White Paper

Introduction to IGMP for IPTV Networks

Understanding IGMP Processing in the Broadband Access Network

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408.745.2000
1.888 JUNIPER
www.juniper.net

Part Number: 200188-003 Oct 2007

Table of Contents

Executive Summary

This document provides an overview of the Internet Group Multicast Protocol (IGMP) and how it can be supported by various elements in an IP-based television (IPTV) network. It is targeted at technical managers looking to introduce IPTV service using a wireline broadband network.

In an IPTV network, 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. This paper focuses on presenting the available alternatives when offering an IGMP service. It does not attempt to position when each should be used. That topic is covered in a related Juniper Networks document, Wireline Broadband Access Considerations.

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

Glossary

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
PIM	Protocol Independent Multicast is a multicast routing protocol for delivery multicast in a routed environment
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

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 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 (or multicast 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 receives multicast groups either through the use of a multicast protocol such as PIM or via static flooding. It is the termination point for IGMP messages, so does not send any IGMP information to its upstream neighbors.

For this discussion, think of the STB as the IGMP host and the BSR as the IGMP router. IGMP provides four 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. Queries may be:
  - Specific Query: Asks whether the host is a member of a specified multicast group
  - General Query: Asks the host to indicate all groups that it belongs to
- MEMBERSHIP REPORT: An IGMP host tells the IGMP host what groups it belongs to. This report can be either:
  - Solicited Membership Report: Sent in response to a QUERY
  - Unsolicited Membership Report: Initiated by the client

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.

IGMP Versions

There are three versions of IGMP. IGMP version 1 (IGMPv1—RFC 1112) is not used for IPTV because it does not include an explicit "LEAVE" capability. The client will continue to receive all requested streams until the multicast router issues the next query. In response, the client will issue a membership report which does not include the multicast group which it wants to leave, so the router will stop sending this stream to this client.

IGMP version 2 (IGMPv2—RFC 2236) and version 3 (IGMPv3—RFC 3376) can both be used for IPTV. Like its predecessor, IGMPv2 supports Any Source Multicast (ASM) networks. In an ASM network, the IGMP host specifies the multicast group that it wishes to join, and receives all traffic with the specified multicast address regardless of who is sending the traffic. Most deployed IPTV clients (set top boxes) support IGMPv2.

The major enhancement in IGMPv3 is support for Source Specific 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 an important driving factor for moving to IGMPv3. IGMPv3 is backwards-compatible with IGMPv2. Emerging standards from the DSL Forum (www.dslforum.org) and ATIS (www.atis.org) use IGMPv3.

Table 1 summarizes the major differences between IGMP v2 and v3:

IGMP v2 (RFC 2236)	IGMP v3 (RFC 3376)	Notes
Joining and Leaving Channels		
JOIN • STB issues a Membership Report. • The destination IP address is the multicast group to be joined.	JOIN • STB issues a Membership Report. • The destination address is the "all IGMPv3 routers" address. • 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 destination IP address are no longer valid • Deep packet inspection is required to read IGMPv3 Group Records. This implies higher processing and memory requirements. An IGMPv2 device may not be upgradeable to support IGMPv3. • Channel changes should occur faster when using IGMPv3 since a single message contains both join and leave information
LEAVE • STB issues a Leave Group • The destination IP address is the "all IGMPv2 multicast routers" address (224.0.0.2).	LEAVE • 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 • Requires two separate messages— Leave Group followed by Membership Report (to join the new channel).	CHANNEL CHANGE • One Membership report includes both Leave and Join information. The IGMP host and router only have to process a single message during each channel change.	
Special IP addresses		
• All IGMPv2 routers: 224.0.0.2 • All IGMP hosts (clients): 224.0.0.1	• All IGMP v3 routers: 224.0.0.22 • All IGMP hosts (clients): 224.0.0.1 (same as IGMPv2)	• All IGMPv3 implementation are backwards compatible with IGMPv2, so recognize the IGMPv2 "all IGMP routers" address also

Table 1: IGMPv2/v3 Comparison

Adding an Intermediate Node

Figure 1 depicts the baseline DSL access network supporting IPTV service. 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 aggregate traffic using Ethernet switches or CO-based DSLAMs. Finally, the broadband services router (BSR, such as Juniper's E320) is the gateway into the backbone network.

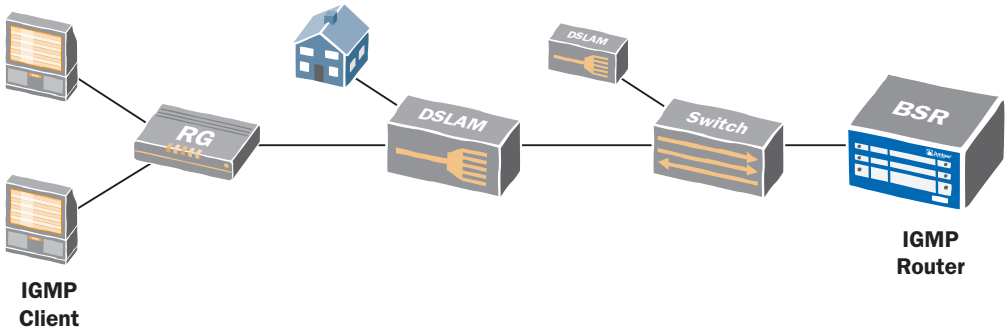

Figure 1: Typical Wireline Broadband Topology Supporting IPTV Service

IGMP networks were not originally designed to have networking equipment between the IGMP client and the IGMP router. Introducing these intermediate devices results in some challenges:

- Excessive multicast traffic: When delivering high bandwidth IPTV (often 4 Mbps or more) over broadband networks (which often support 10 Mbps or less), it is important to ensure that IPTV channels are forwarded only to those subscribers currently viewing them. If the intermediate devices are unaware of IGMP flows, then by default all multicast traffic is broadcast out all ports.
- Excessive IGMP traffic: Extra bandwidth is consumed as IGMP flows are broadcast to across the network. For example, a single IGMP query can result in responses from thousands of IGMP clients.
- Excessive time to LEAVE a group: With little spare bandwidth in broadband networks, leave requests often must be processed before additional join requests can be honored.

To mitigate these issues, several techniques have been developed which are implemented in these intermediary devices. We will discuss the DSLAM, but the discussion is applicable to all devices. These enhancements include:

1. Local replication. The DSLAM can inspect incoming IGMP Join/Leave requests and take appropriate action. If the channel being requested to view is already being received at the DSLAM, it can replicate the stream and forward to the new requester. The DSLAM builds a table to track which channels are being forwarded to each port. Figure 2 shows a simple example.

   Local replication distributes the IPTV forwarding decision to the DSLAM. This ensures good response time, but the DSLAM must ensure that enough bandwidth is available to service the request.

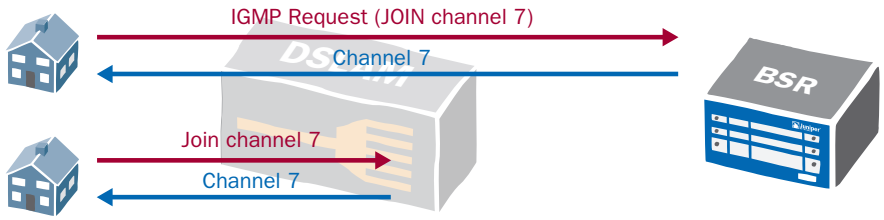

Figure 2: DSLAM with IGMP Snooping

2. Proxy routing. The intermediate device (such as a DSLAM) takes an active role in the network by terminating all IGMP flows. To the hosts (clients), the DSLAM appears to be the IGMP host, subsuming and responding to all incoming requests. To the clients, the DSLAM appears as an IGMP client, terminating and responding to IGMP flows as appropriate. When the DSLAM must forward an incoming IGMP flow, it recreates the IGMP request and uses its own IP address as the source.

   The benefit to proxy routing is scaling. When implemented in the DSLAM (as shown in Figure 3), the BSR does not learn about individual subscribers. This is not a major benefit for most existing operators since the BSR sees each DSLAM as a single subnet. However, implementing proxy routing in the RG hides downstream clients from the DSLAM, potentially increasing scalability (or reducing the cost) of the DSLAM.

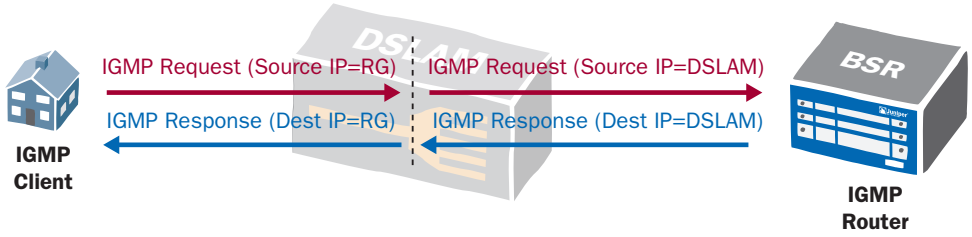

Figure 3: IGMP Proxy Routing

3. Proxy reporting: This function reduces the number of IGMP packets flowing across the network by selectively dropping IGMP packets which do not need to be forwarded. This is subdivided according to the type of IGMP packets being affected:

   - Query suppression. This reduces traffic between the DSLAM and the subscriber premises, by having the DSLAM intercept and respond to IGMP queries sent by the router. This is illustrated in Figure 4. Typically this means that the DSLAM will:
     - Never send a specific query to any client, and
     - Forward general queries only to those clients receiving at least one multicast group.

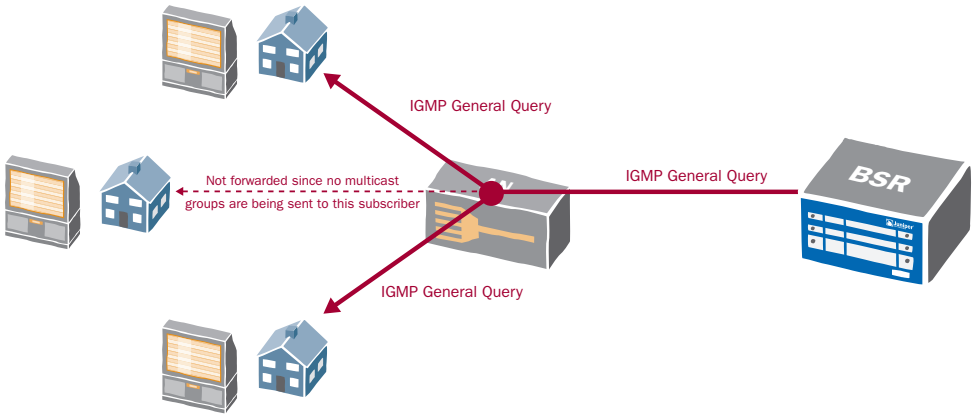

Figure 4: Query Suppression

- Report suppression. This reduces traffic from the DSLAM to the BSR by having the DSLAM aggregate the responses (membership reports) from multiple clients. The DSLAM can intercept IGMP reports coming from IGMP hosts, and forwards a summarized version to the IGMP router only when necessary. Typically this means that the DSLAM will forward IGMP membership reports as follows:
  - Unsolicited membership reports (channel change requests) are forwarded only the first subscriber joins a multicast group, or the last subscriber leaves a multicast group. This tells the IGMP router to begin or stop sending this channel to this DSLAM.
  - Solicited membership reports (sent in response to an IGMP query) are forwarded once per multicast group. The DSLAM may also aggregate multiple responses together into a single membership report, as depicted in Figure 5.

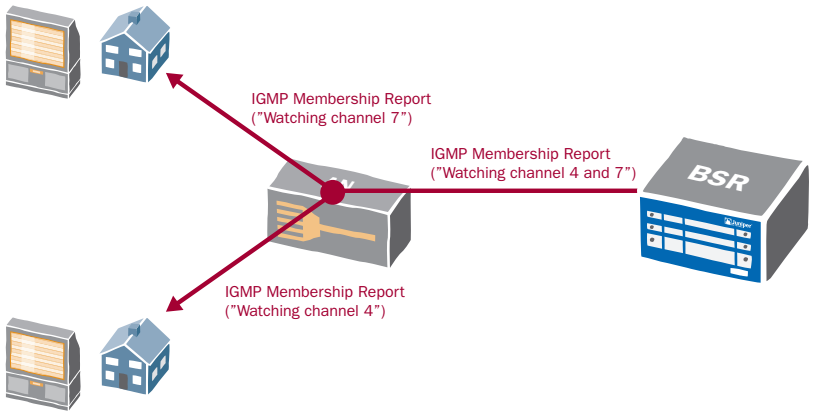

Figure 5: Report Suppression

- Last leave. This is conceptually similar to report suppression, since the AN only forwards Leave requests upstream when the last subscriber (connected to that AN) requests to leave the multicast group. If there are other subscribers still viewing the channel, the AN drops the request after responding locally.

4. Immediate leave: Another challenge concerning IPTV is the time it takes to change the channel. To address this, some multicast routers implement an immediate leave function. This function reduces the time required for a device such as a STB to leave a multicast group. Figure 6 illustrates this function.

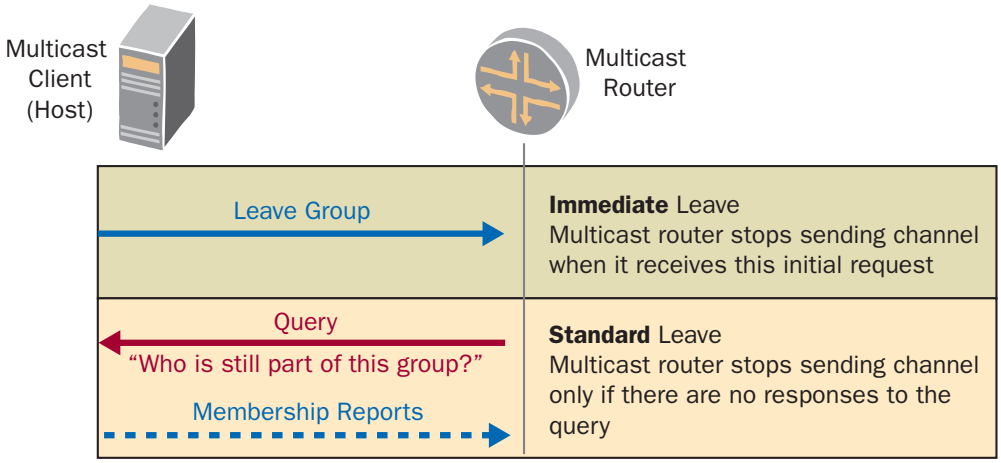

Figure 6: Standard and Immediate Leave

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

- The IGMP host sends a request to leave one multicast group
- The IGMP router responds by issuing Membership Query, effectively asking for confirmation for this request
- The IGMP host responds with a Membership Report which does not include the multicast group

Immediate leave skips steps 2 and 3, overriding the normal checks to see if there are other devices on the local segment interested in the multicast group. Therefore, it typically is only used when there is a single host (or proxy routing device) attached to an interface. It can also be used when the multicast router can track each host joined to a multicast group (on an interface), so can perform the immediate leave when the last remaining host issues a LEAVE.

IGMP Support in IPTV Networks

While each of these options can be implemented independently, there are three common combinations:

- IGMP Passthrough. In this case, the DSLAM does not inspect or process IGMP requests. The BSR forwards channels to each subscriber as required.
- IGMP (Transparent) Snooping. With this method, the DSLAM supports local replication (as described above). It inspects all IGMP packets but does not alter the packet in any way, and all packets are forwarded upstream. This method allows the upstream network to have full visibility into what is occurring in the access network.
- IGMP Proxy (local replication, proxy routing and proxy reporting). This method implements all of the above capabilities. Few flows are forwarded upstream by the DSLAM, and those that are use the source IP address of the DSLAM. Because of this, the upstream network is unaware of most activity occurring in the access network. In addition, most RGs implement this.

Some DSLAMs may implement other mixes of capabilities. For example, the DSLAM may selectively decide whether to forward IGMP membership reports (report suppression) but does not modify the source IP address (no proxy routing).

Figure 7 shows the most common IGMP support in an IPTV network.

- In the RG, IGMP Proxy is typically implemented. Proxy routing hides the downstream clients—TVs, PCs, and gaming devices—from the upstream network. This decreases the size of the forwarding tables, allowing the network to support more subscribers.
- In the aggregation switch, IGMP Passthrough is typically deployed. This provides a consistent operational model regardless of whether the DSLAM is connected to the BSR directly or via an aggregation switch. It is common to directly connect the DSLAM to the BSR in denser areas, while using switches to aggregate traffic from smaller DSLAMs.
- In the DSLAM, any combination of capabilities can be deployed depending upon requirements .

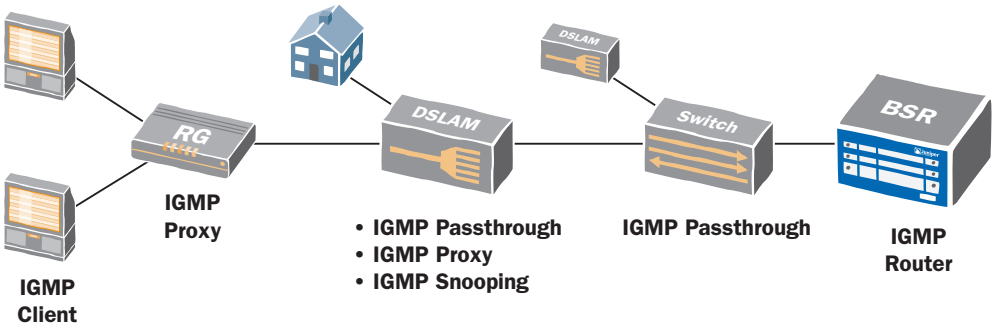

Figure 7: Common IGMP Implementations for IPTV

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 alternative has its own benefits and drawbacks.

As an IGMP (multicast) host, Juniper Networks E-series Broadband Services Routers (BSR) supports all of the relevant capabilities discussed in this document. In addition, they support extensive capabilities to enhance delivering IPTV service, including:

- Multicast router for IPv4 and IPv6 supporting PIM-SSM, PIM-SM and MBGP
- Support for IGMPv2/v3 (for IPv4) and MLDv1/v2 (for IPv6)
- Dynamic QoS adjustment based on IGMP join/leave processing
- IGMP accounting
- Dynamic multicast bandwidth measurement and multicast call admission control (CAC)
- Source mapping, which allows IGMP clients to be used with an SSM backbone.

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.

For further information on these and other design alternatives, contact your Juniper Networks sales representative, visit us at www.juniper.net or call us at 866-298-6428 (USA, Canada and Mexico) or 978-589-0500 (outside USA).

Related Reading

The following materials provide additional information on IGMP and related multicast protocols.

- TR-101: Migration to Ethernet Based DSL Aggregation, available at http://www.dslforum.org/techwork/tr/TR-101.pdf . This DSL Forum Technical Report describes how Ethernet can be used to provide broadband service over DSL lines.
- RFC 2236: Internet Group Management Protocol, Version 2
- RFC 3376: Internet Group Management Protocol, Version 3
- RFC 4541: IGMP and MLD Snooping Switches Considerations. This is the latest version of http://www.ietf.org/internet-drafts/draft-ietf-magmasnoop-12.txt, which is referenced in TR-101.
- RFC 4605: Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding (“IGMP/MLD Proxying”). This is the latest version of http://www.ietf.org/internet-drafts/draft-ietf-magma-igmpproxy-06.txt, which is referenced in TR-101.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.