it will send out a Leave Group Message to the 224.0.0.2 multicast router address listing the group to leave in the Group Address field. The router will immediately respond on the logical interface with a Group-Specific Query to the multicast group address to determine if any hosts are still wishing to receive this specific multicast group. The MRT value will determine the timeout period to wait for a host response to the query. If no response is received within the MRT interval, the router will no longer forward that multicast group out that specific interface.
- QUERY: The General Membership Query is issued periodically to the all-hosts 224.0.0.1 multicast address to determine which groups are currently being used by IGMPv2 hosts. The MRT is used to specify the timeout interval that the router will wait for host responses to the query. If no host responds during the MRT interval, then packet forwarding stops for any groups not required out an interface. The General Membership Query is typically used in IPTV to recover from error conditions such as a STB being powered off and not able to send a Leave Group message or Leave Group messages being dropped by the access elements. This provides a self-healing mechanism for IGMP to synchronize multicast state within the network.

IGMPv3

The major enhancement in IGMPv3 is support for Single Source Multicast (SSM). When using SSM, the host *specifies the source address* that it will listen to. In other words, a multicast group with the IP address of 224.10.10.3 which is receiving traffic from a source device of 192.168.10.1 is a different multicast group than the same group IP address receiving traffic from a different source IP address. This is an important security enhancement since it prevents clients such as a set-top box from receiving traffic generated by other subscribers on the network. The use of source-specific multicast (SSM) will not be discussed in this document, but is a major driving factor for moving to IGMPv3.

In addition, the request to change channels can be done using a single IGMPv3 request instead of requiring separate “leave one channel” and “join another channel” requests. This speeds up the process of changing channels, which is a critical capability in IPTV networks.

IGMPv3 Frames

IGMPv3 is defined by IETF RFC 3376. IGMPv3 uses two types of IGMP messages:

- *Membership Report*, which use a different format than the Membership Reports in IGMPv2. To distinguish these, IGMPv3 Membership Reports set the Type field to 0x22. In IGMPv3, Membership Reports are used to JOIN and LEAVE multicast groups. The Leave Group Message is not used in IGMPv3.
- *Membership Query*, consisting of a General Query, Group-Specific Query, or *Group-and-Source-Specific Query*. The Group-and-Source-Specific Query is new in IGMPv3. Membership queries use the same Type (0x11) as older versions, although the format of the packet has changed.

IGMPv3 makes use of two reserved multicast addresses:

- 224.0.0.1 is the IP address used to send messages to "all multicast hosts". This is the same address used in IGMPv2.
- 224.0.0.22 is the IP address used to send messages to the "multicast router". This is a *different* address than used by IGMPv2.

Figure 4 depicts the frame format for an IGMPv3 Membership Query. The first entries for Type, Max Response Code (MRC, same as IGMPv2 MRT), Checksum, and Group Address match IGMPv2 for backwards compatibility. The packet then extends to include source address information to support source-specific multicast (SSM). The query packet may be targeted at the all hosts 224.0.0.1 multicast group address for General Queries or may be sent to a specific multicast group address when querying for members of that respective group during a Group-Specific or Source-and-Group-Specific query.

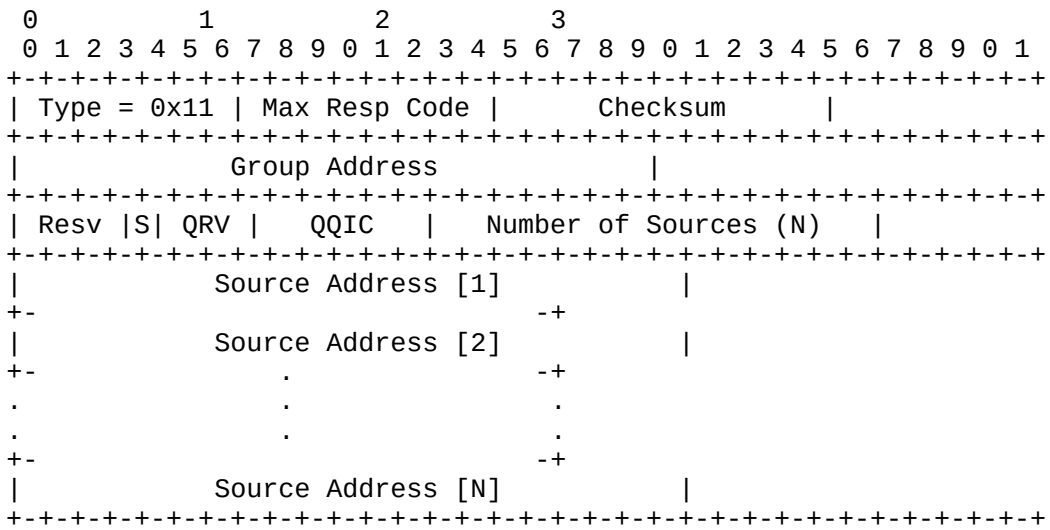

Figure 4: IGMPv3 Membership Query and Packet Format

Figure 5 shows the IGMPv3 Membership Report format. Instead of a single purpose packet, the Membership Report is sent to the 224.0.0.22 group address acting as an out-of-band signaling message. The packet consists of one or more Group Records, each record detailing specific IGMPv3 state information.

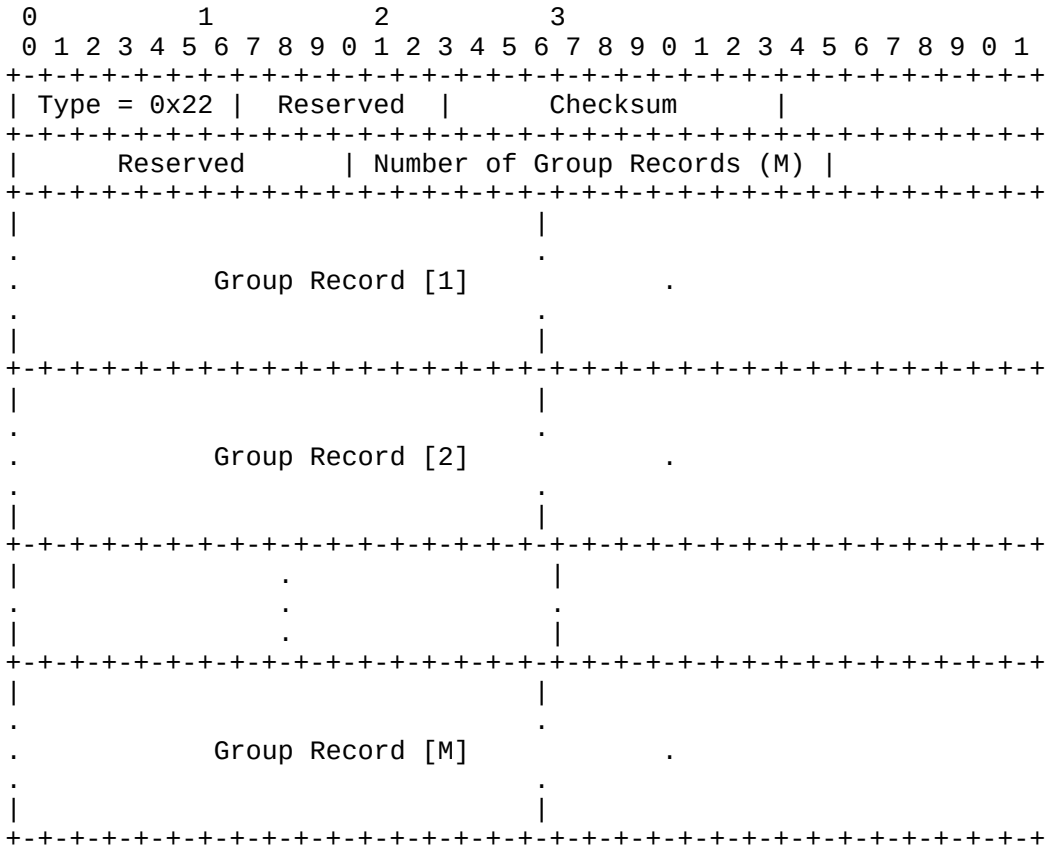

Figure 5: IGMPv3 Membership Report Format

The IGMPv3 Group Record shown below is used for Join, Leave, or Query response requirements. Multiple Group Records can be contained in a single Membership Report. As an example, both a join and leave from an STB may be include in a single Membership Report utilizing two Group Records.

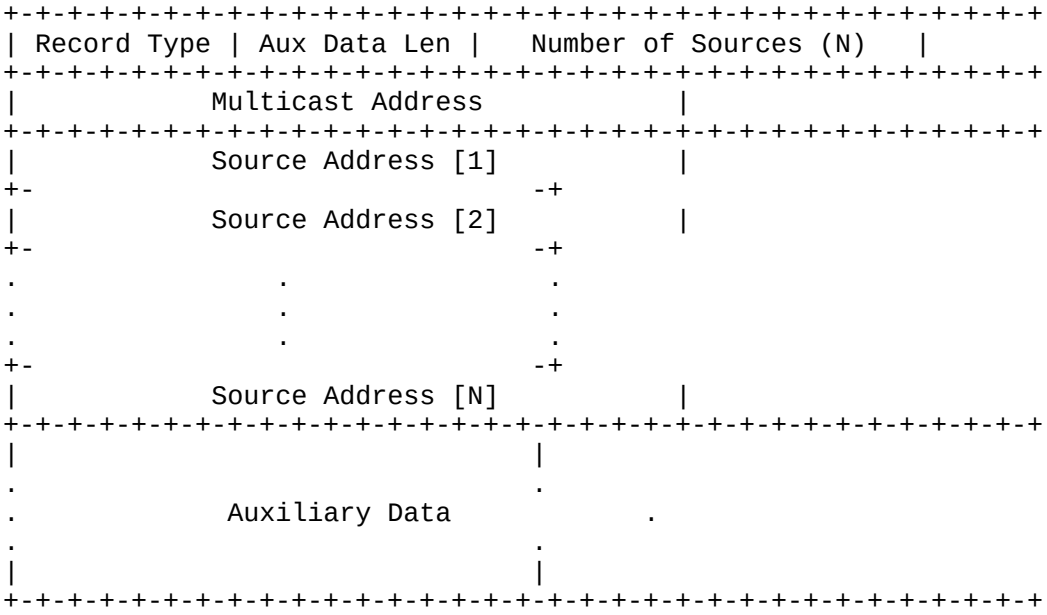

Figure 6: IGMPv3 Group Record Format

IGMPv3 Operation

Figure 7 shows basic IGMP operation when using IGMPv3.

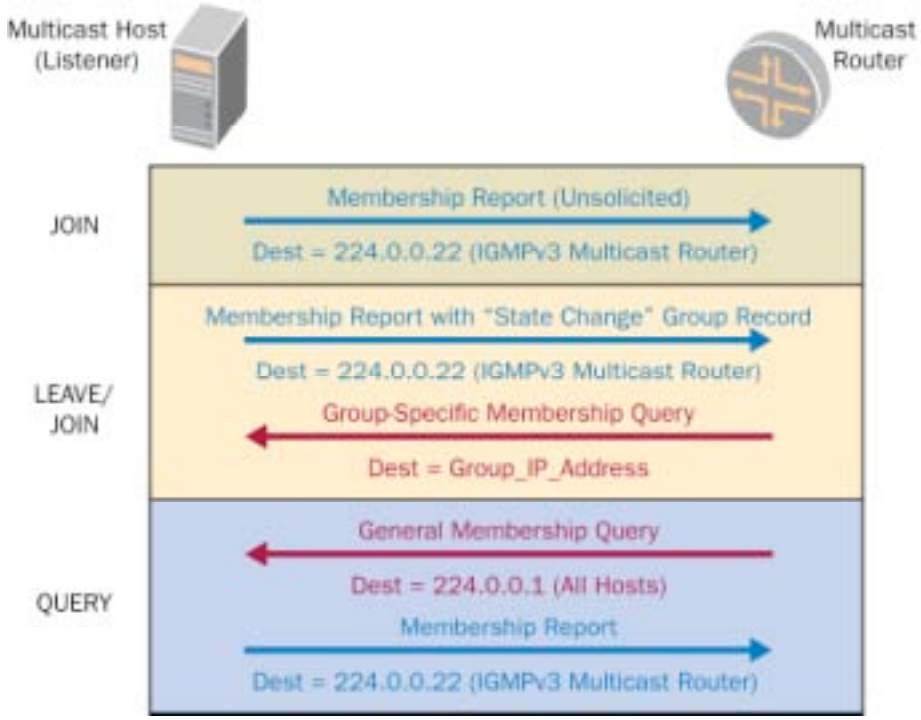

Figure 7: Common Operations Using IGMPv3

- JOIN: Like its predecessor, IGMPv3 sends an unsolicited Membership Report to join a multicast group. Upon receiving the Membership Report, the router will place the IGMP host address in the outgoing interface list for the multicast group and begin forwarding the group out the host's respective interface (if not already being forwarded).
- LEAVE: When an IGMPv3 host wishes to leave a group, it will send out a Membership Report including a State Change Record message to the IGMPv3 multicast address (224.0.0.22) which *excludes* the source address of groups no longer wanting to be received. Excluding the current source for a multicast group will result in that multicast group no longer being joined, similar to an IGMPv2 Leave Group Message.

  The router will immediately respond on the logical interface with a Group-Specific Query or Group-and-Source-Specific Query to the all-hosts 224.0.0.1 multicast address with a defined MRC to determine if any hosts are still wishing to receive this specific multicast group. If no response is received within the MRC period, the router will no longer forward that multicast group out that specific interface.
- QUERY: Like IGMPv2, IGMPv3 periodically issues the General Membership Query to determine which groups are currently being used by IGMPv3 hosts per logical network segment. This is used in IPTV to recover from error conditions such as a STB being powered off and not able to send a "leave" message.

Leaving a Multicast Group

One of the top challenges concerning IPTV is the time it takes to change the channel. To address this, some routers implement an *immediate leave* function.

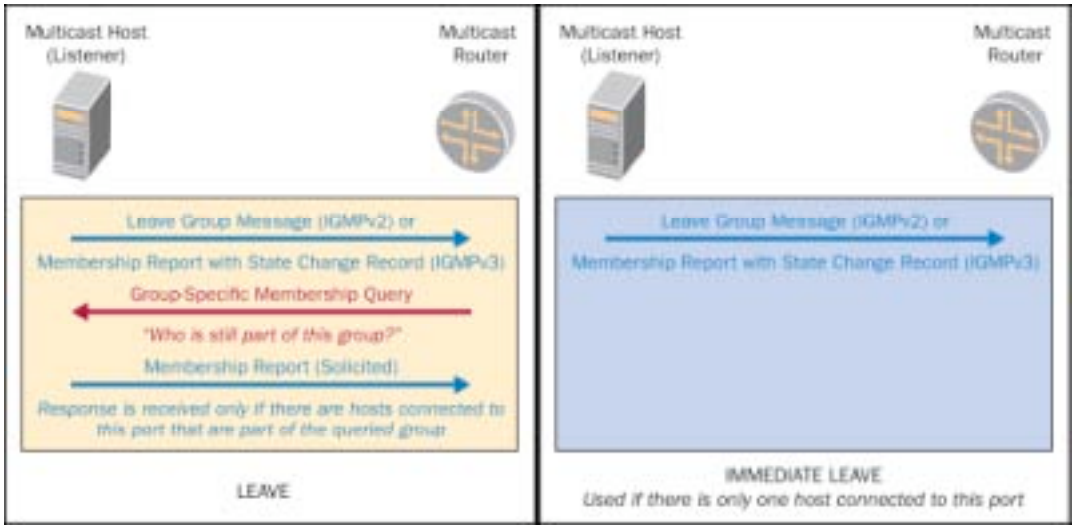

Figure 8: Leaving a Group: Standard and Immediate Leave

Standard LEAVE operation requires the following to occur before the channel is changed:

1. The IGMP host sends a request to leave one multicast group
2. The IGMP router responds by issuing Membership Query, effectively asking for confirmation for this request
3. The IGMP host responds with a Membership Report, confirming the request to the IGMP router
4. The IGMP router receives the Membership Report, interprets it and responds by changing which multicasts groups it forwards out the interface

Immediate leave overrides the normal checks to see if there are other hosts or proxy devices on the local segment interested in the multicast group. In other words, steps 2 and 3 above are skipped. Therefore, immediate leave should only be utilized where only a single host or proxy device is on that interface.

Comparing IGMPv2 and IGMPv3

The following table compares IGMPv2 and IGMPv3:

Function	Messages	Implications
Join	IGMPv2—STB issues a Membership Report message. The destination IP address is the multicast group to be joined. IGMPv3—A Membership Report is issued by the STB. The destination address is the "all IGMPv3 multicast routers" address (224.0.0.22). The "State Change" Group Record field indicates the group(s) to be joined.	• Existing intermediate devices (DSLAMs/RGs) may not support IGMPv3 since packet format is different • Packet filters based on group IP address are no longer valid • Deep packet inspection is required to read Group Records. This implies higher processing and memory requirements. An IGMPv2 device may not be upgradeable to support IGMPv3.
Leave	IGMPv2—STB issues a Leave Group message sent to the "all IGMPv2 multicast routers" address (224.0.0.2). IGMPv3—A Membership Report is issued by the STB. The destination address is the "all IGMPv3 multicast routers" address (224.0.0.22). The "State Change" Group Record field indicates the group(s) to be left.	
Channel Change (Leave and Join)	IGMPv2—Requires two separate messages as above—Leave Group followed by Membership Report. IGMPv3—One Membership report can include both Leave and Join information. The IGMP host and router only have to process a single message during each channel change.	• Channel changes should occur faster when using IGMPv3, assum.1ing there are no processing or memory bottlenecks.
Single Source Multicast	IGMPv2—Listens to all messages sent to the group IP address IMGPv3—Listens only to messages sent to the group IP address AND sent by the desired source IP address	• Can reuse group IP addresses where appropriate. For example, subscribers in New York and Los Angeles can both listen to the same multicast group IP address representing "channel 4" and receive different information (channels).

Table 1: IGMPv2/v3 comparison

Adding an Intermediate Device

In the earliest IGMP networks, devices sitting between the IGMP client and the IGMP router had no understanding of the IGMP flows. Consider the simple example in Figure 9. In this case, the top STB has requested to view channel 4, and the DSLAM forwards the request to the edge router. In response, the edge router begins forwarding the multicast group associated with channel 4. If the DSLAM had no visibility to the IGMP flows, it would not know which downstream port to which this traffic should be sent. By default, most switches broadcast incoming multicast traffic to all ports. In this example, the bottom viewer receives this channel which was not requested.

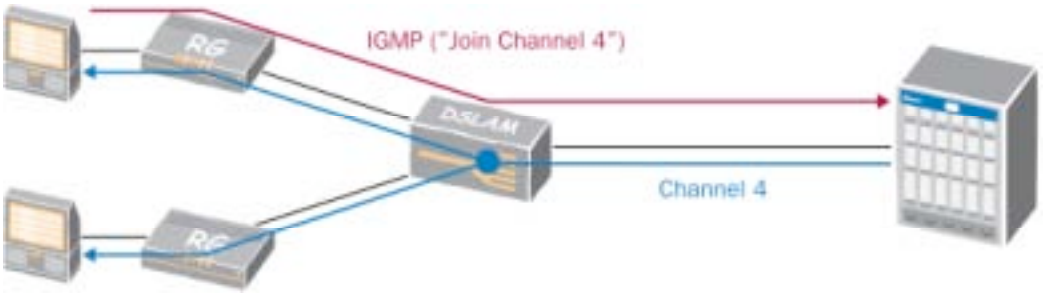

Figure 9: DSLAM without IGMP Visibility

In early networks this was not an issue as multicast usage was low and the intermediate devices were typically LAN switches supporting 100 Mbps Ethernet. However, as IPTV requires large bandwidth (typically 4 Mbps per channel) and bandwidth is limited, it becomes important to ensure that IPTV channels are forwarded only to those subscribers currently viewing them.

To resolve this, all DSLAMs and other intermediate nodes now have some level of understanding of IGMP flows. They examine the flows and build an Outbound Interface (OIF) table. Figure 10 shows a simple example of an OIF table. The OIF table allows the DSLAM to remember where to forward each multicast group.

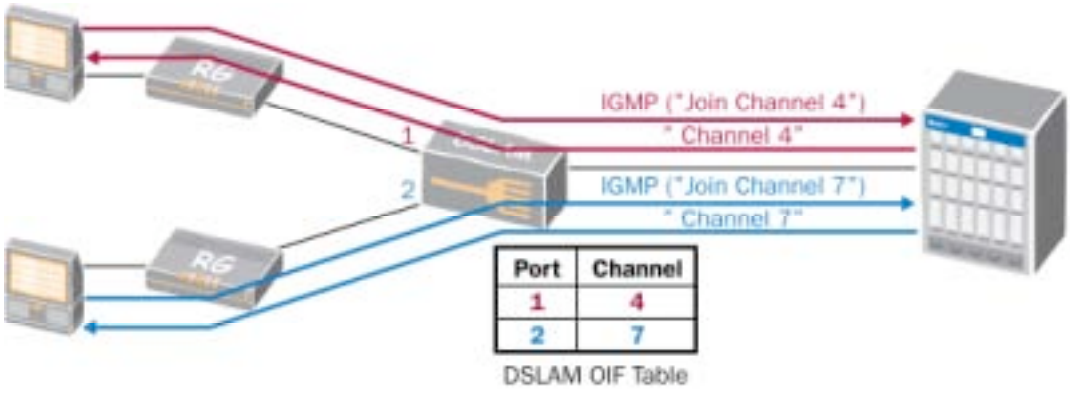

Figure 10: DSLAM with IGMP Visibility

There are two subsystems that an intermediate device such as a DSLAM may use to build this OIF table:

- IGMP (Transparent) Snooping. With this method, the DSLAM "quietly listens" to flows but does not alter the packet in any way. Each flow retains the original source and destination IP address.
- IGMP Proxy. This method makes the DSLAM participate in every IGMP flow. The IGMP host communicates with the DSLAM running IP Proxy, and the DSLAM communicates with the IGMP router.

IGMP Snooping

IGMP snooping allows the intermediate device to monitor IGMP traffic. The snooping agent will add an interface to its OIF table whenever a join request is received. In a similar manner, the snooping agent can stop sending multicast traffic when a leave is received. The snooping agent must also keep some state regarding general Membership Query Maximum Response Time timers in the event a LEAVE message is not issued from an IGMP client (such as when the power cord is pulled out on an IPTV set-top box).

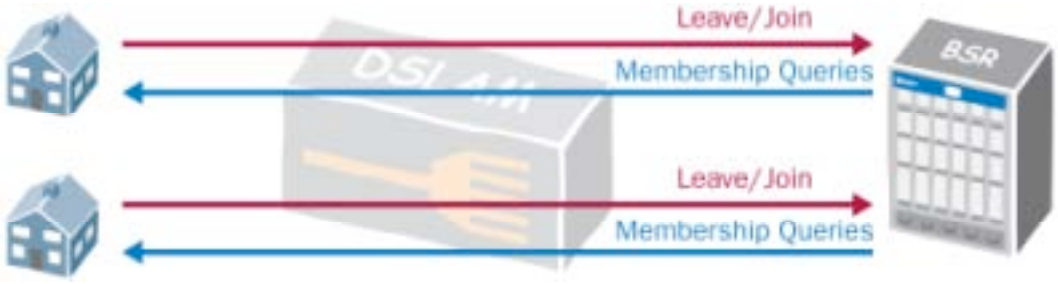

Figure 11: IGMP Transparent Snooping

Since the snooping device is transparent, IGMP packets are read and then forwarded upstream to the multicast router. The snooping device does not participate in the IGMP host messaging and promiscuously listens to transactions between clients and routers to determine when join/leave processing is required to a downstream host. An exception is noted if the snooping agent opts to intercept Membership Reports based on local filters used to prevent a host from joining specific groups. In the case of video, these filters may dictate specific broadcast channels allocated to multicast groups to be blocked from reception by the end STB.

The snooping device will rely on one of multiple mechanisms for the actual reception of the multicast data from its upstream multicast neighbor. The router may be statically configured to flood all multicast groups downstream to the snooping device, the upstream router may only forward groups based on IGMP Membership Reports received from the IGMP hosts, or the snooping agent may invoke an IGMP client process to source its own Membership Reports that are sent to the multicast router. This variation of models leads to a wide array of how a snooping device may be used in the broadband access network.

The later sections that detail the IGMP snooping function in a broadband network will show the various options for a DSLAM to receive its multicast packets to be replicated and forwarded to the respective hosts.

The key defining element of "true" IGMP snooping is transparency. All IGMP packets are forwarded through the intermediate device and these packets are never modified. As a result, the upstream IGMP router has full visibility to each downstream device.

IGMP Proxy

A device performing IGMP proxy acts in a dual mode as an IGMP router and IGMP client.

When the IGMP host issues a join message, the proxy will receive the join and add the interface to its outgoing interface list for a specific multicast group. A General Membership Query timer and state will be used by the proxy to send general queries downstream to all multicast enabled interfaces. When a leave is received, the proxy will be responsible for issuing a group-specific query and removing the interface from the outgoing interface list if no hosts respond within the configured response time interval. ***When interacting with the IGMP hosts, the proxy appears as an IGMP router.***

If the proxy is already receiving a group from an upstream router it will not issue a join message upstream. However, if a downstream host joins a group not currently received by the proxy, the proxy will issue its own join upstream to the multicast router. When a group is no longer required by downstream hosts, the proxy will issue a leave message upstream to stop the flow of packets destination for that multicast group. The proxy will also respond to Membership Queries sent from the multicast router. ***When interacting with the multicast router, the proxy appears as an IGMP client.***

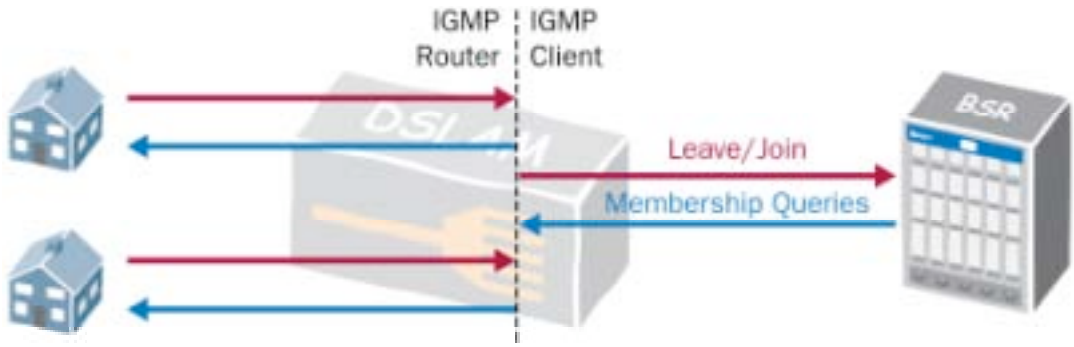

Figure 12: IGMP Proxy

A device functioning as an IGMP Proxy participates in every IGMP flow. This requires much more processing power and memory at the DSLAM, but potentially saves some upstream bandwidth.

To the multicast router, it appears as if one device (the DSLAM) is joining and leaving multicast groups. As a result, the IGMP router has no knowledge of what each subscriber is watching.

Variations

Some DSLAMs implement IGMP subsystems that have characteristics of both snooping and proxy. Most commonly, the DSLAM may decide whether to forward IGMP packets (like IGMP Proxy) but do not modify the source IP address (like IGMP Snooping).

These DSLAMs may refer to their capabilities as either IGMP Snooping or IGMP Proxy. These are generally non-standard implementations which should be avoided.

Juniper Networks Support for IGMP

The Juniper Networks E-series platform supports all of the edge routing capabilities discussed in this document. The E-series has the performance to support IPTV. In addition, they support the following features to enhance delivering IPTV service:

- Full-fledged multicast router for IPv4 and IPv6 supporting multicast routing protocols such as PIM-SSM, PIM-SM and MBGP
- Functions as an multicast router and proxy supporting IGMPv2 and IGMPv3 traffic (for IPv4) as well as MLDv1 and MLDv2 traffic (for IPv6)
- Dynamic QoS adjustment based on IGMP join/leave processing
- IGMP accounting
- Multicast call admission control
- Layer 2 control (L2C)

These features allow a service provider to explore and implement any of the models discussed in this document when using the E-series as the BSR.

Summary

IGMP can be used in many ways to allow the network to be aware of channel change requests and forward the appropriate channels. Each architecture has its own benefits and drawbacks based on the multicast optimization required in the network and the ability of various devices to offering IGMP processing.

Juniper supports each architecture option with the E-series family of edge routing products. The E-series provides both a full feature of multicast capabilities coupled with a robust hardware platform to provide high performance multicast replication and forwarding.

Contact

Marc Bernstein
mbernstein@juniper.net
978-589-0651