

Internet Engineering Task Force
Internet-Draft
Obsoletes:
3913, 2189, 2201, 1584, 1585 (if
approved)
Expires: June 20, 2005

P. Savola
CSC/FUNET
December 20, 2004

**Overview of the Internet Multicast Routing Architecture
draft-savola-mboned-routingarch-01.txt**

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [section 3 of RFC 3667](#). By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with [RFC 3668](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on June 20, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004).

Abstract

The lack of up-to-date documentation on IP multicast routing protocols and procedures has caused a great deal of confusion. To clarify the situation, this memo describes the routing protocols and techniques currently (as of this writing) in use.

Table of Contents

1.	Introduction	3
1.1	Multicast-related Abbreviations	4
2.	Multicast Routing	4
2.1	Setting up Multicast Forwarding State	4
2.1.1	PIM-SM	4
2.1.2	PIM-DM	4
2.1.3	Bidirectional PIM	5
2.1.4	DVMRP	5
2.1.5	Obsolete Protocols	6
2.1.6	Interactions and Summary	6
2.2	Distributing Topology Information	6
2.2.1	Multi-protocol BGP	7
2.2.2	OSPF/IS-IS Multi-topology Extensions	7
2.2.3	Issue: Overlapping Unicast/multicast Topology	7
2.3	Learning (Active) Sources	8
2.3.1	SSM	8
2.3.2	MSDP	9
2.3.3	Embedded-RP	9
2.4	Configuring and Distributing PIM-SM RP Information	9
2.4.1	Manual Configuration with an Anycast Address	9
2.4.2	Embedded-RP	10
2.4.3	BSR and Auto-RP	10
2.5	Mechanisms for Enhanced Redundancy	11
2.5.1	Anycast RP	11
2.5.2	Stateless RP Failover	11
2.5.3	Bidirectional PIM	11
2.6	Interactions with Hosts	11
2.6.1	Hosts Sending Multicast	12
2.6.2	Hosts Receiving Multicast	12
2.7	Restricting Multicast Flooding in the Link Layer	12
2.7.1	Router-to-Router Flooding Reduction	12
2.7.2	Host/Router Flooding Reduction	13
3.	Acknowledgements	13
4.	IANA Considerations	13
5.	Security Considerations	13
6.	References	14
6.1	Normative References	14
6.2	Informative References	15
	Author's Address	17
A.	Other Issues	17
A.1	Multicast Payload Transport Extensions	17
A.1.1	Reliable Multicast	17
A.1.2	Multicast Group Security	18
	Intellectual Property and Copyright Statements	19

1. Introduction

Good, up-to-date documentation of IP multicast is close to non-existent. This issue severely felt with multicast routing protocols and techniques. The consequence is that those who wish to learn of IP multicast and how the routing works in the real world do not know where to begin.

The aim of this document is to provide a brief overview of multicast routing protocols and techniques.

This memo deals with:

- o setting up multicast forwarding state ([Section 2.1](#)),
- o distributing multicast topology information ([Section 2.2](#)),
- o learning active sources ([Section 2.3](#)),
- o configuring and distributing the PIM-SM RP information ([Section 2.4](#)),
- o mechanisms for enhanced redundancy ([Section 2.5](#)),
- o interacting with hosts ([Section 2.6](#)), and
- o restricting the multicast flooding in the link layer ([Section 2.7](#)).

Multicast data delivery issues are also introduced in [Appendix A.1](#).

This memo obsoletes and re-classifies to Historic Border Gateway Multicast Protocol (BGMP), Core-based Trees (CBT), Multicast OSPF (MOSPF) RFCs: [[RFC3913](#)], [[RFC2189](#)], [[RFC2201](#)], [[RFC1584](#)], and [[RFC1585](#)].

1.1 Multicast-related Abbreviations

ASM	Any Source Multicast
BGMP	Border Gateway Multicast Protocol (obsolete)
BSR	Bootstrap Router
CBT	Core-based Trees (obsolete)
CGMP	Cisco Group Management Protocol
DR	Designated Router
DVMRP	Distance Vector Multicast Routing Protocol
GARP	Group Address Resolution Protocol
IGMP	Internet Group Management Protocol
MLD	Multicast Listener Discovery
MSNIP	Multicast Source Notification of Interest Protocol
MOSPF	Multicast OSPF (obsolete)
MBGP	Multi-protocol BGP (*not* "Multicast BGP")
MSDP	Multicast Source Discovery Protocol
PGM	Pragmatic General Multicast
PIM	Protocol Independent Multicast
PIM-DM	Protocol Independent Multicast - Dense Mode
PIM-SM	Protocol Independent Multicast - Sparse Mode
PIM-SSM	Protocol Independent Multicast - (Source-specific) Sparse Mode
RGMP	(Cisco's) Router Group Management Protocol
RP	Rendezvous Point
SSM	Source-specific Multicast

2. Multicast Routing

2.1 Setting up Multicast Forwarding State

The most important part of multicast routing is setting up the multicast forwarding state. This section describes the protocols commonly used for this purpose.

2.1.1 PIM-SM

By far, the most common multicast routing protocol is PIM-SM [[I-D.ietf-pim-sm-v2-new](#)]. The PIM-SM protocol includes both Any Source Multicast (ASM) and Source-Specific Multicast (SSM) functionality; PIM-SSM is a subset of PIM-SM. Most current routing platforms support PIM-SM.

2.1.2 PIM-DM

Whereas PIM-SM is designed to avoid unnecessary flooding of multicast data, PIM-DM [[I-D.ietf-pim-dm-new-v2](#)] operates in a "dense" mode, flooding the multicast transmissions throughout the network ("flood and prune") unless the leaf parts of the network periodically

indicate that they are not interested in that particular traffic.

PIM-DM may be some fit in small and/or simple networks, where setting up an RP would be unnecessary, and possibly in cases where a large number of users is expected to be able to wish to receive the transmission so that the amount of state the network has to keep is minimal. Therefore PIM-DM has typically only been used in special deployments, never currently in, e.g., ISPs' networks.

PIM-DM never really got popular due to its reliance of data plane and potential for loops, and the over-reliance of the complex Assert mechanism. Further, it was a non-starter with high-bandwidth streams.

Many implementations also support so-called "sparse-dense" mode, where Sparse mode is used by default, but Dense is used for configured multicast group ranges (such as Auto-RP in [Section 2.4.3](#)) only. Lately, many networks have been transitioned away from sparse-dense to only sparse mode.

[2.1.3](#) Bidirectional PIM

Bidirectional-PIM [[I-D.ietf-pim-bidir](#)] aims to offer streamlined PIM-SM operation, without data-driven events and data-encapsulation, inside a PIM-SM domain. The usage of bidir-PIM may be on the increase especially inside sites leveraging multicast.

As of this writing, it cannot be applied in Inter-domain multicast or for Embedded-RP. XXX: explain why not, someone send text? (one paragraph please!).

[2.1.4](#) DVMRP

Distance Vector Multicast Routing Protocol (DVMRP) [[RFC1075](#)][[I-D.ietf-idmr-dvmrp-v3](#)][[I-D.ietf-idmr-dvmrp-v3-as](#)] is the first protocol designed for multicasting, and to get around initial deployment hurdles, it also included the tunneling capabilities which were part of its multicast topology functions.

Nowadays, DVMRP is used only very rarely in operator networks, having been replaced with PIM-SM. The most typical deployment of DVMRP is at a leaf network, to run from a legacy firewall only supporting DVMRP to the internal network. However, GRE tunneling [[RFC2784](#)] seems to have overtaken DVMRP in this functionality, and there is relatively little use for DVMRP except in legacy deployments.

2.1.5 Obsolete Protocols

Three protocols which are considered completely obsolete are MOSPF [RFC1584] (which was implemented and deployed at a time, but PIM later replaced it), BGMP [RFC3913] (which did not get sufficient support to get adopted and moved forward), and CBT [RFC2201] (which was never significantly deployed, losing to PIM in the standards process).

2.1.6 Interactions and Summary

It is worth noting that is it is possible to run different protocols with different groups ranges (e.g., treat some groups as dense mode in an other-wise PIM-SM network; this typically requires manual configuration of the groups) or interact between different protocols (e.g., use DVMRP in the leaf network, but PIM-SM upstream). The basics for interactions among different protocols have been outlined in [RFC2715].

The following figure gives a concise summary of the different protocols as of this writing.

	Interdomain	Intradomain	Status
PIM-SM	Yes	Yes	De facto
PIM-DM	Not feasible	Yes	Little use
Bidir-PIM	No	Yes	Wait & see
DVMRP	Not anymore	Yes (stub)	Going out
CBT, MOSPF, BGMP	No	No	Obsolete

2.2 Distributing Topology Information

When unicast and multicast topologies are the same ("congruent"), i.e., use the same routing tables (routing information base, RIB), it has been considered sufficient to just distribute one set of reachability information.

However, when PIM -- which by default built multicast topology based on the unicast topology -- gained popularity, it became apparent that it would be necessary to be able to distribute also non-congruent multicast reachability information in the regular unicast protocols.

The topology information is needed to perform efficient distribution of multicast transmissions and to prevent transmission loops by applying it to Reverse Path Forwarding (RPF) check. XXX: does the

use of RPF need to be expanded?

This subsection introduces these protocols.

[2.2.1](#) Multi-protocol BGP

Multiprotocol Extensions for BGP-4 [[RFC2858](#)] (often referred to as "MBGP"; however, it is worth noting that "MBGP" does *not* stand for "Multicast BGP") specifies a mechanism how BGP can be used to distribute different reachability information for unicast and multicast traffic (using SAFI=2 for multicast). Multiprotocol BGP has been widely deployed for years, and is also needed for routing for example IPv6.

These extensions are in widespread use wherever BGP is used to distribute unicast topology information. Those having multicast infrastructure and using BGP should use Multiprotocol BGP to also distribute multicast reachability information explicitly even if the topologies are congruent (using SAFI=3). A number of significant multicast transit providers even require this, by doing the RPF lookups solely based on explicitly advertised multicast address family.

[2.2.2](#) OSPF/IS-IS Multi-topology Extensions

Similar to BGP, some IGPs also provide the capability for signalling a differing multicast topology, for example IS-IS multi-topology extensions [[I-D.ietf-isis-wg-multi-topology](#)]. Similar work exists for OSPF [[I-D.ietf-ospf-mt](#)].

It is worth noting that interdomain incongruence and intradomain incongruence are orthogonal, so one doesn't require the other. Specifically, interdomain incongruence is quite common, while intradomain incongruence isn't, so you see much more deployments of MBGP than MT-ISIS/OSPF. Commonly deployed networks have managed well without protocols handling intradomain incongruence. However, the availability of multi-topology mechanisms may in part replace the typically used workarounds such as tunnels.

[2.2.3](#) Issue: Overlapping Unicast/multicast Topology

An interesting case occurs when some do not distribute multicast topology explicitly while some do. In particular, this happens when some multicast sites in the Internet are using plain BGP while some use MBGP.

Different implementations deal with this using different means. Sometimes, multicast RPF mechanisms first look up the multicast

routing table, or RIB ("topology database") with longest prefix match algorithm, and if they find any entry (including a default route), that is used; if no match is found, unicast routing table is looked instead.

An alternative approach is to use longest prefix match on the union of multicast and unicast routing tables; an implementation technique here is to copy the whole unicast routing table over to the multicast routing table. The important point to remember here, though, is to not override the multicast-only routes; if the longest prefix match would find both a (copied) unicast route and a multicast-only route, the latter should be treated as preferable.

One implemented approach is to just look up the information in the unicast routing table, and provide the user capabilities to change that as appropriate, using for example copying functions discussed above.

2.3 Learning (Active) Sources

Typically, multicast routing protocols must either assume that the receivers know the IP addresses of the (active) sources for a group a priori, possibly using a rendezvous mechanism (SSM), or the sources must be discovered by the network protocols automatically (ASM).

Learning active sources is a relatively straightforward process with a single PIM-SM domain and with a single RP, but having a single PIM-SM domain is a completely unscalable model for many reasons for the whole Internet. Therefore it is required to be able to split up the multicast routing infrastructures to smaller domains, and there must be a way to share information about active sources using some mechanism, if ASM model is to be supported.

This section discusses the options.

2.3.1 SSM

Source-specific Multicast [[I-D.ietf-ssm-arch](#)] (sometimes also referred to as "single-source Multicast") does not count on learning active sources in the network; it is assumed that the recipients know these using out of band mechanisms, and when subscribing to a (S,G) channel, indicate toward which source(s) the multicast routing protocol should send the Join messages.

As of this writing, there are attempts to analyze and/or define out-of-band source discovery functions which would help SSM in particular.

[2.3.2](#) MSDP

Multicast Source Discovery Mechanism [[RFC3618](#)] was invented as a stop-gap mechanism, when it became apparent that multiple PIM-SM domains (and RPs) were needed in the network, and information about the active sources needed to be propagated between the PIM-SM domains using some other protocol.

MSDP is also used to share the state about sources between multiple RPs in a single domain for, e.g., redundancy purposes [[RFC3446](#)]. There is also work in progress to achieve the same using PIM extensions [[I-D.ietf-pim-anycast-rp](#)]. See [Section 2.5](#) for more.

There is no intent to define MSDP for IPv6, but instead use only SSM and Embedded-RP instead [[I-D.ietf-mboned-ipv6-multicast-issues](#)].

[2.3.3](#) Embedded-RP

Embedded-RP [[RFC3956](#)] is an IPv6-only technique to map the address of the RP to the multicast group address. Using this method, it is possible to avoid the use of MSDP while still allowing multiple multicast domains (in the traditional sense). The model works by defining a single RP for a particular group for all of the Internet, so there is no need to share state about that with any other RPs (except, possibly, for redundancy purposes with Anycast-RP using PIM).

[2.4](#) Configuring and Distributing PIM-SM RP Information

For PIM-SM, there exist configuration mechanisms which are used to configure the RP addresses and which groups are to use those RPs in the routers. This section outlines the approaches.

[2.4.1](#) Manual Configuration with an Anycast Address

It is often the easiest to just manually configure the RP information on the routers when PIM-SM is used.

Originally, static RP mapping was considered suboptimal since it required explicit configuration changes every time the RP address changed. However, with the advent of anycast RP addressing, the RP address is unlikely to ever change. Therefore, the administrative burden is generally limited to initial configuration. Since there is usually a fair amount of multicast configuration required on all routers anyway (eg, PIM on all interfaces), adding the RP address statically isn't really an issue. Further, static anycast RP mapping provides the benefits of RP load balancing and redundancy (see [Section 2.5](#)) without the complexity found in dynamic mechanisms like

Auto-RP and Bootstrap Router (BSR).

With such design, an anycast RP uses a "portable" address, which is configured on a loopback interfaces of the routers currently acting as RPs, as described in [[RFC3446](#)].

Using this technique, each router might only need to be configured with one, portable RP address.

[2.4.2](#) Embedded-RP

Embedded-RP provides the information about the RP's address in the group addresses which are delegated to those who use the RP, so unless no other ASM than Embedded-RP is used, one only needs to configure the RP routers themselves.

While Embedded-RP in many cases is sufficient for IPv6, other methods of RP configuration are needed if one needs to provide ASM service for other than Embedded-RP group addresses. In particular, service discovery type of applications may need hard-coded addresses that are not dependent on local RP addresses.

As the RP's address is exposed to the users and applications, it is very important to ensure it does not change often, e.g., by using manual configuration of an anycast address.

[2.4.3](#) BSR and Auto-RP

BSR [[I-D.ietf-pim-sm-bsr](#)] is a mechanism for configuring the RP address for groups. It may no longer be in as wide use with IPv4 it was earlier, and for IPv6, Embedded-RP will in many cases be sufficient.

Cisco's Auto-RP is an older, proprietary method for distributing group to RP mappings, similar to BSR. Auto-RP has little use today.

Both Auto-RP and BSR require some form of control at the routers to ensure that only valid routers are able to advertise themselves as RPs. Further, flooding of BSR and Auto-RP messages must be prevented at PIM borders. Additionally, routers require monitoring that they are actually using the RP(s) the administrators think they should be using, for example if a router (maybe in customer's control) is advertising itself inappropriately. All in all, while BSR and Auto-RP provide easy configuration, they also provide very significant configuration and management complexity.

It is worth noting that both Auto-RP and BSR were deployed before the use of a manually configured anycast-RP address became relatively

commonplace, and there is actually relatively little use for them today.

[2.5](#) Mechanisms for Enhanced Redundancy

A couple of mechanisms, already described in this document, have been used as a means to enhance redundancy, resilience against failures, and to recover from failures quickly. This section summarizes these techniques explicitly.

[2.5.1](#) Anycast RP

As mentioned in [Section 2.3.2](#), MSDP is also used to share the state about sources between multiple RPs in a single domain for, e.g., redundancy purposes [[RFC3446](#)]. The purpose of MSDP in this context is to share the same state information on multiple RPs for the same groups to enhance the robustness of the service.

There is also work in progress to achieve the same using PIM extensions [[I-D.ietf-pim-anycast-rp](#)]. This is a required method to be able to use Anycast RP with IPv6.

[2.5.2](#) Stateless RP Failover

It is also possible to use some mechanisms for smaller amount of redundancy as Anycast RP, without sharing state between the RPs. A traditional mechanism has been to use Auto-RP or BSR (see [Section 2.4.3](#)) to select another RP when the active one failed. However, the same functionality could be achieved using a shared-unicast RP address ("anycast RP without state sharing") without the complexity of a dynamic mechanism. Further, Anycast RP offers a significantly more extensive failure mitigation strategy, so today there is actually very little need to use stateless failover mechanisms, especially dynamic ones, for redundancy purposes.

[2.5.3](#) Bidirectional PIM

Bidirectional PIM (see [Section 2.1.3](#)) was designed with faster multicast convergence in mind, and if fast failover is required, it may seem like an attractive solution.

XXX: someone want to send a bit more text about the convergence differences wrt PIM-SM?

[2.6](#) Interactions with Hosts

Previous sections have dealt with the components required by routers to be able to do multicast routing. Obviously, the the real users of

multicast are the hosts: either sending or receiving multicast.

This section describes the required interactions with hosts.

[2.6.1](#) Hosts Sending Multicast

Hosts don't need to do any signalling prior to sending multicast to a group; they just send to the link-layer multicast address, and the designated router will listen to the multicast packets and start forwarding them as appropriate.

There is a proposal, Multicast Source Notification of Interest Protocol (MSNIP) [[I-D.ietf-magma-msnip](#)] which could be used for both the sending hosts to register as having interest in sending if there would be receivers, and for the receivers (or routers) to signal their interest in receiving the multicast transmission. XXX: to be removed if MSNIP is not finalized before this goes forward.

[2.6.2](#) Hosts Receiving Multicast

Hosts signal their interest in receiving a multicast group or channel by the use of IGMP [[RFC3376](#)] and MLD [[RFC3810](#)]. IGMPv2 and MLDv1 are also commonplace, but most new deployments support the latest specifications.

[2.7](#) Restricting Multicast Flooding in the Link Layer

Multicast transmission in the link layer, for example Ethernet, typically includes some form of flooding the packets through a LAN. This causes unnecessary bandwidth usage and discarding unwanted frames on those nodes which did not want to receive the multicast transmission.

Therefore a number of techniques have been developed, to be used in Ethernet switches between routers, or between routers and hosts, to limit the flooding.

These options are discussed in this section.

[2.7.1](#) Router-to-Router Flooding Reduction

Only a proprietary solution, Cisco's RGMP [[RFC3488](#)] is known to have been developed to reduce the amount of router-to-router flooding on a LAN. This is typically only considered a problem in some Ethernet-based Internet Exchange points.

There have been similar proposals to snoop PIM messages [[I-D.tsenevir-pim-sm-snoop](#)][[I-D.serbest-l2vpn-vpls-mcast](#)] to achieve

the same effect.

2.7.2 Host/Router Flooding Reduction

There are a number of techniques to help reduce flooding both from a router to hosts, and from a host to the routers (and other hosts).

Cisco's proprietary CGMP [[CGMP](#)] provides a solution where the routers notify the switches, but also allows the switches to snoop IGMP packets to enable faster notification of hosts no longer wishing to receive a group. IPv6 is not supported.

IEEE specifications mention Group Address Resolution Protocol (GARP) [[GARP](#)] as a link-layer method to perform the same functionality. Its implementation status is unknown.

IGMP snooping [[I-D.ietf-magma-snoop](#)] appears to be the most widely implemented technique, often used alongside with IGMP proxying [[I-D.ietf-magma-igmp-proxy](#)]. IGMP snooping requires that the switches implement a significant amount of IP-level packet inspection; this appears to be something that is difficult to get right, and often the upgrades are also a challenge. To allow the snooping switches to identify at which ports the routers reside (and therefore where to flood the packets) instead of requiring manual configuration, Multicast Router Discovery protocol is being specified [[I-D.ietf-magma-mrdisc](#)].

3. Acknowledgements

Tutoring a couple multicast-related papers, the latest by Kaarle Ritvanen [[RITVANEN](#)] convinced the author that the up-to-date multicast routing and address assignment/allocation documentation is necessary.

Leonard Giuliano, James Lingard, Jean-Jacques Pansiot, Dave Meyer, and Stig Venaas provided good comments, helping in improving this document.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

This memo only describes different approaches to multicast routing, and this has no security considerations; the security analysis of the mentioned protocols is out of scope of this memo.

6. References

6.1 Normative References

- [I-D.ietf-idmr-dvmrp-v3]
Pusateri, T., "Distance Vector Multicast Routing Protocol", [draft-ietf-idmr-dvmrp-v3-11](#) (work in progress), December 2003.
- [I-D.ietf-idmr-dvmrp-v3-as]
Pusateri, T., "Distance Vector Multicast Routing Protocol Applicability Statement", [draft-ietf-idmr-dvmrp-v3-as-01](#) (work in progress), May 2004.
- [I-D.ietf-isis-wg-multi-topology]
Przygienda, T., Shen, N. and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in IS-IS", [draft-ietf-isis-wg-multi-topology-07](#) (work in progress), June 2004.
- [I-D.ietf-ospf-mt]
Psenak, P., "MT-OSPF: Multi Topology (MT) Routing in OSPF", [draft-ietf-ospf-mt-00](#) (work in progress), October 2004.
- [I-D.ietf-pim-bidir]
Handley, M., Kouvelas, I., Speakman, T. and L. Vicisano, "Bi-directional Protocol Independent Multicast (BIDIR-PIM)", [draft-ietf-pim-bidir-07](#) (work in progress), July 2004.
- [I-D.ietf-pim-dm-new-v2]
Adams, A., Nicholas, J. and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", [draft-ietf-pim-dm-new-v2-05](#) (work in progress), June 2004.
- [I-D.ietf-pim-sm-v2-new]
Fenner, B., Handley, M., Holbrook, H. and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode PIM-SM): Protocol Specification (Revised)", [draft-ietf-pim-sm-v2-new-11](#) (work in progress), October 2004.
- [I-D.ietf-ssm-arch]
Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", [draft-ietf-ssm-arch-06](#) (work in progress), September 2004.

- [RFC2858] Bates, T., Rekhter, Y., Chandra, R. and D. Katz, "Multiprotocol Extensions for BGP-4", [RFC 2858](#), June 2000.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B. and A. Thyagarajan, "Internet Group Management Protocol, Version 3", [RFC 3376](#), October 2002.
- [RFC3618] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", [RFC 3618](#), October 2003.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), June 2004.
- [RFC3956] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", [RFC 3956](#), November 2004.

6.2 Informative References

- [CGMP] "Cisco Group Management Protocol", <<http://www.javvin.com/protocolCGMP.html>>.
- [GARP] Tobagi, F., Molinero-Fernandez, P. and M. Karam, "Study of IEEE 802.1p GARP/GMRP Timer Values", 1997.
- [I-D.ietf-magma-igmp-proxy] Fenner, B., He, H., Haberman, B. and H. Sandick, "IGMP/MLD-based Multicast Forwarding ('IGMP/MLD Proxying')", [draft-ietf-magma-igmp-proxy-06](#) (work in progress), April 2004.
- [I-D.ietf-magma-mrdisc] Haberman, B. and J. Martin, "Multicast Router Discovery", [draft-ietf-magma-mrdisc-03](#) (work in progress), October 2004.
- [I-D.ietf-magma-msnip] Fenner, B., "Multicast Source Notification of Interest Protocol (MSNIP)", [draft-ietf-magma-msnip-05](#) (work in progress), March 2004.
- [I-D.ietf-magma-snoop] Christensen, M., Kimball, K. and F. Solensky, "Considerations for IGMP and MLD Snooping Switches", [draft-ietf-magma-snoop-11](#) (work in progress), May 2004.
- [I-D.ietf-mboned-ipv6-multicast-issues] Savola, P., "IPv6 Multicast Deployment Issues",

[draft-ietf-mboned-ipv6-multicast-issues-01](#) (work in progress), September 2004.

[I-D.ietf-pim-anycast-rp]

Farinacci, D., "Anycast-RP using PIM",
[draft-ietf-pim-anycast-rp-02](#) (work in progress), June 2004.

[I-D.ietf-pim-sm-bsr]

Fenner, B., "Bootstrap Router (BSR) Mechanism for PIM",
[draft-ietf-pim-sm-bsr-04](#) (work in progress), July 2004.

[I-D.serbest-l2vpn-vpls-mcast]

Serbest, Y., "Supporting IP Multicast over VPLS",
[draft-serbest-l2vpn-vpls-mcast-01](#) (work in progress),
November 2004.

[I-D.tsenevir-pim-sm-snoop]

Senevirathne, T. and S. Vallepali, "Protocol Independent
Multicast-Sparse Mode (PIM-SM) Snooping",
[draft-tsenevir-pim-sm-snoop-00](#) (work in progress), April
2002.

[RFC1075] Waitzman, D., Partridge, C. and S. Deering, "Distance
Vector Multicast Routing Protocol", [RFC 1075](#), November
1988.

[RFC1584] Moy, J., "Multicast Extensions to OSPF", [RFC 1584](#), March
1994.

[RFC1585] Moy, J., "MOSPF: Analysis and Experience", [RFC 1585](#), March
1994.

[RFC2189] Ballardie, T., "Core Based Trees (CBT version 2) Multicast
Routing -- Protocol Specification --", [RFC 2189](#), September
1997.

[RFC2201] Ballardie, T., "Core Based Trees (CBT) Multicast Routing
Architecture", [RFC 2201](#), September 1997.

[RFC2715] Thaler, D., "Interoperability Rules for Multicast Routing
Protocols", [RFC 2715](#), October 1999.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D. and P. Traina,
"Generic Routing Encapsulation (GRE)", [RFC 2784](#), March
2000.

[RFC3208] Speakman, T., Crowcroft, J., Gemmell, J., Farinacci, D.,

Lin, S., Leshchiner, D., Luby, M., Montgomery, T., Rizzo, L., Tweedly, A., Bhaskar, N., Edmonstone, R., Sumanasekera, R. and L. Vicisano, "PGM Reliable Transport Protocol Specification", [RFC 3208](#), December 2001.

[RFC3446] Kim, D., Meyer, D., Kilmer, H. and D. Farinacci, "Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)", [RFC 3446](#), January 2003.

[RFC3488] Wu, I. and T. Eckert, "Cisco Systems Router-port Group Management Protocol (RGMP)", [RFC 3488](#), February 2003.

[RFC3740] Hardjono, T. and B. Weis, "The Multicast Group Security Architecture", [RFC 3740](#), March 2004.

[RFC3913] Thaler, D., "Border Gateway Multicast Protocol (BGMP): Protocol Specification", [RFC 3913](#), September 2004.

[RITVANEN]

Ritvanen, K., "Multicast Routing and Addressing", HUT Report, Seminar on Internetworking, May 2004,
<<http://www.tml.hut.fi/Studies/T-110.551/2004/papers/>>.

Author's Address

Pekka Savola
CSC - Scientific Computing Ltd.
Espoo
Finland

E-Mail: psavola@funet.fi

[Appendix A](#). Other Issues

[A.1](#) Multicast Payload Transport Extensions

A couple of mechanisms have been, and are being specified, to improve the characteristics of the data that can be transported over multicast.

These go beyond the scope of multicast routing, but as reliable multicast has some relevance, these are briefly mentioned.

[A.1.1](#) Reliable Multicast

Reliable Multicast Working Group has been working on experimental

specifications so that applications requiring reliable delivery characteristics, instead of simple unreliable UDP, could use multicast as a distribution mechanism.

One such mechanism is Pragmatic Generic Multicast (PGM) [[RFC3208](#)]. This does not require support from the routers, but PGM-aware routers may act as helpers delivering missing data.

[A.1.2](#) Multicast Group Security

Multicast Security Working Group has been working on methods how the integrity, confidentiality, and authentication of data sent to multicast groups can be ensured using cryptographic techniques [[RFC3740](#)].

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.