PIM
Internet-Draft
Obsoletes: 1112 (if approved)
Updates: 791 (if approved)
Intended status: Standards Track
Expires: 11 January 2024

S. E. Deering
Retired
T. Eckert, Ed.
Futurewei Technologies USA

10 July 2023

Host Extensions for "Any Source" IP Multicasting (ASM)
draft-eckert-pim-rfc1112bis-02

Abstract

This memo specifies the extensions required of a host implementation of the Internet Protocol (IP) to support Any Source Multicast (ASM) IP Multicasting or abbreviated IP Multicast. Distribution of this memo is unlimited.

This document replaces [RFC1112] for anything but its specification of the IGMP version 1 protocol.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on 11 January 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components

extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. STATUS OF THIS MEMO

This memo specifies the extensions required of a host implementation of the Internet Protocol (IP) to support Any Source Multicast (ASM) IP Multicasting or abbreviated IP Multicast. Distribution of this memo is unlimited.

This document replaces [RFC1112] for anything but its specification of the IGMP version 1 protocol.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. INTRODUCTION

The host extensions defined in this memo are called Any Source Multicast (ASM) IP multicast or abbreviated IP multicast. The term Any Source Multicast is used to distinguish these extensions from Source Specific Multicast (SSM) IP multicast as defined by [SSM]. The abbreviation IP multicast always refers to this memo's extensions.

This memo applies to both IP and IPv6. When it uses the term IP it implies either or both version of the IP protocol. It uses the terms IP and/or IPv6 explicitly when referring to functions applicable to only a specific version of the IP protocol.

This document replaces [RFC1112] for anything but the specification of IGMP version 1 in Appendix I. of [RFC1112]. See Section 8 and Section 9 for a detailled list of changes from that memo.

IP multicasting is the transmission of an IP datagram to a "host group", a set of zero or more hosts identified by a single IP destination address. A multicast datagram is delivered to all members of its destination host group with the same "best-efforts" reliability as regular unicast IP datagrams, i.e., the datagram is not guaranteed to arrive intact at all members of the destination group or in the same order relative to other datagrams.

The membership of a host group is dynamic; that is, hosts may join and leave groups at any time. There is no restriction on the location or number of members in a host group. A host may be a member of more than one group at a time. A host need not be a member of a group to send datagrams to it.

A host group may be permanent or transient. A permanent group has a well-known, administratively assigned IP address. It is the address, not the membership of the group, that is permanent; at any time a permanent group may have any number of members, even zero. Those IP multicast addresses that are not reserved for permanent groups are available for dynamic assignment to transient groups which exist only as long as they have members.

Internetwork forwarding of IP multicast datagrams is handled by "multicast routers" which may be co-resident with, or separate from, internet gateways. A host transmits an IP multicast datagram as a local network multicast which reaches all immediately-neighboring members of the destination host group. If the datagram has an IP time-to-live greater than 1, the multicast router(s) attached to the local network take responsibility for forwarding it towards all other networks that have members of the destination group. On those other member networks that are reachable within the IP time-to-live, an attached multicast router completes delivery by transmitting the datagram as a local multicast.

This memo specifies the extensions required of a host IP implementation to support IP multicasting, where a "host" is any internet host or gateway other than those acting as multicast routers. The algorithms and protocols used within and between multicast routers are transparent to hosts and will be specified in

separate documents. This memo also does not specify how local network multicasting is accomplished for all types of network, although it does specify the required service interface to an arbitrary local network and gives an Ethernet specification as an example. Specifications for other types of network will be the subject of future memos.

3. LEVELS OF CONFORMANCE

There are three levels of conformance to this specification:

3.1. Level 0: no support for IP multicasting.

There is, at this time, no requirement that all IP implementations support IP multicasting. Level 0 hosts will, in general, be unaffected by multicast activity. The only exception arises on some types of local network, where the presence of level 1 or 2 hosts may cause misdelivery of multicast IP datagrams to level 0 hosts. Such datagrams can easily be identified by the presence of a class D IP address in their destination address field; they SHOULD be quietly discarded by hosts that do not support IP multicasting. Class D addresses are described in section 4 of this memo.

3.2. Level 1: support for sending but not receiving multicast IP datagrams.

Level 1 allows a host to partake of some multicast-based services, such as resource location or status reporting, but it does not allow a host to join any host groups. An IP implementation may be upgraded from level 0 to level 1 very easily and with little new code. Only sections 4, 5, and 6 of this memo are applicable to level 1 implementations.

3.3. Level 2: full support for IP multicasting.

Level 2 allows a host to join and leave host groups, as well as send IP datagrams to host groups. Most IPv6 hosts require Level 2 support because IPv6 Neighbor Discovery ([RFC4861], as used on most link types) depends on multicast and requires that nodes join Solicited Node multicast addresses.

Level 2 requires implementation of the Internet Group Management Protocol (IGMP) for IP and the equivalent Multicast Listener Discovery Protocol (MLD) for IPv6 and extension of the IP and local network service interfaces within the host.

The current protocol versions are [IGMPv3] and [MLDv2] or lightweight versions of either protocol [IGMPv3LITE].

All of the following sections of this memo are applicable to level 2 implementations.

4. HOST GROUP ADDRESSES

IP Host groups are identified by class D IP addresses, i.e., those with "1110" as their high-order four bits. Class E IP addresses, i.e., those with "1111" as their high-order four bits, are reserved for future addressing modes.

In Internet standard "dotted decimal" notation, host group addresses range from 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is guaranteed not to be assigned to any group, and 224.0.0.1 is assigned to the permanent group of all IP hosts (including gateways). This is used to address all IP multicast hosts on the directly connected network. There is no multicast address (or any other IP address) for all hosts on the total Internet. The addresses of other well-known, permanent groups are to be published in "Assigned Numbers".

IPv6 Host groups are identified by IPv6 addresses as defined in [RFC4291] section 2.7 and updated by [RFC7346], [RFC7371].

IP and IPv6 addresses as specified in [SSM] are not used for ASM IP multicast and are not considered IP host groups. They are instead only the destination address part G of Source Specific Multicast (SSM) IP multicast (S,G) channels.

Appendix I contains some background discussion of several issues related to host group addresses.

5. MODEL OF A HOST IP IMPLEMENTATION

The multicast extensions to a host IP implementation are specified in terms of the layered model illustrated below in Figure 1. In this model, ICMP/ICMPv6 and (for level 2 hosts) IGMP/MLD are considered to be implemented within the IP module, and the mapping of IP addresses to local network addresses is considered to be the responsibility of local network modules. This model is for expository purposes only, and should not be construed as constraining an actual implementation.

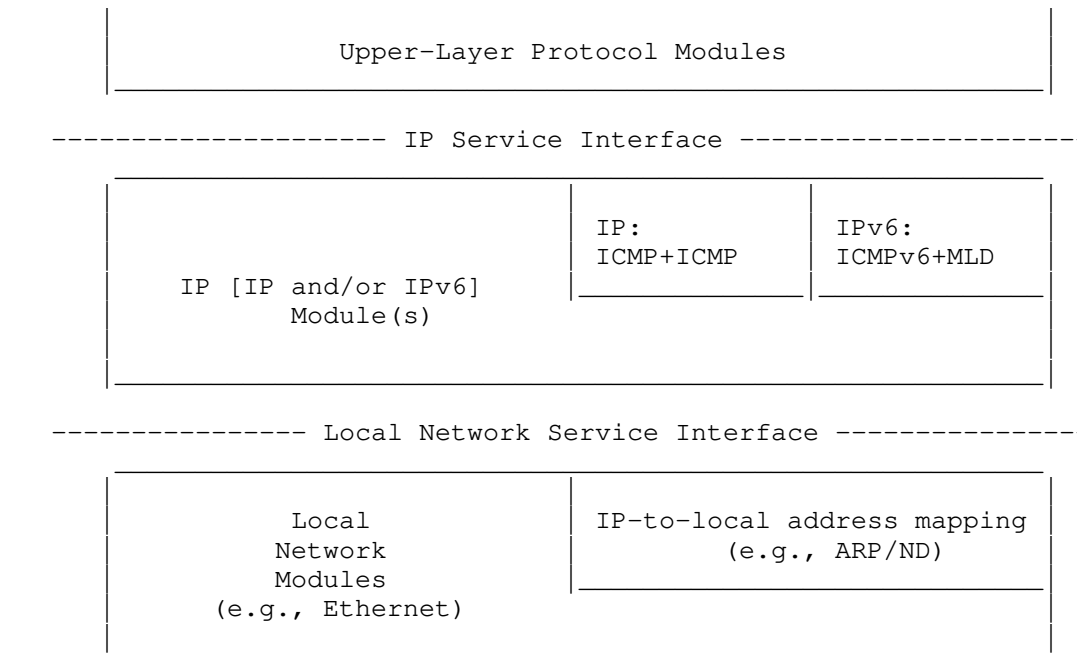

Figure 1: multicast extensions to a host IP implementation

To provide level 1 multicasting, a host IP implementation MUST support the transmission of multicast IP datagrams. To provide level 2 multicasting, a host MUST also support the reception of multicast IP datagrams. Each of these two new services is described in a separate section, below. For each service, extensions are specified for the IP service interface, the IP module, the local network service interface, and an Ethernet local network module. Extensions to local network modules other than Ethernet are mentioned briefly, but are not specified in detail.

6. SENDING MULTICAST IP DATAGRAMS

6.1. Extensions to the IP Service Interface

Multicast IP datagrams are sent using the same "Send IP" operation used to send unicast IP datagrams; an upper-layer protocol module merely specifies an IP host group address, rather than an individual IP address, as the destination. However, a number of extensions may be necessary or desirable.

First, the service interface SHOULD provide a way for the upper-layer protocol to specify the IP time-to-live of an outgoing multicast datagram, if such a capability does not already exist. If the upper-

layer protocol chooses not to specify a time-to-live, it SHOULD default to 1 for all multicast IP datagrams, so that an explicit choice is required to multicast beyond a single network.

Second, for hosts that may be attached to more than one network, the service interface SHOULD provide a way for the upper-layer protocol to identify which network interface is be used for the multicast transmission. Only one interface is used for the initial transmission; multicast routers are responsible for forwarding to any other networks, if necessary. If the upper-layer protocol chooses not to identify an outgoing interface, a default interface SHOULD be used, preferably under the control of system management.

Third (level 2 implementations only), for the case in which the host is itself a member of a group to which a datagram is being sent, the service interface SHOULD provide a way for the upper-layer protocol to inhibit local delivery of the datagram; by default, a copy of the datagram is looped back. This is a performance optimization for upper-layer protocols that restrict the membership of a group to one process per host (such as a routing protocol), or that handle loopback of group communication at a higher layer (such as a multicast transport protocol).

IPv6 socket extensions supporting these functions are defined in [RFC3493], section 5.2.

6.2. Extensions to the IP Module

To support the sending of multicast IP datagrams, the IP module MUST be extended to recognize IP host group addresses when routing outgoing datagrams. Most IP implementations include the following logic:

```
if IP-destination is on the same local network,
   send datagram locally to IP-destination
else
   send datagram locally to GatewayTo( IP-destination )
```

To allow multicast transmissions, the routing logic MUST be changed to:

```
if IP-destination is on the same local network
or IP-destination is a host group,
   send datagram locally to IP-destination
else
   send datagram locally to GatewayTo( IP-destination )
```

If the sending host is itself a member of the destination group on the outgoing interface, a copy of the outgoing datagram MUST be looped-back for local delivery, unless inhibited by the sender. (Level 2 implementations only.)

The IP source address of the outgoing datagram MUST be one of the individual addresses corresponding to the outgoing interface.

A host group address MUST never be placed in the source address field or anywhere in a source route or record route option of an outgoing IP datagram. These packets are not IP Multicast packets but simply invalid packets.

6.3. Extensions to the Local Network Service Interface

No change to the local network service interface is required to support the sending of multicast IP datagrams. The IP module merely specifies an IP host group destination, rather than an individual IP destination, when it invokes the existing "Send Local" operation.

6.4. Extensions to an Ethernet Local Network Module

The Ethernet directly supports the sending of local multicast packets by allowing multicast addresses in the destination field of Ethernet packets. All that is needed to support the sending of multicast IP datagrams is a procedure for mapping IP host group addresses to Ethernet multicast addresses.

An IP host group address is mapped to an Ethernet multicast address by placing the low-order 23-bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01-00-5E-00-00-00 (hex). Because there are 28 significant bits in an IP host group address, more than one host group address may map to the same Ethernet multicast address.

Mapping of IPv6 host group addresses to Ethernet is defined in [RFC2464] and [RFC6085].

6.5. Extensions to Local Network Modules other than Ethernet

Other networks that directly support multicasting, such as rings or buses conforming to the IEEE 802.2 standard, may be handled the same way as Ethernet for the purpose of sending multicast IP datagrams. For a network that supports broadcast but not multicast, such as the Experimental Ethernet, all IP host group addresses may be mapped to a single local broadcast address (at the cost of increased overhead on all local hosts). For a point-to-point link joining two hosts (or a host and a multicast router), multicasts SHOULD be transmitted

exactly like unicasts. For a store-and-forward network like the ARPANET or a public X.25 network, all IP host group addresses might be mapped to the well-known local address of an IP multicast router; a router on such a network would take responsibility for completing multicast delivery within the network as well as among networks.

7. RECEIVING MULTICAST IP DATAGRAMS

7.1. Extensions to the IP Service Interface

Incoming multicast IP datagrams are received by upper-layer protocol modules using the same "Receive IP" operation as normal, unicast datagrams. Selection of a destination upper-layer protocol is based on the protocol field in the IP header, regardless of the destination IP address. However, before any datagrams destined to a particular group can be received, an upper-layer protocol must ask the IP module to join that group. Thus, the IP service interface MUST be extended to provide two new operations:

JoinHostGroup (group-address, interface)

LeaveHostGroup (group-address, interface)

The JoinHostGroup operation requests that this host become a member of the host group identified by "group-address" on the given network interface. The LeaveGroup operation requests that this host give up its membership in the host group identified by "group-address" on the given network interface. The interface argument may be omitted on hosts that support only one interface. For hosts that may be attached to more than one network, the upper-layer protocol may choose to leave the interface unspecified, in which case the request will apply to the default interface for sending multicast datagrams (see section 6.1).

It is permissible to join the same group on more than one interface, in which case duplicate multicast datagrams may be received. It is also permissible for more than one upper-layer protocol to request membership in the same group.

Both operations SHOULD return immediately (i.e., they are non-blocking operations), indicating success or failure. Either operation may fail due to an invalid group address or interface identifier. JoinHostGroup may fail due to lack of local resources. LeaveHostGroup may fail because the host does not belong to the given group on the given interface. LeaveHostGroup may succeed, but the membership persist, if more than one upper-layer protocol has requested membership in the same group.

IPv6 socket extensions supporting these functions are defined in [RFC3493], section 5.2. [RFC3678] specifies these functions for IP and IPv6 (as well as for SSM). Note that these are UDP socket extions (and not IP/IPv6 socket extensions due to the absence of widely available/used IP/IPv6 level socket APIs).

7.2. Extensions to the IP Module

To support the reception of multicast IP datagrams, the IP module MUST be extended to maintain a list of host group memberships associated with each network interface. An incoming datagram destined to one of those groups is processed exactly the same way as datagrams destined to one of the host's individual addresses.

Incoming datagrams destined to groups to which the host does not belong are discarded without generating any error report or log entry. On hosts with more than one network interface, if a datagram arrives via one interface, destined for a group to which the host belongs only on a different interface, the datagram is quietly discarded. (These cases should occur only as a result of inadequate multicast address filtering in a local network module.)

An incoming datagram is not rejected for having an IP time-to-live of 1 (i.e., the time-to-live should not automatically be decremented on arriving datagrams that are not being forwarded). An incoming datagram with an IP host group address in its source address field is quietly discarded. An ICMP/ICMPv6 error message (Destination Unreachable, Time Exceeded, Parameter Problem, Source Quench, or Redirect) is never generated in response to a datagram destined to an IP host group.

The list of host group memberships is updated in response to JoinHostGroup and LeaveHostGroup requests from upper-layer protocols. Each membership should have an associated reference count or similar mechanism to handle multiple requests to join and leave the same group. On the first request to join and the last request to leave a group on a given interface, the local network module for that interface is notified, so that it may update its multicast reception filter (see section 7.3).

The IP module MUST also be extended to implement the IGMP protocol for IP and/or the MLD protocol for IPv6 (depending on the version of IP to be supported). IGMP/MLD are used to keep neighboring multicast routers informed of the host group memberships present on a particular local network.

7.3. Extensions to the Local Network Service Interface

Incoming local network multicast packets are delivered to the IP module using the same "Receive Local" operation as local network unicast packets. To allow the IP module to tell the local network module which multicast packets to accept, the local network service interface is extended to provide two new operations:

```
JoinLocalGroup  ( group-address )

LeaveLocalGroup ( group-address )
```

where "group-address" is an IP host group address. The JoinLocalGroup operation requests the local network module to accept and deliver up subsequently arriving packets destined to the given IP host group address. The LeaveLocalGroup operation requests the local network module to stop delivering up packets destined to the given IP host group address. The local network module is expected to map the IP host group addresses to local network addresses as required to update its multicast reception filter. Any local network module is free to ignore LeaveLocalGroup requests, and may deliver up packets destined to more addresses than just those specified in JoinLocalGroup requests, if it is unable to filter incoming packets adequately.

The local network module MUST NOT deliver up any multicast packets that were transmitted from that module; loopback of multicasts is handled at the IP layer or higher.

7.4. Extensions to an Ethernet Local Network Module

To support the reception of multicast IP datagrams, an Ethernet module MUST be able to receive packets addressed to the Ethernet multicast addresses that correspond to the host's IP host group addresses. It is highly desirable to take advantage of any address filtering capabilities that the Ethernet hardware interface may have, so that the host receives only those packets that are destined to it.

Unfortunately, many current Ethernet interfaces have a small limit on the number of addresses that the hardware can be configured to recognize. Nevertheless, an implementation MUST be capable of listening on an arbitrary number of Ethernet multicast addresses, which may mean "opening up" the address filter to accept all multicast packets during those periods when the number of addresses exceeds the limit of the filter.

For interfaces with inadequate hardware address filtering, it may be desirable (for performance reasons) to perform Ethernet address filtering within the software of the Ethernet module. This is not mandatory, however, because the IP module performs its own filtering based on IP destination addresses.

7.5. Extensions to Local Network Modules other than Ethernet

Other multicast networks, such as IEEE 802.2 networks, can be handled the same way as Ethernet for the purpose of receiving multicast IP datagrams. For pure broadcast networks, such as the Experimental Ethernet, all incoming broadcast packets can be accepted and passed to the IP module for IP-level filtering. On point-to-point or store-and-forward networks, multicast IP datagrams will arrive as local network unicasts, so no change to the local network module should be necessary.

8. Normative changes

8.1. Moving RFC1112 and IGMPv1 to historic status

This document moves [RFC1112] to historic status so that the IGMP version 1 protocol as specified in Appendix 1 of [RFC1112] is moved to historic status. This protocol is not included in the text of this document anymore, which hence renders IGMPv1 historic.

All other aspects of [RFC1112] beside IGMPv1 are inherited and enhanced by this document and maintain their current Internet Standard designation from [RFC1112] through the normative status of this document.

8.2. Backward compatibility with IGMPv1

Newer versions of IGMP or other protocols/mechisms including but not necessary limited to [IGMPv2], [IGMPv3] or [IGMPv3LITE] do or may (such as in [IGMPsnooping]) include backward compatibility with IGMPv1, which requires the [RFC1112] specification of IGMPv1.

This document does not ask for any change to any specifications or implementations that includes any form of support for IGMPv1 for backward compatibility reasons as long as it also includes compatibility with a newer version of IGMP starting with [IGMPv2]. Any new or updated specification with such backward compatibility needs to continue to reference the specification of IGMPv1 in [RFC1112]. Any future reference for new or updated work to any other definition from [RFC1112] needs to refer to this document instead.

9. Changes from RFC1112

Beyond the normative changes described in Section 8, this document introduces the following changes over [RFC1112].

9.1. Normative language

This document introduces the use of normative language through capitalization. [RFC1112] preceeded this method and hence did not have this.

TBD: This version is an initial run across the text to find the appropriate places. It may be incomplete.

9.2. Superceeding references to IGMPv1

References to IGMPv1 in [RFC1112] are replaced by references to [IGMPv3] in this text.

9.3. Introduction of the term Any-Source Multicast (ASM)

This update introduces the term "ASM IP multicast" (ASM) as another term for "Host Extensions for IP multicast". This term was introduced when [SSM] introduced another service model for IP Multicast called "Source Specific Multicast" (SSM), and hence, the service described in [RFC1112] and this update is more precisely called Any Source Multicast (ASM) IP multicast.

[RFC1112] defines and uses the term "host group". This term is not applicable to IP/IPv6 multicast group addresses that are not used for ASM but SSM according to [SSM]. New text in this document explains this.

No functional changes to the IP Multicast service are incurred by these changes.

9.4. Applicability to both IP and IPv6

This update is written to apply to both IP and IPv6 by adding equivalent detail for IPv6 where [RFC1112] only covered IP: addressing and protocols in support of the service – Multicast Listener Discovery [MLDv2] for IPv6 versus IGMP for IP.

Note: IPv6 documents such as [RFC1883] and all its updates (e.g.: [RFC8200]) are defining multicasting in the assumption of the service of [RFC1112] for IPv6, but without being able to refer to [RFC1112], as it was only defined for IP. Future documents can refer to this document as the IP Multicast / ASM service for both IP and IPv6.

Additional text provides references for IETF UDP socket API specifications that instantiate the abstract APIs defined in this document.

No functional changes to the IP Multicast service are incurred by these text changes.

9.5. Standard for IP multicasting in controlled networks

This document removes the claim in the abstract of [RFC1112], that these host extensions are "... the recommended standard for IP multicasting in the Internet."

The reason for this is that [RFC8815] deprecated the ASM Service across the Internet because there is no Internet Standard solution for protocols to support interdomain ASM except for [RFC3956], which is only applicable to IPv6, and even that solution does not resolve the challenges to source access control in interdomain deployments.

In result, ASM is today "only" a recommended solution for controlled networks including controlled federated networks for applications for which SSM is not usable.

However, these limitations to the applicability of ASM to no impact the applicability of most of the host stack described in this document for other forms of IP Multicast, specifically "Source Specific Multicast", [SSM], which inherits all aspects of ASM specified in this document, especially the sending (Section 6, Section 6.2) of IP Multicast packets as well as the mapping to ethernet (Section 6.4). It only amends the joining of IP Multicast traffic on IP Multicast receivers with additional procedures fitting into the host stack described in this document.

10. References

10.1. Normative References

[IGMPv2] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, DOI 10.17487/RFC2236, November 1997, <https://www.rfc-editor.org/rfc/rfc2236>.

[IGMPv3] Haberman, B. and J. Martin, "Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction", RFC 5186, DOI 10.17487/RFC5186, May 2008, <https://www.rfc-editor.org/rfc/rfc5186>.

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <https://www.rfc-editor.org/rfc/rfc1112>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <https://www.rfc-editor.org/rfc/rfc2464>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/rfc/rfc4291>.

[RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/rfc/rfc791>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/rfc/rfc8200>.

[SSM] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, DOI 10.17487/RFC4607, August 2006, <https://www.rfc-editor.org/rfc/rfc4607>.

10.2. Informative References

[I-D.ietf-taps-interface]
Trammell, B., Welzl, M., Enghardt, R., Fairhurst, G., Kühlewind, M., Perkins, C., Tiesel, P. S., and T. Pauly, "An Abstract Application Layer Interface to Transport Services", Work in Progress, Internet-Draft, draft-ietf-taps-interface-22, 6 July 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-taps-interface-22>.

[IGMPsnooping]
Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol

(IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, DOI 10.17487/RFC4541, May 2006, <https://www.rfc-editor.org/rfc/rfc4541>.

[IGMPv3LITE] Liu, H., Cao, W., and H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790, DOI 10.17487/RFC5790, February 2010, <https://www.rfc-editor.org/rfc/rfc5790>.

[MLDv2] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <https://www.rfc-editor.org/rfc/rfc3810>.

[RFC1045] Cheriton, D., "VMTP: Versatile Message Transaction Protocol: Protocol specification", RFC 1045, DOI 10.17487/RFC1045, February 1988, <https://www.rfc-editor.org/rfc/rfc1045>.

[RFC1883] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, DOI 10.17487/RFC1883, December 1995, <https://www.rfc-editor.org/rfc/rfc1883>.

[RFC1972] Crawford, M., "A Method for the Transmission of IPv6 Packets over Ethernet Networks", RFC 1972, DOI 10.17487/RFC1972, August 1996, <https://www.rfc-editor.org/rfc/rfc1972>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <https://www.rfc-editor.org/rfc/rfc2460>.

[RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", RFC 3493, DOI 10.17487/RFC3493, February 2003, <https://www.rfc-editor.org/rfc/rfc3493>.

[RFC3678] Thaler, D., Fenner, B., and B. Quinn, "Socket Interface Extensions for Multicast Source Filters", RFC 3678, DOI 10.17487/RFC3678, January 2004, <https://www.rfc-editor.org/rfc/rfc3678>.

[RFC3956] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3956, DOI 10.17487/RFC3956, November 2004, <https://www.rfc-editor.org/rfc/rfc3956>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/rfc/rfc4861>.

[RFC6085] Gundavelli, S., Townsley, M., Troan, O., and W. Dec, "Address Mapping of IPv6 Multicast Packets on Ethernet", RFC 6085, DOI 10.17487/RFC6085, January 2011, <https://www.rfc-editor.org/rfc/rfc6085>.

[RFC7346] Droms, R., "IPv6 Multicast Address Scopes", RFC 7346, DOI 10.17487/RFC7346, August 2014, <https://www.rfc-editor.org/rfc/rfc7346>.

[RFC7371] Boucadair, M. and S. Venaas, "Updates to the IPv6 Multicast Addressing Architecture", RFC 7371, DOI 10.17487/RFC7371, September 2014, <https://www.rfc-editor.org/rfc/rfc7371>.

[RFC8507] Deering, S. and R. Hinden, Ed., "Simple Internet Protocol (SIP) Specification", RFC 8507, DOI 10.17487/RFC8507, December 2018, <https://www.rfc-editor.org/rfc/rfc8507>.

[RFC8815] Abrahamsson, M., Chown, T., Giuliano, L., and T. Eckert, "Deprecating Any-Source Multicast (ASM) for Interdomain Multicast", BCP 229, RFC 8815, DOI 10.17487/RFC8815, August 2020, <https://www.rfc-editor.org/rfc/rfc8815>.

Appendix A. HOST GROUP ADDRESS ISSUES

This appendix is not part of the IP multicasting specification, but provides background discussion of several issues related to IP host group addresses.

A.1. Group Address Binding

The binding of IP host group addresses to physical hosts may be considered a generalization of the binding of IP unicast addresses. An IP unicast address is statically bound to a single local network interface on a single IP network. An IP host group address is dynamically bound to a set of local network interfaces on a set of IP networks.

It is important to understand that an IP host group address is NOT bound to a set of IP unicast addresses. The multicast routers do not need to maintain a list of individual members of each host group. For example, a multicast router attached to an Ethernet need associate only a single Ethernet multicast address with each host group having local members, rather than a list of the members' individual IP or Ethernet addresses.

A.2. Allocation of Transient Host Group Addresses

This memo does not specify how transient group address are allocated. It is anticipated that different portions of the IP transient host group address space will be allocated using different techniques. For example, there may be a number of servers that can be contacted to acquire a new transient group address. Some higher-level protocols (such as VMTP, specified in [RFC1045]) may generate higher-level transient "process group" or "entity group" addresses which are then algorithmically mapped to a subset of the IP transient host group addresses, similarly to the way that IP host group addresses are mapped to Ethernet multicast addresses. A portion of the IP group address space may be set aside for random allocation by applications that can tolerate occasional collisions with other multicast users, perhaps generating new addresses until a suitably "quiet" one is found.

In general, a host cannot assume that datagrams sent to any host group address will reach only the intended hosts, or that datagrams received as a member of a transient host group are intended for the recipient. Misdelivery must be detected at a level above IP, using higher-level identifiers or authentication tokens. Information transmitted to a host group address should be encrypted or governed by administrative routing controls if the sender is concerned about unwanted listeners.

Appendix B. Discussion and Explanations (TO BE REMOVED)

[RFC-editor: Please remove this section]

Please refer to Section 9 for the non-process disucssion of the goals of this document.

B.1. Goals of this document

The goal of this document is to allow for IETF to declare [RFC1112] historic and inherit the full INTERNET STANDARD status of [RFC1112] with this document immediately - without going through the otherwise necessary long process.

The reason why [RFC1112] needs to be declared historic is so that the IGMP version 1 protocol specified in it can be declared obsolete. This update removes IGMPv1 text.

The reason why this document is still needed (as an Internet Standard), is because the IP Multicast service specified in [RFC1112] has since its inception been the Internet Standard for the IP Multicasting service.

To allow for this document to gets immediately the intended Internet Standard status, it introduces no functional changes and it deliverately avoids also any unnecessary textual changes. This includes the deliberade non-upgrade of the [RFC1112] language to use [RFC2119] terminology. While the use of that language might be preferred for new work/text, the success of IP Multicasting as defined in [RFC2119] seems to indicate that the existing text was more than sufficient.

B.2. Internet Standard status

Note that the removal of the IGMPv1 protocol may raise the question whether the document in its current form still contains specifications sufficient for Internet Standard as opposed to Informational.

The core aspects that impacts interoperability (and hence qualifies the document for Internet Standard) is the format of IP packets when IP Multicast service is used, e.g.: IP Multicast addressing and binding to Multicast Ethernet MAC addresses. There is no other RFC that introduces these specifications for IP, because there was never another update to [RFC791] to do this. [SSM], another standards track document bilding on [RFC1112], defining the SSM service / host stack.

This update also includes the necessary text for IPv6. Note that for IPv6 the ethernet MAC address mapping of IPv6 multicast packets was later (after [RFC1112]) specified in [RFC1972] and its updates, but scattering the aspects of IPv6 multicast across (currently) [RFC2464], {RFC4291}} and [RFC8200] makes it arguably more difficult for implementers to understand the technology than this document that coalesces all these services aspects - from ethernet bonding to application interfac.

Beyond those packet format/ethernet aspects, historically, the Multicast service (API) related text in [RFC1112] would not have been considered to be an Internet Standards scope definition because this classification was not extended to (abstract) APIs, even though they do of course define an interoperability interface between e.g.: operating system providing the API and applications using it.

Recently, the IETF has changed its stance on this issue though and is working on [I-D.ietf-taps-interface] with the intent for it to become Internet Standard. With this in mind, all that text of [RFC1112] can also be considered appropriate for Internet Standard.

B.3. Authors email

This document does include Steve Deering as the original author of the base rfc [RFC1112] in the same way as [RFC8200] does for [RFC2460] ([RFC8507] is also similar). The mayority of text originates from his original RFC, hence he is the primary co-author. He is not actively involved in editing this -bis document but is in support of the work. Changes are edited by the co-author(s).

Being retired and not actively involved, he does not want for his personal email address to be included in the work. Any notifications where all authors need to provide feedback will be forwarded through the co-author.

Whereas at the time of [RFC8200] it was possible for a co-author to not have to include an email address, policies on Datatracker have since changed and therefore this draft include a placeholder email address for Steve solely to allow uploading to Datatracker.

B.4. Changelog

B.5. Open Issues

Need to revisit in more details the logic of upding RFC791. RFC1112 did not claim to be such an update even though it does effectively update RFC791 because it exempts IP Multicast packets from RFC791 processing. And also introduces invalid packets (source address IP Multicast which are neither unicast nor multicast).

Likewise the same would apply to RFC8200 which does not specify these details either.

B.5.1. draft-eckert-pim-rfc1112bis-02

Changed core references from numbered style to name style .

Changed copyright clause to pre5378Trust200902, which is the same as used for RFC8200 due to the presence of text with similar early status.

To resolve Dinos concerns at IETF116 with -01: Added hopefully extensive explanation wrt. to how to treat IGMPv1 based on Dino's feedback from IETF117: This document does not ask for any removal of IGMPv1 in any IETF specs which include it for backward compatibility reasons, it only effectively causes it to become historic once RFC1112 would be declared historic.

To resolve Alvaros concerns at IETF1116 with -01: Added normative language (MUST/SHOULD). Seems as if this is quite easy given how "must" was written appropriately in the original text. The logic of applying MUST/MUST-NOT was based on understanding by the author how none of the MUST would actually put existing working implementations out of compliance.

Added explicit text to move rfc1112 to historic status.

Moved explanation of changes from rfc1112 from appendix to main text as this seem to the common practice for document updates.

Added claim for this document to be an update to rfc791. See open issues section though.

B.5.2. draft-eckert-pim-rfc1112bis-01

Changed all use of IPv4 back to IP. Seems standard in IETF specs. Only IPv6 has in IETF specs the distinction of including the version.

Changed Steve Deerings address to a pseudo-email address at IETF. See prior section.

Converted document into kramdownrfc2629 format for easier editing.

Claims that rfc2119 language is not desired/used (to maintain maximum original text without changes).

Rewrote section for updates to rfc1112 to hopefully better motivate/ explain the reason for this document and detail what its changes are.

B.5.3. draft-eckert-pim-rfc1112bis-00

Initial version based on [RFC1112] text version, edited.

Authors' Addresses

Stephen E. Deering
Retired
Vancouver, British Columbia
Canada
Email: deering@noreply.ietf.org

Toerless Eckert (editor)
Futurewei Technologies USA
United States of America
Email: tte@cs.fau.de