Internet Engineering Task Force
INTERNET-DRAFT
draft-ietf-magma-igmp-proxy-00.txt

Bill Fenner, AT&T Research
Haixiang He, Nortel Networks
Brian Haberman, Nortel Networks
Hal Sandick, Nortel Networks

Expire: May, 2002
November, 2001

IGMP-based Multicast Forwarding ("IGMP Proxying")

Status of this Memo

Abstract

In certain topologies, it is not necessary to run a multicast routing protocol. It is sufficient to learn and proxy group membership information and simply forward based upon that information. This draft describes a mechanism for forwarding based solely upon IGMP membership information.

This document is a product of the IDMR working group within the Internet Engineering Task Force. Comments are solicited and should be addressed to the working group's mailing list at idmr@cs.ucl.ac.uk and/or the authors.

1. Introduction

This document applies spanning tree multicast routing[Deering91] to an IGMP-only environment. The topology is limited to a tree, since we specify no protocol to build a spanning tree over a more complex topology. The root of the tree is assumed to be connected to a wider multicast infrastructure.

1.1. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [Bradner97].

2. Definitions

2.1. Upstream Interface

A router's interface in the direction of the root of the tree. Also called the "Host interface".

2.2. Downstream Interface

Each of a router's interfaces that is not in the direction of the root of the tree. Also called the "Router interfaces".

2.3. Group Mode

For each multicast group, a group is in IGMPv1 mode if an IGMPv1 report is heard. A group is in IGMPv2 mode if an IGMPv2 report is heard but no IGMPv1 report is heard. A group is in IGMPv3 mode if an IGMPv3 is heard but no IGMPv1 or IGMPv2 report is heard.

2.4. Subscription

When a group is in IGMPv1 or IGMPv2 mode, the subscription is a group membership on an interface. When a group is in IGMPv3 mode, the subscription is an IGMPv3 state entry (i.e. a (multicast address, group timer, filter-mode, source-element list) tuple) on an interface.

2.5. Membership Database

The database maintained at each router into which the membership information of each of its downstream interfaces is merged.

3. Abstract protocol definition

A router performing IGMP-based forwarding has a single upstream interface and one or more downstream interfaces. These designations are explicitly configured; there is no protocol to determine what type each interface is. It performs the router portion of the IGMP [Deering89, Fenner97, CDFKT01] protocol on its downstream interfaces, and the host portion of IGMP on its upstream interface. The router MUST NOT perform the router portion of IGMP on its upstream interface.

The router maintains a database consisting of the merger of all subscriptions on any downstream interface. Refer to section 4 for the details about the construction and maintenance of the membership database.

The router sends IGMP membership reports on the upstream interface when queried, and sends unsolicited reports or leaves when the database changes.

When the router receives a packet destined for a multicast group, it uses a list consisting of the upstream interface and any downstream interface which has a subscription pertaining to this packet and on which it is the IGMP Querier. This list may be built dynamically or cached. It removes the interface on which this packet arrived from the list and forwards the packet to the remaining interfaces.

Note that the rule that a router must be the querier in order to forward packets restricts the IP addressing scheme used; in particular, the IGMP-based forwarding routers must be given the lowest IP addresses of any potential IGMP Querier on the link, in order to win the IGMP Querier election. If another device wins the IGMP Querier election, no packets will flow.

Forwarder election is necessary for links which are considered to be downstream links by multiple IGMP-based forwarders. This rule "piggy-backs" forwarder election on IGMP Querier election. On a link with only one IGMP-based forwarding router, this rule MAY be disabled (i.e. the router MAY be configured to forward packets to an interface on which it is not the querier). However, the default configuration MUST include the querier rule.

Note that this does not protect against an "upstream loop". For example, as shown in the figure below:

B will unconditionally forward packets from LAN 1 to LAN 2, and A will unconditionally forward packets from LAN 2 to LAN 1. This will cause an upstream loop. A multicast routing protocol which employs a tree building algorithm is requried to resolve loops like this.

3.1. Topology Restrictions

This specification describes a protocol that works only in a simple tree topology. The tree must be manually configured by designating upstream and downstream interfaces on each router, and the root of the tree is expected to be connected to a wider multicast infrastructure.

3.2. Supporting Senders

In order for senders to send from inside the proxy tree, all traffic is forwarded towards the root. The multicast router(s) connected to the wider multicast infrastructure should be configured to treat all systems inside the proxy tree as though they were directly connected -- e.g., for PIM-SM, these routers should Register-encapsulate traffic from new sources within the proxy tree just as they would directly-connected sources.

This information is likely to be manually configured; IGMP-based multicast forwarding provides no way for the routers upstream of the proxy tree to know what networks are connected to the proxy tree. If the proxy topology is congruent with some routing topology, this information MAY be learned from the routing protocol running on the topology; e.g. a router may be configured to treat multicast packets from all prefixes learned from routing protocol X via interface Y as though they were from a directly connected system.

4. Router Behavior

This section describes an IGMP-based multicast forwarding router's actions in more detail.

4.1. Membership Database

The router performs the router portion of the IGMP protocol on each downstream interface. For each interface, the version of IGMP used is explicitly configured and default to the highest version supported by the system.

The output of this protocol is a set of subscriptions; this set is maintained separately on each downstream interface. In addition, the subscriptions on each downstream interface are merged into the membership database.

The membership database is a set of membership records of the form:

(multicast-address, filter-mode, source-list)

Each record is the result of the merge of all subscriptions for that record's multicast-address on downstream interfaces. If some subscriptions are IGMPv1 or IGMPv2 subscriptions, these subscriptions are converted to IGMPv3 subscriptions. The IGMPv3 subscriptions and the converted subscriptions are merged using the merging rules for multiple memberships on a single interface specified in the IGMPv3 specification[CDFKT01] to create the membership record. For example, there are two downstream interfaces I1 and I2 that have subscriptions for multicast address G. I1 has an IGMPv2 subscription that is (G). I2 has an IGMPv3 subscription that is (G, INCLUDE, (S1, S2)). The I1's subscription is converted to (G, EXCLUDE, NULL). Then the subscriptions are merged and final membership record is (G, EXCLUDE, NULL).

The router performs the host portion of the IGMP protocol on upstream interface. If there is an IGMPv1 or IGMPv2 querier on upstream network, then the router will perform IGMPv1 or IGMPv2 on upstream interface accordingly. Otherwise, it will perform IGMPv3.

If the router performs IGMPv3 on upstream interface, then when the composition of the membership database changes, the change in the database is reported on the upstream interface as though this router were a host performing the action. If the router performs IGMPv1 or IGMPv2 on upstream interface, then when the membership records are created or deleted, the changes are reported on the upstream interface. All other changes are ignored. When the router reports using IGMPv1 or IGMPv2, only the multicast address field in the membership record is used.

4.2. Forwarding Packets

A router forwards packets received on its upstream interface to each downstream interface based upon the downstream interface's subscriptions and whether or not this router is the IGMP Querier on each interface. A router forwards packets received on any downstream interface to the upstream interface, and to each downstream interface other than the incoming interface based upon the downstream interfaces' subscriptions and whether or not this router is the IGMP Querier on each interface. A router MAY use a forwarding cache in order not to make this decision for each packet, but MUST update the cache using these rules any time any of the information used to build it changes.

4.3. SSM Considerations

To support Source-Specific Multicast (SSM), the router should be compliant with the specification about using IGMPv3 for SSM [HC01]. Note that the router should be compliant with both the IGMP Host Requirement and the IGMP Router Requirement for SSM since it performs IGMP Host Portion on upstream interface and IGMP Router Portion on each downstream interface.

An interface can be configured to perform IGMPv1 or IGMPv2. In this scenario, the SSM semantic will not be maintained for that interface. However, a router that supports this document should ignore those IGMPv1 or IGMPv2 subscriptions sent to SSM addresses. And more importantly, the packets with source-specific addresses SHOULD not be forwarded to interfaces with IGMPv2 or IGMPv1 subscriptions for these addresses.

5. Security Considerations

Since only the Querier forwards packets, the IGMP Querier election process may lead to black holes if a non-forwarder is elected Querier. An attacker on a downstream LAN can cause itself to get elected Querier resulting in no packets being forwarded.

References

Bradner97 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119/BCP 14, Harvard University, March 1997.

CDFKT01 Cain, B., S. Deering, B. Fenner, I. Kouvelas and A. Thyagarajan, "Internet Group Management Protocol,

Version 3", Work in progress. (draft-ietf-idmr-igmp-v3-07.txt)

Deering91 Deering, S., "Multicast Routing in a Datagram Internetwork", Ph.D. Thesis, Stanford University, December 1991.

Fenner97 Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, Xerox PARC, November 1997.

Deering89 Deering, S., "Host Extensions for IP Multicasting", RFC 1112, August 1989.

HC01 Holbrook, H., and Cain, B., "Using IGMPv3 For Source-Specific Multicast", draft-holbrook-idmr-igmpv3-ssm-01.txt, March 2001.

Author's Address:

William C. Fenner
AT&T Labs - Research
75 Willow Rd
Menlo Park, CA 94025
Phone: +1 650 330 7893
Email: fenner@research.att.com

Haixiang He
Nortel Networks
600 Technology Park Drive
Billerica, MA 01821
Phone: 978-288-7482
Email: haixiang@nortelnetworks.com

Brian Haberman
Nortel Networks
300 Perimter Park
Morrisvile, NC 27560
Email: haberman@nortelnetworks.com

Hal Sandick
Nortel Networks
300 Perimter Park
Morrisvile, NC 27560
Email: hsandick@nortelnetworks.com