

Network Working Group
Request for Comments: 2337
Category: Experimental

D. Farinacci
Cisco Systems
D. Meyer
Cisco Systems
Y. Rekhter
Cisco Systems
April 1998

Intra-LIS IP multicast among routers over ATM using Sparse Mode PIM

Status of this Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

2. Abstract

This document describes how intra-LIS IP multicast can be efficiently supported among routers over ATM without using the Multicast Address Resolution Server (MARS). The method described here is specific to Sparse Mode PIM [PIM-SM], and relies on the explicit join mechanism inherent in PIM-SM to notify routers when they should create group specific point-to-multipoint VCs.

3. Overall model

This document focuses on forwarding of multicast traffic among PIM-SM routers connected to an ATM network. Routers on an ATM network are partitioned into Logical IP Subnets, or LISs. This document deals with handling multicast within a single LIS. Handling inter-LIS multicast traffic, including handling shortcuts, is outside the scope of this document. In addition, this document does not address forwarding of multicast traffic to or from hosts connected to an ATM network.

4. Router behavior

This document requires that each router within a LIS knows IP and ATM addresses of all other routers within the LIS. The mapping between IP and ATM addresses may be provided by an ARP server [RFC2225], or by any other means (e.g., static configuration).

Each PIM router within a LIS is required to maintain a single (shared) point-to-multipoint distribution VC rooted at the router with all other PIM routers in the LIS as the leaf nodes. The VC is expected to be used for forwarding of multicast traffic (both data and control) among routers within the LIS. For example, this VC would be used for distributing PIM [PIM-SM] control messages (Join/Prune messages).

In addition, if a PIM router receives a IGMP report from a non-PIM neighbor, then the router may add the reporter to the existing shared distribution VC or to the group specific distribution VC (if it exists). The PIM router may also create a specific VC for this IGMP proxy.

4.1. Establishing Dedicated, Per Group Point-to-Multipoint VCs

Routers may also maintain group specific, dedicated point-to-multipoint VCs. In particular, an upstream router for a group may choose to become the root of a group specific point-to-multipoint VC whose leaves are the downstream routers that have directly connected or downstream receivers for the group. While the criteria for establishing a group specific point-to-multipoint VC are local to a router, issues such as the volume of traffic associated with the group and the fanout factor within the LIS should be considered. Finally, note that a router must minimally support a single shared point-to-multipoint VC for distribution of control messages and data (to all group addresses).

A router can choose to establish a dedicated point-to-multipoint VC (or add another leaf to an already established dedicated point-to-multipoint VC) when it receives a PIM Join or IGMP report messages from another device in the same LIS. When a router that is the root of a point-to-multipoint VC receives PIM Prune message or IGMP leave, it removes the originator of the message from its dedicated point-to-multipoint VC.

4.2. Switching to a Source-Rooted Tree

If at least one of the routers within a LIS decides to switch to a source-rooted tree (by sending (S,G) PIM Joins), then all other routers within the LIS that have downstream members for G should switch to that source-rooted tree as well. Since a router that switches to a source-rooted tree sends PIM Join messages for (S,G) over its shared point-to-multipoint VC, the other routers within the LIS are able to detect this. Once a router that has downstream members for G detects this, the router should send (S,G) PIM Join message as well (otherwise the router may receive duplicate traffic from S).

Note that it is possible for a non-PIM router in the LIS to fail to receive data if the injection point moves to router to which there is not an existing VC.

4.2.1. Adding New Members to a Source-Rooted Tree

As mentioned above, this document requires that once one router in a LIS decides to switch to the source tree for some (S,G), all routers in the LIS that have downstream members must also switch to the (S,G) source tree. Now, when a new router wants to receive traffic from G, it starts sending (*,G)-Joins on it's shared point-to-multipoint VC toward the RP for G. The root of the (S,G)-source-rooted tree will know to add the new router to the point-to-multipoint VC servicing the (S,G)-source-rooted tree by observing the (*,G)-joins on it's shared point-to-multipoint VC. However, the new router must also switch to the (S,G)-source-rooted tree. In order to accomplish this, the newly added router must:

- (i). Notice that it has been added to a new point-to-multipoint VC
- (ii). Notice (S,G) traffic coming down this new point-to-multipoint VC
- (iii). Send (S,G) joins toward S, causing it to switch to the source-rooted tree. The router learns that the VC is used to distribute (S,G) traffic in the previous steps.

4.3. Handling the "Packet Reflection" Problem

When a router receives a multicast packet from another router in its own LIS, the router should not send the packet on any of the routers distribution point-to-multipoint VCs associate with the LIS. This eliminates the problem of "packet reflection". Sending the packet on the routers' distribution VCs associated with other LISs is controlled by the multicast routing procedures.

5. Brief Comparison with MARS

The intra-LIS multicast scheme described in this document is intended to be a less complex solution to an important subset of the functionality provided by the Multicast Address Resolution Server, or MARS [MARS]. In particular, it is designed to provide intra-LIS multicast between routers using PIM-SM, and does not consider the case of host-rooted point-to-multicast multicast distribution VCs.

Although MARS supports both of the current schemes for mapping the IP multicast service model to ATM (multicast server and meshes of point-to-multipoint VCs), it does so at at cost and complexity higher than of the scheme described in this document. In addition, MARS requires new encapsulations, whereas this proposal works with either LLC/SNAP or with NLPID encapsulation. Another important difference is that MARS allows point-to-multipoint VCs rooted either at a source or at a multicast server (MCS). The approach taken here is to constrain complexity by focusing on PIM-SM (taking advantage of information available in explicit joins), and by allowing point-to-multipoint VCs to be rooted only at the routers (which is roughly analogous to the complexity and functionality of rooting point-to-multipoint VCs at the sources).

In summary, the method described in this document is designed for the router-to-router case, and takes advantage of the explicit-join mechanism inherent in PIM-SM to provide a simple mechanism for intra-LIS multicast between routers. MARS, on the other hand, accepts different tradeoffs in complexity-functionality design space. In particular, while the MARS paradigm provides a general neighbor discovery mechanism, allows host to participate, and is protocol independent, it does so at considerable cost.

6. Security Considerations

In general, the security issues relevant to the proposal outlined in the memo are subsumed by those faced by PIM-SM. While work in proceeding on security for PIM-SM, it is worthwhile noting that several issues have been raised in conjunction with multicast routing and with PIM-SM in particular. These issues include but are not limited to:

- (i). Unauthorized Senders
- (ii). Unauthorized Receivers
- (iii). Unauthorized use of the RP
- (iv). Unauthorized "last hop" switching to shortest path tree.

6.1. General Comments on Multicast Routing Protocol Security

Historically, routing protocols used within the Internet have lacked strong authentication mechanisms [RFC1704]. In the late 1980s, analysis revealed that there were a number of security problems in Internet routing protocols then in use [BELLOVIN89]. During the early 1990s it became clear that adversaries were selectively attacking various intra-domain and inter-domain routing protocols (e.g. via TCP session stealing of BGP sessions) [CERTCA9501, RFC1636]. More recently, cryptographic authentication mechanisms have been developed for RIPv2, OSPF, and the proprietary EIGRP routing protocols. BGP protection, in the form of a Keyed MD5 option for TCP, has also become widely deployed.

At present, most multicast routing protocols lack strong cryptographic protection. One possible approach to this is to incorporate a strong cryptographic protection mechanism (e.g. Keyed HMAC MD5 [RFC2104]) within the routing protocol itself. Alternately, the routing protocol could be designed and specified to use the IP Authentication Header (AH) [RFC1825, RFC1826, RFC2085] to provide cryptographic authentication.

Because the intent of any routing protocol is to propagate routing information to other parties, confidentiality is not generally required in routing protocols. In those few cases where local security policy might require confidentiality, the use of the IP Encapsulating Security Payload (ESP) [RFC1825, RFC1827] is recommended.

Scalable dynamic multicast key management is an active research area at this time. Candidate technologies for scalable dynamic multicast key management include CBT-based key management [RFC1949] and the Group Key Management Protocol (GKMP) [RFC2093,RFC2094]. The IETF IP Security Working Group is actively working on GKMP extensions to the standards-track ISAKMP key management protocol being developed in the same working group.

7. References

- [BELLOVIN89] S. Bellovin, "Security Problems in the TCP/IP Protocol Suite", ACM Computer Communications Review, Volume 19, Number 2, pp. 32-48, April 1989.
- [CERTCA9501] CERT, "IP Spoofing Attacks and Hijacked Terminal Connections", ftp://ftp.cert.org/cert_advisories/, January 1995.
- [MARS] Armitage, G., "Support for Multicast over UNI 3.0/3.1 based ATM Networks.", RFC 2022, November 1996.
- [PIM-SM] Estrin, D, et. al., "Protocol Independent Multicast Sparse Mode (PIM-SM): Protocol Specification", Work in Progress.
- [RFC1636] Braden, R., Clark, D., Crocker, S., and C. Huitema, "Report of IAB Workshop on Security in the Internet Architecture February 8-10, 1994", RFC 1636, June 1994.
- [RFC1704] Haller, N., and R. Atkinson, "On Internet Authentication", RFC 1704, October 1994.
- [RFC1825] Atkinson, R., "IP Security Architecture", RFC 1825, August 1995.
- [RFC1826] Atkinson, R., "IP Authentication Header", RFC 1826, August 1995.
- [RFC1827] Atkinson, R., "IP Encapsulating Security Payload", RFC 1827, August 1995.
- [RFC1949] Ballardie, A., "Scalable Multicast Key Distribution", RFC1949, June 1996.
- [RFC2085] Oehler, M., and R. Glenn, "HMAC-MD5 IP Authentication with Replay Prevention", RFC 2085, February 1997.

- [RFC2093] Harney, H., and C. Muckenhirn, "Group Key Management Protocol (GKMP) Specification", RFC 2093, July 1997.
- [RFC2094] Harney, H., and C. Muckenhirn, "Group Key Management Protocol (GKMP) Architecture", RFC 2094, July 1997.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed Hashing for Message Authentication", RFC 2104, February 1997.
- [RFC2225] Laubach, M., and J. Halpern, "Classical IP and ARP over ATM", RFC 2225, April 1998.

8. Acknowledgments

Petri Helenius provided several insightful comments on earlier versions of this document.

9. Author Information

Dino Farinacci
Cisco Systems
170 Tasman Dr.
San Jose, CA 95134

Phone: (408) 526-4696
EMail: dino@cisco.com

David Meyer
Cisco Systems
170 Tasman Dr.
San Jose, CA 95134

Phone: (541) 687-2581
EMail: dmm@cisco.com

Yakov Rekhter
cisco Systems, Inc.
170 Tasman Dr.
San Jose, CA 95134

Phone: (914) 528-0090
EMail: yakov@cisco.com

10. Full Copyright Statement

Copyright (C) The Internet Society (1998). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.