PIM
Internet-Draft
Obsoletes: 1112 (if approved)
Updates: 791, 1122 (if approved)
Intended status: Standards Track
Expires: 31 August 2026

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

27 February 2026

Host Extensions for IP Multicasting and "Any Source Multicasting" (ASM) IP service

draft-ietf-pim-rfc1112bis-08

Abstract

This memo specifies the extensions required of a host implementation of the Internet Protocol (IP) to support IP multicast with the IP service interface "Any Source Multicast" (ASM). This specification applies to both versions 4 and 6 of the Internet Protocol. Distribution of this memo is unlimited.

This document replaces RFC1112 for everything 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 31 August 2026.

Copyright Notice

Table of Contents

1. STATUS OF THIS MEMO

[To be removed before publication: Summary of considerations for reviews by different groups:

This -bis is intended to replace RFC1112 maintaining it internet standard designation, but extending it for IPv6, additional terminology (ASM/SSM), and refining the specification with established industry practices.

The core parts of the document are changed as little as possible to maintain all original rfc1112 text (except IGMPv1) as much as possible - given how it has very well stood the test of time: all well-known IP multicast host stack implementations including IPv6 - even though unspecified there - are based on the principles of rfc1112. New sections and existing, minimally changed sections can easily be recognized by using rfcdiff against RFC1112.

All changes/enhancements are meticulously matched against implementation and operational practices that have evolved and are detailled in this memo: this -bis should match the ubiquitously deployed IP multicast service better than rfc1112.

SECDIR is asked primarily to review section 12 (Security Considerations).

INTDIR: This document would logically belong to INT as it extends the IPv4/IPv6 host stack for IP Multicast (and references to SSM). It simply evolved as a PIM document due to PIM-WG ongoing ownership of all of IP multicast below application layer. IPv6 is added mostly "by-reference", because in the absence of an earlier attempt to add IPv6 support into an rfc1112bis, all normatively necessary aspects of IPv6 multicast where added to a scattered set of RFCs, which are now comprehensively referenced in this memo.

TSVDIR: Consider this document to be normative for all "UDP" independent service and abstract API aspects of datagram IP multicast service. Its hence related to the work by TAPS. The Security Considerations sections specifically discusses challenges of adopting the socket model from unicast to multicast.

IOTDIR: IP multicast is widely in IOT, often without IP multicast routing just locally in LANs, radio-LANs. This memo should be the best common reference for the quirks of IP multcast host stacks, specifically with the added discussion of link-local addresses and socket (security) challenges.

]

This memo specifies the extensions required of a host implementation of the Internet Protocol (IP) to support IP multicast with the IP service interface "Any Source Multicast" (ASM). This specification applies to both versions 4 and 6 of the Internet Protocol. Distribution of this memo is unlimited.

This document replaces RFC1112 for everything except for 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

2.1. Summary

This memo specifies the extensions required of a host implementation of the Internet Protocol (IP) to support IP multicast. It replaces [RFC791] for everything except for the specification of the protocol IGMP version 1 in Appendix I. of RFC1112. This document declares RFC1112 including IGMP version 1 historic.

RFC1112 specified IP multicast for version 4 of the IP protocol (IPv4, [RFC791]), and refers to that version as IP. This document applies both to version 4 of the IP protocol and version 6 of the IP protocol (IPv6, [RFC8200]).

THE TERM IP IS USED IN THIS DOCUMENT FOR TEXT APPLYING EQUALLY TO IPv4 AND IPv6.

Where specifications in support of IP multicast for version 6 of the IP protocol where already provided by other RFCs, this document provides references to those pre-existing specifications, so that this document can serve as a complete single point of reference for the host extensions for IP multicast with either versions of IP.

"Source Specific Multicast", (SSM, [SSM]) introduced a complementary extension to the IP service from the one specified here. It relies on all aspects of the host stack extensions specified here, such as Section 6.4, and uses or extends them. The service specified here is called "Any Source Multicast" (ASM) to distinguish it explicitly from SSM. This document also describes, where SSM changes specifications from RFC1112.

Due to the existence of both ASM and SSM, the term "IP multicast" best refers to the complete set of IP host extensions in support of either service options: this specification for ASM plus [SSM]). When the term IP multicast is used to refer to the IP multicast service without further qualification, then ASM is to be implied. See also Section 10.14.

This specification aims to maintain all the original text of RFC1112 where technically appropriate. This incurs the use of some historic language, such as "(internet) gateway" to sometimes refer to IP routers, and capitalization of chapter headings.

[RFCeditor: please remove this remark before publication. Reviewers: Please use rfcdiff to easier recognize the sections inherited from RFC1112 and distinguish them from new chapters and sections. The pre-existing text attempts to include only necessary technical enhancements but not other editorial enhancements.]

See Section 9 and Section 10 for a detailed list of changes from RFC1112.

2.2. Overview

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 IPv4 time-to-live or IPv6 hop limit 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 IPv4 time-to-live or IPv6 hop limit, 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 are 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 four levels of conformance to this specification. They apply independently for IPv4 and IPv6.

All Internet hosts and gateways are RECOMMENDED to conform to Level 2 for the versions of IP that they support.

Hosts or gateways supporting IPv4 that can not conform to Level 2 for it are RECOMMENDED to conform to Level 2L.

Hosts or gateways supporting IPv6 that can not conform to Level 2 for IPv6 are REQUIRED to conform to Level 2L. This option is introduced in support of the requirements from [RFC4291], section 2.8.

See also Appendix A.3 for further explanations of the use of link local addresses.

3.1. Level 0: no support for 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 an IP multicast address in their destination address field; they SHOULD be quietly discarded by hosts that do not support IP multicasting. Class D addresses in support of multicasting with IPv4 are described in Section 4, IPv6 addresses for IP multicasting are described in section 2.7 of [RFC4291] and [RFC7371].

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, see [RFC8504], section 5.4), depends on multicast and requires that nodes join Solicited Node multicast addresses.

Level 2 requires implementation of the host side of the Internet Group Management Protocol (IGMP) for IPv4 and the equivalent host side of the Multicast Listener Discovery Protocol (MLD) for IPv6 and extension of the IP and local network service interfaces within the host as specified or referred to in the following sections.

The current protocol versions for full Level 2 support of IP multicasting are [IGMPv3] and [MLDv2] or lightweight versions of either protocol [RFC5790].

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

3.4. Level 2L: support for only link local IP multicasting.

Level 2L has the same functionality as Level 2 except that it does not include the implementation of IGMP for IPv4 or MLD for IPv6. Level 2L hosts can only send/receive IP multicast to their local network.

Level 2L hosts SHOULD only join/leave Link-Local host groups (see Section 4) and send IP datagrams to Link-Local host groups - but not other host groups.

4. HOST GROUP ADDRESSES

IPv4 Host groups are identified by class D IPv4 addresses, i.e., those with "1110" as their high-order four bits. Class E IPv4 addresses, i.e., those with "1111" as their high-order four bits, are reserved for future addressing modes.

In Internet standard "dotted decimal" notation, IPv4 host group addresses range from 224.0.0.0 to 239.255.255.255. IPv4 host group addresses in the "Local Network Control Block", 224.0.0.0 - 224.0.0.255 are called Link-Local IPv4 host group addresses. IP datagrams with a Link-Local destination address are called Link-Local multicast packets. The IPv4 link local 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 IPv4 hosts (including gateways). It is called the all-nodes group. This is used to address all IP multicast hosts (including gateways) 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 well-known, permanent IPv4 multicast groups are to be published in "Assigned Numbers", see [RFC3232], currently through the IANA "IPv4 Multicast Address Space Registry" [IANA.MASR]. [RFC5771] and [RFC6034] refine more detailed allocation and uses of different sub-blocks of 224.0.0.0/4.

Allocation guidelines for Link-Local IPv6 multicast group addresses are specified in [RFC5771]. The IPv6 Link-Local all-nodes group address is ff02::1. IPv6 Host groups are identified by IPv6 addresses as defined in [RFC4291] section 2.7 and updated by [RFC7346], [RFC7371]. The addresses of other groups are currently published via the IANA "IPv6 Multicast Address Space Registry".

IP addresses as specified in [SSM] are not used for ASM IP multicast and are not considered host groups by [SSM], Terminology section, third paragraph. They are instead only the destination address part G of Source Specific Multicast (SSM) IP multicast (S,G) channels. The term IP multicast address covers both ASM host group addresses and SSM destination addresses.

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.

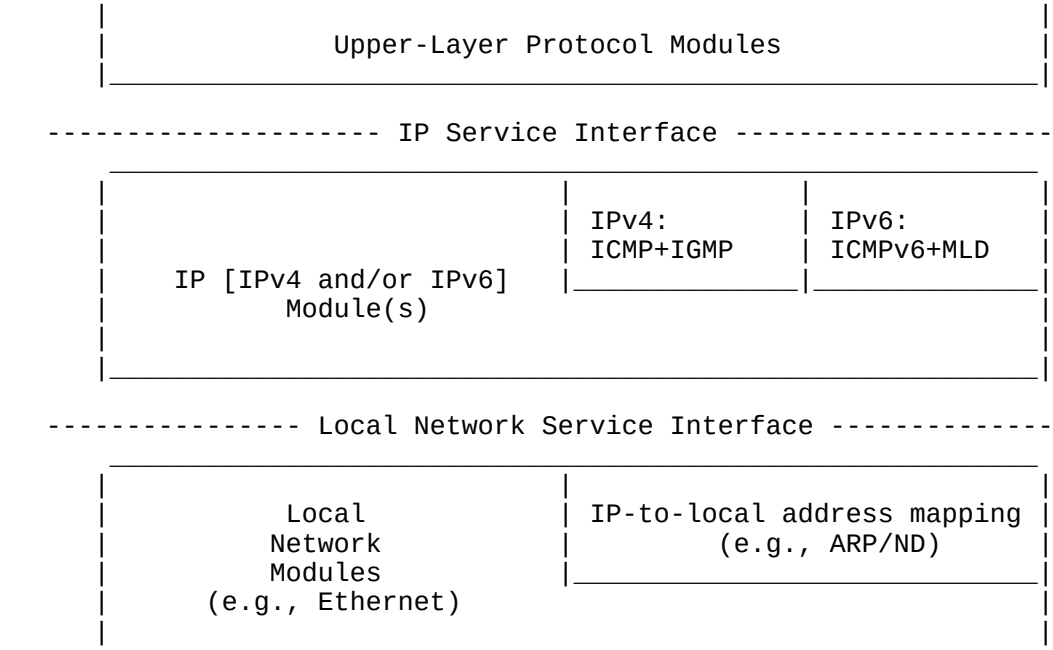

Figure 1: multicast extensions to a host IP implementation

Note that as described in Section 3.3, ND ([RFC4861]) itself operates on top of the IPv6 Service Interface as extended by this document because it relies on sending/receiving IPv6 multicast packets. However, it is shown as part of the Local Network Module because that is the component in this host stack model that relies on ND to perform its operation.

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 IPv4 time-to-live or IPv6 hop limit 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/hop limit, 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 to 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 2/2L 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 2/2L implementations only.)

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

An IP multicast 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 IPv4 host group address is mapped to an Ethernet multicast address by placing the low-order 23-bits of the IPv4 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 IPv4 host group address, more than one host group address may map to the same Ethernet multicast address.

These address mappings for IP addresses do apply not only to host group addresses, but also to IP multicast addresses which are SSM destination addresses.

Mapping of IPv6 multicast addresses (both host group addresses and SSM destination addresses) to Ethernet addresses is defined in [RFC2464] and [RFC6085]. Note that [RFC9542] establishes an "IANA OUI Ethernet Numbers" registry covering the IPv4 and IPv6 multicast MAC address ranges.

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 IPv4 header or the next header field in the IPv6 header or IPv6 extension header preceeding the upper-layer protocol header (when IPv6 extension headers are used). This is 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 socket options for these functions for ASM and also includes socket options in support of SSM. See also Section 12.

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 MUST be 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 IPv4 time-to-live of 1 or IPv6 Hop Limit of 1. This field MUST not automatically be decremented on arriving datagrams that are not being forwarded. An incoming datagram with an IP multicast 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 or SSM range destination IP address.

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).

When supporting Level 2, the IP module MUST also be extended to implement the IGMP protocol for IPv4 and the MLD protocol for IPv6 depending on the version(s) 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.

Level 2 hosts and gateways MAY omit the sending of IGMP messages to report membership for Link-Local IPv4 host group addresses, especially on networks known not to (be able to) use any form of IGMP snooping. This does also apply for the IPv6 Link-Local all-nodes group ff02::1, but not to other Link-Local IPv6 host group addresses. See Section 10.7 and Appendix A.3.

Level 2/2L hosts and gateways SHOULD permanently join to the Link-Local all-nodes group for the version of IP they implement. See Section 10.11.

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 multicast addresses (host group addresses or SSM destination 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. ROUTING MULTICAST IP DATAGRAMS

IP multicast routers are recommended to support the IP host stack extensions as specified in this document especially to support applications using the IP Service Interface Section 5 to send/receive IP multicast packets including those commonly required for IPv6 ([RFC4861]).

Given how IP multicast routers behavior and their behavior for IGMP/MLD differs from non IP multicast routers, Local Network Module layer and IGMP/MLD protocol requirements MAY be optimized/changed from what is required by this document. See Appendix A.4 for more details/examples.

IPv4 datagrams with a Link-Local destination address MUST never be forwarded to a different link by multicast routers, regardless of their time-to-live. See Section 10.10 for explanations.

The equivalent requirement are specified for IPv6 in [RFC4291], section 2.5.6.

Rules for forwarding of non Link-Local IP multicast packets are outside the scope of this document.

9. Status changes

9.1. Moving RFC1112 and IGMPv1 to historic status

This document moves RFC1112 to historic status which also moves the IGMP version 1 protocol as specified in Appendix 1 of RFC1112 to historic status, as it is not included into this document anymore.

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

9.2. Backward compatibility with IGMPv1

Current versions of IGMP ([IGMPv2], [IGMPv3]) and other protocols/mechanisms including, but not limited to [RFC5790] or [IGMPsnooping] do include backward compatibility with IGMPv1. This requires them to refer to RFC1112 as the specification for IGMPv1. Backward compatibility is when a specification also includes support for any newer version of IGMP starting with [IGMPv2] and prefers it over IGMPv1.

This document does not ask for any change to any current or future specifications or implementations that includes any form of support for IGMPv1 for backward compatibility reasons.

Any new or updated specification that wants to maintain such backward compatibility with IGMPv1 need to continue to reference RFC1112 as the specification of IGMPv1.

Any future reference for new or updated work to any other definition from RFC1112 (host extensions for IP multicast and/or Any Source Multicast service) needs to refer to this document instead of RFC1112.

9.3. Update to RFC 791

This document is an update to [RFC791] because none of the core procedures to send and receive IP multicast packets described in this document match those defined for IP unicast packets in [RFC791]. Instead, IP multicast is carving out parts of the IP address space to trigger completely new forwarding for completely new entities: host groups in ASM, channels in SSM). See Appendix B.3 for further discussions.

9.4. Update to RFC 1122

This document updates [RFC1122] section 3.2.3 by making support for Level 2 conformance and hence support for IGMP recommended instead of optional as required by [RFC1122]. See Section 3.

9.5. Update to STD 5

This document replaces RFC1112 in [STD5] which defines IPv4 ([RFC791]) including its core extensions.

Note: As there is no precedent for STD86 (IPv6) to include any specifications for extension of IPv6, this document is not asked to become part of STD86.

10. Changes from RFC1112

Beyond the status changes described in Section 9, this document introduces the following changes over RFC1112.

All requirements changes are intended to make this specification aligned with long-term, most widely implemented, deployed and standardised RFCs for IP multicast, so that this document does not create the need to change existing implementations or deployments, as could be the case if RFC1112 (without IGMPv1) was to be implemented today.

10.1. Normative language

This document introduces the use of normative language through capitalization. RFC1112 preceded [RFC2119] and hence did not include this language.

10.2. References to IGMPv1

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

10.3. New summary

The new Section 2.1 summarizes the scope of this document and the core new changes over RFC1112.

10.4. Any-Source Multicast (ASM)

This update introduces the term "ASM IP multicast" (ASM) as a new term for the IP service interface specified in this document (and previously in RFC1112) as explained in Section 2.1.

10.5. SSM

Section 2.1 explains the relationship of this document to SSM ([SSM]).

Section 4 adds the specification that the term host groups specified in this document does not apply to destination addresses used for SSM. IP multicast address applies to both host group address and SSM destination addresses.

No functional changes to the IP multicast service are incurred by these changes, except that it acknowledges the existence of SSM which reduces the range of host group addresses used for ASM.

10.6. Applicability to both IPv4 and IPv6

This document is written to apply to both IPv4 and IPv6 by adding detail for IPv6 where RFC1112 only covered IPv4. This includes addressing and protocols in support of the service - Multicast Listener Discovery [MLDv2] for IPv6 versus IGMP for IPv4.

IPv6 documents such as [RFC1883] and all its updates (e.g.: [RFC8200]) are defining the necessary wire encoding aspects of IP multicast in the assumption of the service of RFC1112 for IPv6, but without being able to refer to RFC1112, as it was only defined for IPv4. Future documents can refer to this document as the IP multicast / ASM service for both IPv4 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 changes.

10.7. RFC1122 and Level 2L

[RFC1122] did not require support for IPv4 multicasting ("there is at this time no requirement that all IP implementations support IP multicasting"). Instead, [RFC1122] recommends support for IPv4 multicast (according to RFC1112), but support for IGMP to be optional, specifying that sending/receiving IPv4 multicast from/to the local networks works without IGMP and that that is the recommended form to support IPv4 multicasting. See also Appendix A.3.

Whereas [RFC1122] was not even specifying the combination of supporting sending/receiving IPv4 multicast but not supporting IGMP, this document now adds that option by specifying it as conformance Level 2L. Introduction of this text does also not change long-term deployment practices but only formalizes them.

10.8. RFC4291 and Level 2L

According to [RFC4291], IPv6 nodes must support a variety of Link-Local IPv6 multicast address. This translates into the requirement for IPv6 hosts to at least support Level 2L, which is sufficient to support Link-Local IPv6 multicast. Choosing to support only Level 2L is also the only option in which an IPv6 host or gateway will not need to send MLD messages for Link-Local groups because the [MLDv2] specification (unlike IGMP) choose to mandate the sending of MLD messages even for Link-Local host groups. See Appendix A.3 for more

details.

10.9. IP multicast support

With [IGMPv3] now being Internet Standard, there is sufficient experience to also make support for conformance Level 2 of IPv4 multicasting recommended through this document. This is also documented as an update to the IGMP support requirement in [RFC1122] from optional to recommended. See Section 9.4).

Unlike [RFC1122], [RFC8504] does not directly raise a requirement against support for MLD for every node supporting IPv6. Instead, it explains the dependencies against IPv6 multicast and hence MLD for core IPv6 protocols used on most link types (ND, SLAAC).

With [MLDv2] now being Internet standard, and over two decades of experience with IPv6 multicast availability and use on almost all IPv6 implementations, this documents now also recommends support for Level 2 conformance for IPv6 multicast, see Section 3. Note that this is not declared as an update to [RFC8504], because it is outside that BCP documents scope.

10.10. IPv4 Local Network Control Block

RFC1112 defines the requirement for IPv4 datagrams to the all-nodes group 224.0.0.1 to never be forwarded beyond a single network. In later RFCs, this behavior became the BCP for the whole IPv4 Local Network Control Block 224.0.0.0 - 224.0.0.255, making it the Link-Local host group address block for IPv4 multicast. [RFC2365] and [RFC5771], section 4 are the BCPs covering this requirement.

This document formalizes this BCP behavior as a standard requirement in Section 8, superseding and encompassing the more specific requirement for just 224.0.0.1 from RFC1112, and mirroring the same standardized behavior for IPv6 link local addresses. Because this is actually a requirement against IP multicast routers and not hosts, this is now also accordingly described in a separate section.

This requirement does not incur changes over how IP multicast is implemented or deployed.

10.11. Permanent membership for Link-Local all-nodes groups

RFC1112, section 7.2 introduced the requirements for hosts to permanently join 224.0.0.1. Its explains this requirement to be in support of IGMP (version 1).

[IGMPv2], section 6. and [IGMPv3], section 5. inherits this requirement, and [MLDv1], section 6. and [MLDv2] section 6. also define the same requirement for the IPv6 Link-Local all-nodes address ff02::1.

RFC1112 explains this choice by being "(1) it is simpler", and "(3) the all-nodes address may serve other routing-oriented purposes, such as advertising the presence of gateways or resolving local addresses."

Technically, there is no necessity to permanently join the Link-Local all-nodes group. Like any other group, reception of packets could enabled through the JoinHostGroup()/LeaveHostGroup(), as described in Section 7.1. However, all known host implementations that support IP multicast since RFC1112 are based on its definitions and there is no obvious benefit in changing this. Hence this functionality is a should requirement in this document.

Note that one simplification that this requirement enables is to avoid supporting the JoinHostGroup()/LeaveHostGroup() API inside an operating system kernel, but still allow kernel level protocols to receive packets to the Link-Local all-nodes group. This is for example common in support of ICMP/ICMPv6 echo: "ping 224.0.0.1" to discover IP hosts with IP multicast support on the local network. However, this functionality is not enabled by default anymore in modern systems for security reasons (e.g.: linux: net.ipv4.icmp_echo_ignore_broadcasts=1 default configuration).

The requirements text in this spec therefore does not incur any requirements changes for implementations of these existing versions of IGMP/MLD. By making the requirement only a should, it is also clear that future versions of IGMP/MLD or new host stack implementations may change this if they find good reasons to do so - without requiring to update this specification.

Note that [RFC5790] omits this requirement.

10.12. IGMP/MLD messages for Link-Local IP host group addresses

RFC1112, Appendix I. (IGMPv1), [IGMPv2], [IGMPv3], [MLDv1], [MLDv2] require hosts to not send IGMP/MLD messages for the all-nodes group. This would be in conflict with the general rules of RFC1112 (outside of its IGMPv1 specific definitions) and equally this specification if it was not enhanced. This specification therefore contains new text that makes it compatible with existing IGMP/MLD specifications, and with long term established and deployed implementation practices.

New text in Appendix A.3 explains how after RFC1112, it became a common place implementation choice to not send IGMP messages for any IPv4 Link-Local host group address, and explains how this was done with good technical reason at the time. This behavior is so common, that [IGMPsnooping] mandates to explicit support it in IGMP snooping implementations.

Referring to that explanation, a new MAY requirement in Section 7.2 allowing (but not recommending) this behavior makes existing specifications and deployments compatible with this documents specifications. It is only a MAY even though it is common in IPv4, because the experience with IPv6 shows that it does work (of course) equally well if this is not done, and can then support better MLD snooping than IGMP snooping.

10.13. 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 preferable.

However, these limitations to the applicability of ASM do not impact the applicability of any parts of the host stack described in this document for other IP multicast service interfaces, 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.

10.14. Terminology

In RFC1112, all IPv4 multicast addresses where designated to be used with ASM and where thus host group addresses. In result, the term multicast and host group address could be used interchangeably. With the introduction of SSM in [SSM], subsets of IP multicast addresses where carved out for use with SSM instead of ASM. Since then, not every multicast address is a host group address, but every host group address is still a multicast address.

In [SSM], the equivalent to a host group is the SSM channel. It is addressed by the packets (S,G) - the combination of the unicast source and multicast destination address. Multicast addresses used for SSM by themselves are called SSM destination address or SSM multicast address.

Terms like "SSM channel address" or "SSM multicast address" are ambiguous and should be avoided: They could either refer to a specific (S,G) channel, in which case it should be called an SSM channel (S,G) address pair, or it could refer to a multicast destination address G used for SSM, which can be part of many different SSM (Si,G) channels, in which case it should be called an SSM destination address.

Specifications whose behavior does not differ between ASM and SSM can continue to refer to multicast addresses, implying the meaning of multicast to be the superset of ASM and SSM. This is for example true for [RFC4291] and [RFC8200].

New documents should explicitly indicate whether they apply to only ASM and/or SSM even if their behavior applies to both ASM and SSM identically. Else it is left to the reader to guess whether the text does also apply to SSM. If the term multicast is used to indicate behavior that only applies to ASM, this should equally be called out explicitly. Behavior applying only to ASM may use the terms host group address or ASM multicast address.

11. IANA Considerations

11.1. Protocol Numbers registry

IANA is asked to replace the Reference field for the IGMP protocol in the Protocol Numbers registry (https://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml) from RFC1112 to [THIS-RFC].

Explanation: This protocol number is used by all versions of IGMP, including [IGMPv2] and [IGMPv3] and is unaffected by making IGMP version 1 historic.

11.2. Internet Group Management Protocol (IGMP) Type Numbers Registry

IANA is asked to replace the Reference to RFC1112 for the 0x11 / "IGMP Membership Query" entry in the "Internet Group Management Protocol (IGMP) Type Numbers Registry" (https://www.iana.org/assignments/igmp-type-numbers/igmp-type-numbers.xhtml) with "RFC1112, [RFC2236], [RFC3376]".

Explanation: These type code messages where introduced by RFC1112 but modified versions thereof where also introduced by [RFC2236] and [RFC3376], so that it is clearer if all three RFCs are indicated. All other references to RFC1112 in this registry are specifically referring to that RFC in its role of defining IGMP version 1 and thus need to continue to refer to RFC1112 and not [THIS-RFC].

11.3. Multicast 48-bit MAC Addresses registry

IANA is asked to replace the Reference field for the IPv4 Multicast range entry in the "IANA Multicast 48-bit MAC Addresses" (https://www.iana.org/assignments/ethernet-numbers) from RFC1112 to [THIS-RFC].

11.4. IPv4 Address range registries

IANA is asked to replace the Reference field for the 240.0.0.0/4 entry in the "IANA IPv4 Special-Purpose Address Registry" (https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml) from RFC1112 to [THIS-RFC]. The Section 4 text stays unchanged.

IANA is asked to replace the Reference to RFC1112 in the "IANA IPv4 Address Space Registry" (https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml) with [THIS-RFC].

11.5. IPv4 Multicast Address Space registry

IANA is asked to replace the three references to RFC1112 in the "IPv4 Multicast Address Space Registry" (https://www.iana.org/assignments/multicast-addresses/multicast-addresses.xhtml) with [THIS-RFC].

11.6. IP Flow Information Export registry

IANA is asked to replace the two references to RFC1112 in the "IPFIX Information Elements" registry (https://www.iana.org/assignments/ipfix/ipfix.xhtml) with [THIS-RFC].

11.7. IANA OUI Ethernet Numbers

IANA is asked to replace the RFC1112 reference in the IPv4 Multicast entry of the "IANA Multicast 48-bit MAC Addresses" registry table with [THIS-RFC].

12. Security Considerations

This section may repeat a few core observations from elsewhere in the document to make it easier for security interested readers to understand the context without having to understand the whole document.

Application Socket Security Considerations are outside the scope of this document yet important for secure operations of an IP multicast host stack. They are hence covered in Appendix A.5.

12.1. Network forwarding issues

Security issues exists in an internetwork when sending IP multicast packets or when joining IP multicast groups leads to internetwork state. Nevertheless, those issues are not caused by the ASM service model itself but are the result of specific choices of forwarding of ASM traffic across routers.

For example, these issues do not exist if the internetwork is simply a stateless broadcast domain such as a (non-switched) ethernet or wifi network, or if the network uses a stateless forwarding model in routers such as Bit Index Explicit Replication ([BIER]). Therefore the remainder of this section focusses on isues directly linked to the aspects specified in this document: ASM service model, host stack and some relevant L2 network technologies.

12.2. Receiver control

Senders in ASM can not control who receives their traffic because any host can join the group that the sender sends to. The larger address space of IPv6 multicast groups may make it harder for an IPv6 multicast address from being successfully discovered by undesired receivers, but many IPv6 multicast addresses are not random but well-defined. Encryption of ASM traffic and sharing of keys with only desired receivers is another solution against this challenge. For example, [GDOI] specifies a key management mechanism for secure sharing of symmetric group communication keys for ASM (which could also be applied to SSM).

Some types of deployed IP multicast based application services such as multicasting of high-value content do not consider such group encryption keys as secure enough alone, especially when they are shared between a large number of legitimate but not necessarily trustworthy receivers. A single impaired receiver may be set up to extract the shared key and pass it on to illegitimate receivers in real-time.

This has wideley happened in deployed solutions in the past with multicast/broadcast media content transmitted via IP multicast. In these cases, additional, per receiver, per host group authorization can be used to limit what IP multicast traffic is forwarded by the network to each host.

These receiver control options are often available in IP multicast implementations in network equipment but are not IETF standardized. Likewise, hardware and/or software solutions on hosts to prohibit such key extraction can be used. These are commonly called "Trusted Execution Environments" (TEE) and solutions applying them to prohibit content leakage are called "Digital Rights Managmeent" (DRM).

12.3. Sender control

Receivers in ASM can not control who is sending traffic to them.

Especially in IPv6 with its larger address space, random multicast group addresses (see Section 12.2) may help to limit undesired senders if all allowed senders and receivers can be trusted not to leak the secret address, and the network towards such legitimate senders and receivers can not easily be observed by attackers to determine the secret random address.

If deployed, network filtering may aid in restricting unexpected or unauthorized traffic.

This sender control problem is the same in unicast except that the methods or likelyhoods to keep destination host unicast addresses and ASM group addresses private vary significantly. There is no analysis of ASM group address privacy comparable to [RFC7721].

The [SSM] service model eliminates the sender control challenge by requiring receivers to explicitly indicate the desired sender of the multicast traffic. Using an appropriate forwarding method across the network, [SSM] is better than unicast in protecting against undesired traffic as it can often stop unwanted SSM traffic from even entering the network, whereas in unicast undesired traffic can only be discarded at the receiver. Note too, that an [SSM] host stack is an extension of the host-stack specified in this document. It only enhances further what is specified here but does not replace it.

12.4. Packet spoofing

Unless sender control is performed, packet spoofing may not even be necessary to perform equivalent attacks as outlined in Section 12.3. The ease of spoofing a sender IP source address and its layer 2 sender address (like sender MAC-address on ethernet) highly depends on the (inter)network between sender and receiver.

In a simple broadcast domain without active switches between sender and receiver, IP multicast packets are as easily spoofed as IP packets. If switches are introduced, without [IGMPsnooping], then IP multicast packets are equally easy to be spoofed because they are still broadcast, whereas IP packets become more difficult to spoof because attackers may not even see IP exchanges between a sender to spoof and its receivers, nor may it know their IP addresses.

Introducing [IGMPsnooping] somewhat levels the playing field and makes spoofing IP multicast packets more difficult, but as long as an attacker can be a valid receiver of IP multicast packets from a sender it wants to spoof and can guess the IP multicast group(s), it can also learn the source IP address and layer 2 address of the sender it wants to spoof by simply joining to its IP multicast traffic.

[Note: In internetworks, routers do typically perform RPF check for IP multicast packets as part of stateful forwarding of IP multicast packets, but this varies by the IP multicast routing / tree building protocol and is, as mentioned in Section 12.1 out of scope.]

Authentication of ASM/SSM traffic with mechanisms relying on symmetric group keys, such as [GDOI], can protect against many cases of spoofing, but it can not effectively prohibit sender spoofing by any of the legitimate receivers which could potentially be millions. This is, because each legitimate receiver knows the symmetric key required to become a sender. Asymmetric (public) key cryptography resolves this issue but is significantly more compute expensive than symmetric key cryptography. More advanced mechanisms tackling this issue, include TESLA [RFC4082] and its followup documents in [MSEC] as well as [I-D.ietf-mboned-ambi], [I-D.krose-mboned-alta] and [I-D.moskowitz-tesla-update-gnss-sbas].

12.5. Address management

Receiving IP multicast packets from undersired senders may not be malicious but can simply be a result of absent or incorrect IP multicast group address management that needs to assign unique group addresses to every application instance that needs them. Static allocation of IP multicast groups is the most widely used option in deployment today. Early proposals for dynamic address allocation protocols, including [MASC] and [MADCAP] have not gained traction and most options do not consider IPv6. See [RFC2908], [RFC6308].

12.5.1. Waste traffic in the absence of address management

While it is possible to forego address management and (randomnly) share IP multicast groups across multiple application instances simply by de-multiplexing at higher layers such as UDP ports and/or encryption layer selectors, relying solely on those higher layers for traffic separation is highly undesirable.

Assume an IP multicast application on host H1 joins to IP Multicast group G with traffic on UDP port P1. Other applications on other hosts are receivers for other IP Multicast applications that all (randomnly) also use G, but each uses a separate UDP Port P2, ... PN. H1 will receive traffic for all applications and discard the received packets in the UDP/socket layer because of their UDP ports.

This "waste traffic" can result in overload of resources in H1 and possible unexpected discarding of packets due to such overload. In switched networks with IGMP/MLD snooping and internetworks with IP multicast routers it can even lead to overload of network path segments towards H1 and discarding of packets to other hosts when traffic is admission controlled and this waste traffic is not taken into account.

12.5.2. Waste traffic due to layer 2 to layer 3 mapping

Hosts may need to receive and discard IP multicast packets in their IP module (typically in software) for host groups that they have not joined because of possible N:1 mapping issues in the layer 2 mapping of IP multicast. As described in Section 6.4, in IPv4 224.x.y.z, 224.(x+128).y.z, ..., 239.x.y.z, 239.(x+128).y.z will map to the same MAC address 01-00-5E-xx-yy-zz for x=0..127/xx=hex(x), y=0..255/yy=hex(y), z=0..255/zz=hex(z). For IPv6 over ethernet, similar mapping issues exist.

An only slightly overstated example is a broadcast network where few high-speed hosts receive a high bitrate IPv4 multicast video stream to address 239.128.0.251 and a very small, low-end CPU alarm siren has to be discovered via [mDNS] on 239.0.0.251. Both addresses map to Ethernet address 01-00-5E-00-00-FB. The software infrastructure (CPU, buffers) on the alarm siren gets overloaded by the high-bitrate IP multicast video stream because those packets are not filtered in the MAC hardware filter, and [mDNS] fails to discover the alarm siren when a fire in the building is discovered by a fire sensor.

These issues are resolved by avoiding the use of multiple IP multicast group addresses that map to the same ethernet MAC addresses. In practice, industry recommendations primarily focus on avoiding the use of IP multicast group addresses that map to statically assigned link-local IP multicast group addresses to avoid impacting key protocols such as [mDNS] in the example.

12.5.3. Multiple application instances

If two or more instances of the same (or similar enough in packet format) applications that do not well enough distinguish their instances through higher layer methods (transport layer ports, security selectors, application layer identification of instance) are instantiated and (erroneously) re-use the same IP multicast group, then this will not only cause the aforementioned waste traffic problems, that waste traffic can also leak into the application where it causes malfunction or other application security issues.

An example of this issue are protocols like [OSPFv2] which do not have instance differentiation in their packet format, so when supposedly separate instances of OSPF are incorrectly wired together, routing problems occur.

In [OSPFv2], the common solution against this issue is to rely on the authentication option and simply distinguish instances through separate passwords. This is a practical separation strategy, providing an instance identification to protect against accidental incorrect wiring.

Applications using well-known transport layer ports are likewise easily subject to this issue.

12.6. MAC filters

Joining to ASM multicast groups uses ressources in the host. The challenges in managing resource exhaustion and/or fair share across multiple applications are similar to those for unicast sockets except that filtering of packet reception at layer 2 will typically consume additional hardware limited filtering resources ("MAC filters").

13. Acknowledgements

Many thanks to Stig Veenas for his thorough review (WG chair). Many thanks to Brian Haberman, Dino Farinnacci, Alvaro Retana (RTG AD) and Jim Stevens, Pascal Thubert (INTDIR), Zheng Zhang (RTGDIR), Erik Nordmark (IOTDIR). Special thanks to Rob Hinden. Many thanks to Brian Weis (SECDIR), Kyle Rose and Rob Moskowitz for multicast security input.

14. References

14.1. 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., Ed., "Internet Group Management Protocol, Version 3", STD 100, RFC 9776, DOI 10.17487/RFC9776, March 2025, <https://www.rfc-editor.org/rfc/rfc9776>.

[MLDv2] Haberman, B., Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", STD 101, RFC 9777, DOI 10.17487/RFC9777, March 2025, <https://www.rfc-editor.org/rfc/rfc9777>.

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/rfc/rfc1122>.

[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>.

[RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <https://www.rfc-editor.org/rfc/rfc8504>.

[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>.

[STD5] Internet Standard 5, <https://www.rfc-editor.org/info/std5>. At the time of writing, this STD comprises the following:

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

Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <https://www.rfc-editor.org/info/rfc792>.

Mogul, J., "Broadcasting Internet Datagrams", STD 5, RFC 919, DOI 10.17487/RFC0919, October 1984, <https://www.rfc-editor.org/info/rfc919>.

Mogul, J., "Broadcasting Internet datagrams in the presence of subnets", STD 5, RFC 922, DOI 10.17487/RFC0922, October 1984, <https://www.rfc-editor.org/info/rfc922>.

Mogul, J. and J. Postel, "Internet Standard Subnetting Procedure", STD 5, RFC 950, DOI 10.17487/RFC0950, August 1985, <https://www.rfc-editor.org/info/rfc950>.

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

14.2. Informative References

[BIER] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <https://www.rfc-editor.org/rfc/rfc8279>.

[GDOI] Smyslov, V. and B. Weis, "Group Key Management Using the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 9838, DOI 10.17487/RFC9838, November 2025, <https://www.rfc-editor.org/rfc/rfc9838>.

[I-D.ietf-mboned-ambi] Holland, J., Rose, K., and M. Franke, "Asymmetric Manifest Based Integrity", Work in Progress, Internet-Draft, draft-ietf-mboned-ambi-05, 17 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-mboned-ambi-05>.

[I-D.ietf-pim-gaap] Farinacci, D. and M. McBride, "Group Address Allocation Protocol (GAAP)", Work in Progress, Internet-Draft, draft-ietf-pim-gaap-10, 25 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-pim-gaap-10>.

[I-D.ietf-pim-zeroconf-mcast-addr-alloc-ps] Karstens, N., Farinacci, D., and M. McBride, "Zeroconf Multicast Address Allocation Problem Statement and Requirements", Work in Progress, Internet-Draft, draft-

ietf-pim-zeroconf-mcast-addr-alloc-ps-13, 17 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-pim-zeroconf-mcast-addr-alloc-ps-13>.

[I-D.krose-mboned-alta]
Rose, K. and J. Holland, "Asymmetric Loss-Tolerant Authentication", Work in Progress, Internet-Draft, draft-krose-mboned-alta-01, 8 July 2019, <https://datatracker.ietf.org/doc/html/draft-krose-mboned-alta-01>.

[I-D.moskowitz-tesla-update-gnss-sbas]
Moskowitz, R. and R. Canetti, "TESLA Update for GNSS SBAS Authentication", Work in Progress, Internet-Draft, draft-moskowitz-tesla-update-gnss-sbas-01, 2 November 2025, <https://datatracker.ietf.org/doc/html/draft-moskowitz-tesla-update-gnss-sbas-01>.

[IANA.MASR]
IANA, "IPv6 Multicast Address Space Registry", Web https://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xhtml, n.d..

[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>.

[MADCAP] Hanna, S., Patel, B., and M. Shah, "Multicast Address Dynamic Client Allocation Protocol (MADCAP)", RFC 2730, DOI 10.17487/RFC2730, December 1999, <https://www.rfc-editor.org/rfc/rfc2730>.

[MASC] Radoslavov, P., Estrin, D., Govindan, R., Handley, M., Kumar, S., and D. Thaler, "The Multicast Address-Set Claim (MASC) Protocol", RFC 2909, DOI 10.17487/RFC2909, September 2000, <https://www.rfc-editor.org/rfc/rfc2909>.

[mDNS] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <https://www.rfc-editor.org/rfc/rfc6763>.

[MLDv1] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, DOI 10.17487/RFC2710, October 1999, <https://www.rfc-editor.org/rfc/rfc2710>.

[MPL] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", RFC 7731, DOI 10.17487/RFC7731, February 2016, <https://www.rfc-editor.org/rfc/rfc7731>.

[MSEC] "MSEC WG documents", Web https://datatracker.ietf.org/wg/msec/documents/.

[NTP] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <https://www.rfc-editor.org/rfc/rfc5905>.

[OSPFv2] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <https://www.rfc-editor.org/rfc/rfc2328>.

[PGM] 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, DOI 10.17487/RFC3208, December 2001, <https://www.rfc-editor.org/rfc/rfc3208>.

[PIM-SM] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <https://www.rfc-editor.org/rfc/rfc7761>.

[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>.

[RFC1884] Hinden, R., Ed. and S. Deering, Ed., "IP Version 6 Addressing Architecture", RFC 1884, DOI 10.17487/RFC1884, December 1995, <https://www.rfc-editor.org/rfc/rfc1884>.

[RFC2365] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, DOI 10.17487/RFC2365, July 1998, <https://www.rfc-editor.org/rfc/rfc2365>.

[RFC2908] Thaler, D., Handley, M., and D. Estrin, "The Internet Multicast Address Allocation Architecture", RFC 2908, DOI 10.17487/RFC2908, September 2000, <https://www.rfc-editor.org/rfc/rfc2908>.

[RFC2974] Handley, M., Perkins, C., and E. Whelan, "Session Announcement Protocol", RFC 2974, DOI 10.17487/RFC2974, October 2000, <https://www.rfc-editor.org/rfc/rfc2974>.

[RFC3171] Albanna, Z., Almeroth, K., Meyer, D., and M. Schipper, "IANA Guidelines for IPv4 Multicast Address Assignments", RFC 3171, DOI 10.17487/RFC3171, August 2001, <https://www.rfc-editor.org/rfc/rfc3171>.

[RFC3232] Reynolds, J., Ed., "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, DOI 10.17487/RFC3232, January 2002, <https://www.rfc-editor.org/rfc/rfc3232>.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <https://www.rfc-editor.org/rfc/rfc3376>.

[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>.

[RFC3810] 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>.

[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>.

[RFC4082] Perrig, A., Song, D., Canetti, R., Tygar, J. D., and B. Briscoe, "Timed Efficient Stream Loss-Tolerant Authentication (TESLA): Multicast Source Authentication Transform Introduction", RFC 4082, DOI 10.17487/RFC4082, June 2005, <https://www.rfc-editor.org/rfc/rfc4082>.

[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>.

[RFC5771] Cotton, M., Vegoda, L., and D. Meyer, "IANA Guidelines for IPv4 Multicast Address Assignments", BCP 51, RFC 5771, DOI 10.17487/RFC5771, March 2010, <https://www.rfc-editor.org/rfc/rfc5771>.

[RFC5790] 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>.

[RFC6034] Thaler, D., "Unicast-Prefix-Based IPv4 Multicast Addresses", RFC 6034, DOI 10.17487/RFC6034, October 2010, <https://www.rfc-editor.org/rfc/rfc6034>.

[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>.

[RFC6308] Savola, P., "Overview of the Internet Multicast Addressing Architecture", RFC 6308, DOI 10.17487/RFC6308, June 2011, <https://www.rfc-editor.org/rfc/rfc6308>.

[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>.

[RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", RFC 7721, DOI 10.17487/RFC7721, March 2016, <https://www.rfc-editor.org/rfc/rfc7721>.

[RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <https://www.rfc-editor.org/rfc/rfc8085>.

[RFC8313] Tarapore, P., Ed., Sayko, R., Shepherd, G., Eckert, T., Ed., and R. Krishnan, "Use of Multicast across Inter-domain Peering Points", BCP 213, RFC 8313, DOI 10.17487/RFC8313, January 2018, <https://www.rfc-editor.org/rfc/rfc8313>.

[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>.

[RFC8866] Begen, A., Kyzivat, P., Perkins, C., and M. Handley, "SDP: Session Description Protocol", RFC 8866, DOI 10.17487/RFC8866, January 2021, <https://www.rfc-editor.org/rfc/rfc8866>.

[RFC9466] Liu, Y., Ed., Eckert, T., Ed., McBride, M., and Z. Zhang, "PIM Assert Message Packing", RFC 9466, DOI 10.17487/RFC9466, October 2023, <https://www.rfc-editor.org/rfc/rfc9466>.

[RFC9542] Eastlake 3rd, D., Abley, J., and Y. Li, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", BCP 141, RFC 9542, DOI 10.17487/RFC9542, April 2024, <https://www.rfc-editor.org/rfc/rfc9542>.

[RIPv2] Malkin, G., "RIP Version 2 - Carrying Additional Information", RFC 1723, DOI 10.17487/RFC1723, November 1994, <https://www.rfc-editor.org/rfc/rfc1723>.

[RPL] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <https://www.rfc-editor.org/rfc/rfc6550>.

[RTP] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <https://www.rfc-editor.org/rfc/rfc3550>.

[TAPS] "TAPS WG documents", Web https://datatracker.ietf.org/wg/taps/documents/.

[UDP] Sluizer, S. and J. Postel, "Mail Transfer Protocol: ISI TOPS20 MTP-NIMAIL interface", RFC 786, DOI 10.17487/RFC0786, July 1981, <https://www.rfc-editor.org/rfc/rfc786>.

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

A.2.1. Original RFC1112 text

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. See Section 12 for more details.

A.2.2. Evolution since RFC1112

Historically (1990th), SDP [RFC8866] over SAP ([RFC2974] was used to multicast the session information of application sessions with transient IPv4 multicast addresses via the MBone's sdr tool. When creating a new application instance, sdr would simply avoid picking any of the already assigned IPv4 multicast addresses as known from other SDP/SAP announcements.

[RFC6308] section 3.5 summaries and explains the lack of adoption of mechanisms specified in RFCs since RFC1112 for allocation of transient host group addresses.

Current evolving mechanisms for zeroconf dynamic host group addreses are based on [I-D.ietf-pim-zeroconf-mcast-addr-alloc-ps]. It specifically requires solutions to allocate addresses so that different application do not use IP multicast addresses that map to the same MAC address. Such mapping would defeat the filtering of most IGMP/MLD snooping switches which most often only operate on MAC level. One protocol supporting these requirements is [I-D.ietf-pim-gaap].

A.3. Link local considerations

IGMP/MLD are required for Level 2 hosts on a subnet so that IP multicast routers on the same subnet can forward traffic from sender on another subnet. IGMP/MLD are technically not required for IP multicast packets with link local addresses on a broadcast subnet. Such link local IP multicast can not have senders from a different subnets (Section 8), and IP multicast traffic from senders on the same subnet is forwarded to all hosts on the subnet on a broadcast subnet.

If IGMP/MLD snooping are used in a subnet, it looses its broadcast nature for IP multicast traffic and sending of IGMP/MLD messages is often also required to receive IP multicast traffic from local senders in the same subnet. Subject to how the IGMP/MLD snooping switch operates.

For IPv4 link local addresses, IGMP snooping must not filter traffic due to the historic non-implementation of IGMP in early hosts (such as IP routers using [OSPFv2]) and the limited number of link local groups in IPv4 make it quite unimportant to improve.

In IPv6, there are as many link local addresses as there are in other scopes. Therefore link local IPv6 is well usable for non network control traffic, such as any IP multicast application that wants its traffic to be constrained to a single subnet (such as large, single-subnet campus networks with MLD snooping switches). Hence the requirement in [MLDv2] to also signal MLD for link local IPv6 multicast addresses.

However, subnets can be known to never have IGMP/MLD snooping switches, such as radio subnets in IoT mesh networks. And those subnets can also be known to never require any IP multicast traffic other than link local IP multicast protocol packets that use the host stack. An example of this are radio subnets solely between [RPL] nodes.

For such subnets, nodes can act as Level 2L hosts, hence avoiding the unnecessary complexity of IGMP/MLD and the radio energy it would cost.

A.4. IP multicast router considerations

Nodes that are IP multicast routers do typically not use IGMP/MLD to indicate the multicast groups or channels that they need to receive from a subnet. Nor could they in many cases.

IGMP/MLD snooping can not constrain IP multicast traffic to any port with such an IP multicast router connected to it because of this. This is true even in the most simple subnet setup with only IP multicast hosts and just one IP multicast router acting as the IGMP/MLD querier. That gateway needs to see all IP multicast traffic sent by any host onto the subnet - to determine which of that traffic to forward to receivers on other subnets.

However, IP multicast routers may also use the IP multicast host stack specified in this document. Consider the case where an IP multicast router sends and/or receives IP multicast packets not because of its IP multicast forwarding function as an IP multicast router, but because these IP multicast packets are sent or received by an application on the gateway, including but not limited to IGMP/MLD protocol packets or IP multicast routing protocol packets such as [PIM-SM] IP multicast protocol packets. Without further gateway considerations, these packets are logically subject to the host stack requirements of this document.

For example, an IP multicast router running IGMPv3 would need to indicate to its own IP host stack the desire to receive packets for 224.0.0.22, resulting in it also sending IGMP membership reports for that address.

However, as explained before, even in the presence of IGMP/MLD snooping switches, IP multicast routers need to receive any IP multicast packet on a subnet without itself sending IGMP/MLD messages to join the traffic. And IGMP/MLD snoopping switches support this by manualor automatic detection of ports connected to IP multicast routers. Hence an IP multicast router can safely forego sending IGMP/MLD membership messages for any IP multicast addresses it is joined to as a host: It will receive the traffic anyhow, even in the presence of IGMP/MLD snooping switches, and it will anyhow use its non-IGMP/MLD multicast routing protocol to ensure traffic from other subnets gets forwarded to the subnet.

In summary: For hosts (or gateways) that are also IP multicast routers the Level 2 host stack may skip sending IGMP/MLD membership reports to receive IP multicast packets when this is deemed specifically beneficial. This can simply be justified as a case where the behavior of an IP multicast router (which is outside of scope of this document) superceeds the requirements of the host stack as specified here, even if the host stack of the gateway is devices from the specification of this document.

Note that this document gives no recommendations to do this, this appendix purely explains how this could work and be justified when needed - without violating this specifications requirements. Given how most IP multicast routers are just optionally configurable as IP gateways, they would need to conform to the full L2 host stack reqiurements whenever they do not act as an IP multicast gateway, hence optimizing the host stack purely to reduce the amount of code is not an option in those cases.

A.5. Application Socket Security Considerations

The following section addresses socket security issues beyond the scope of this document. While they are in general independent of the transport protocol used, they most often happen for UDP because of the prevalence of using IP multicast with UDP and because even if other applications for IP multicast exist on hosts (such as [OSPFv2]), in most hosts, only UDP can be used for IP multicast by unprivileged and hence more likely malicious applications. The following considerations are not covered by [RFC8085] or resolved trough the requirements specified by [TAPS] RFCs.

Even with correct IP multicast group address management (Section 12.5), or when using SSM: with just the methods specified in this document for the host stack, application sockets may still receive unexpected IP multicast traffic destined to other IP multicast addresses than they joined to.

This problem can exist because like RFC1112, this memo only specifies the host stack up to the IP layer and hence does not include the specification that ASM group membership (or SSM channel membership) has to be per (transport layer) application socket.

In result, early host stacks for IPv4 multicast did indeed have the problem that two UDP sockets each joining to a different IPv4 multicast address but the same UDP port would receive traffic destined to either IPv4 multicast addresses. And could accordingly cause application malfunctions or other security issues. Such port re-use can easily happen when applications define the use of a well-known UDP port number and just expect (like they should), that different application instances can just use different IP multicast addresses.

A.5.1. IGMPv3/MLDv2

In current host stacks for Level 2 hosts, this problem is usually eliminated when implementations correctly implement the following sentence present in IGMP/MLD specifications since [RFC3376]/[RFC3810].

After a multicast packet has been accepted from an interface by the IP layer, its subsequent delivery to the application or process that listens on a particular socket depends on the multicast listening state of that socket...

A.5.2. Level 2L

Level 2L implementation would equally have to implement their host stack using such per-socket membership even in the absence of IGMP to support equivalent demultiplexing replication and filtering on a per socket basis for received IP multicast packets. Otherwise this filtering would be left up to the application, not only violating reasonable per-socket expectations but also incurring unnecessary overhead: Unnecessary replication and process-level processing of such unnecessary packet copies.

A.6. Application socket issues

The following issues relate to the current behavior of known (transport layer) application sockets across various operating systems. These behaviors evolved by simply not improving the behavior of BSD sockets for IP multicast from a security perspective and proliferation of that socket model across other operating systems and POSIX standard.

Host stacks by default do not allow multiple application sockets to bind() to the same transport layer port (TCP, UDP or other). This is highly desirable in IP unicast because it guarantees the application with the socket that no other application can be a responder/"server" for that port on the same host/IP-address(es). Likewise, any responder/"client" application can (implicitly or explicitly) bind() to a dynamic, unused port due to the nature of IP unicast initiator/ responder protocol exchanges.

In IP multicast the default for socket operations is the same, but the impact on IP multicast applications is different. In [UDP], [PGM] or any other IP multicast capable transport protocols using the notion of Source Port and Destination Port, the port that a socket binds to is like for IP unicast traffic the Source Port for packets sent and the Destination Port for packets received.

When an IP multicast receiver application binds to a port, by default no other application on the same host can receive the same IP multicast traffic. This is not only undesirable when multiple receiver applications for the IP multicast application instance are desired to be to run on the same host simultaneously, but a malicious attacker application started before a legitimate receiver application can perform a DoS attack against these IP multicast receiver ("client") applications by binding to the known transport layer port that the sender(s) sends to.

The comparable attack is not possible in IP (unicast) because the as mentioned above, the client application (unicast initiator) can bind to any free port and then negotiate with the sender that it sends to that Destination Port. In IP multicast the sender of course can not negotiate with every receiver a separate receiver Destination Port. It must send IP multicast to one port common for all receivers, which then makes that port subject to the attack.

Enabling re-binding to the same UDP port on sockets used to receive IP multicast traffic (SO_REUSEADDR/SO_REUSEPORT) allows benevolent applications on the same host to receive the same IP multicast traffic, but known host stacks have no option to force this option on all (receiver) IP multicast sockets to prohibit the aforementioned

attack. Simply because there is no concept of an IP multicast receiver only socket, and forcing re-use of ports would in most cases be wrong for other type of sockets.

For an IP multicast sender application, the attack is different. A malicious application binding to a socket can not prohibit a legitimate sender application to send to the same port. Which it could do in IP (unicast). However, an IP multicast sender binding to a port can not rely on the fact that there is no malicious application on the same host sending to the same IP multicast group and Destination Port because the bind only guarantees exclusive use of the Source Port, which is irrelevant in most IP multicast application stacks, for example when using [RTP]. Arguably, the IP multicast problem is bigger because an IP server application will know at bind() time when it can not exclusively use the relevent port because of the prior presence of a malicious application on the same host, whereas in IP multicast, the server can not prohibit that a later started malicious application on the same host is impersonating packets with the same Source IP address, IP multicast address and Destination Port number as the legitimate server application.

IP multicast applications could recognize the attacking application based on its Source Port instead of only its Source IP address, but that is not common in IP multicast applications / specifications today, such as when using [RTP]. Even worse, the legitimate sender applications itself may not even be able to recognize packets from the malicious sender on the same host if the socket interface allows to prohibit looping back of IP multicast packets from one socket to any other socket on the local host (IP_MULTICAST_LOOP). Which is a commonly supported option in todays socket APIs.

In summary, malicious local applications do pose different and potentially more severe risks to IP multicast sender and receiver applications than malicious IP multicast applications running on other hosts with todays application socket semantics.

Appendix B. Discussion and Explanations (TO BE REMOVED)

[RFC-editor: Please remove this Appendix after observing the following section addressed to you]

Please refer to Section 10 for the non-process discussion of the goals of this document.

B.1. RFC-Editor notes

The kramdown tooling did not allow to have references for both STD5 and RFC1112, those fail because the STD5 reference creates an "RFC1112" anchor. Thus there is no separate reference for RFC1112 in this version of the document. This needs to be fixed in XML by adding a full reference to RFC1112 and removing the RFC1112 anchor from the STD5 reference.

B.2. Goals and evolution of this document

The initial goal of this document was to allow for IETF to declare the IGMPv1 protocol historic which today is a Full Internet Standard due to it being defined in RFC1112. This should be achieved without changing the Full Internet Standard status of the IP Host Extensions for IP multicast and ASM IP service interface specified in RFC1112 because those specification are as fundamental to the definition of IP multicast as RFC791 is for IP (unicast).

The best way to achieve this seemed to be an update to RFC1112 which removes all of IGMPv1, but maintains the rest of the document. None of these removal of IGMPv1 changes changed the applicability or requirements to existing IP multicast (plus its protocols) implementations or other specifications.

The next refinement was to rectify the situation that there is no specification explaining the same details as RFC1112 for IPv6 multicast even though RFC8200 (full internet standard) even explicitly includes IPv6 multicast, and a range of other RFC define necessary code-points (such as for ethernet mapping) for IPv6 multicast.

Most of the text of this specification can hence can simply talk about "IP" which in this specification implies both IPv4 and IPv6, and only in places where IPv6 differs, does the document now include new explicit text, most often pointing to pre-existing RFCs specifying the necessary details for IPv6. Again, none of these changes impact other specs or deployments.

The third step of refinement was add the necessary verbiage to explain the differences between SSM and the specifications in this document. None of these text enhancements incur any functional changes of long-term established practices. Instead, they are only resulting in references to SSM RFCs, introduction of the term ASM (which was previously only defined in SSM RFCs), and the limitation of applicability of terms in this document (such as host group) to their use with ASM.

The last round of changes added and refined details to be in-line with long-term established practices and removing any possible contradictions between the original RFC1112 text and newer standards track specification such as IGMPv2/MLDv3 or long term established implementation practices. This includes the limitation of scope of ASM to controlled networks and the definition of the IPv4 Link-Local address range, which so far had only been defined through BCP RFC, unlike in IPv6, where it’s part of the architecture, as well as permitting (but not recommending) non-use of IGMP for them.

In summary, all changes in the document will make this document a replacement of rfc1112 which much more reflects the full internet standard nature of the technology than rfc1112 did as of recent.

B.3. Update to RFC791

This version of the text proposes that this spec is declared to be an update to RFC791.

The argument made in Section 9.3 to support this classification may not be persuasive enough (because the according rfc791 text may be read as a perfectly good extension point specification), in which case the update status and related text should be deleted.

However, If anyone where to come up with a re-use of 224.0.0.0/4 for any non-IP multicast purposes, havoc might ensure with devices that do assume IP multicast semantics, so it may simply be prudent to include this declaration. It would also make the relationship between IPv4 and IPv4 multicast be more aligned with IPv6, where IPv6 multicast is included in RFC8200.

B.4. Changelog

This document is hosted at https://github.com/toerless/rfc1112bis. Please submit issues with this text as issues to that github and report them on pim@ietf.org.

B.4.1. draft-ietf-pim-rfc1112bis-08

Revision including fixes for nits uncovered by directorate reviews. Eric Nordmark (INTDIR), Brian Weis feedback (SECDIR), Sandy Zhang (RTGDIR), Pascal Thubert (INTDIR)

Several textual nits fixed, not detailling.

Section 6.4: Extensions to Ethernet module

Reordered expand last two paragraphs. Added reference to very recent relevant RFC9542 IANA registry for MAC addresses (from Pascal).

Section 9.2: Compatibility with IGMPv1

Added definition for "backward compatibility" (with IGMPv1) and refined wording.

Section 10: changes over RFC1112:

10.8 removed second paragraph, just pointing to A.3

10.14: Added sub section 10.14 to more comprehensively discuss the correct terms "host group" / "host group address" vs "SSM destination address" vs (ambiguous "SSM multicast address"). RFCs applying equally to ASM/SSM can just use "IP multicast address". RFCs applying to only ASM should use "IP host group address" - etc. pp.

11.7 (IANA asks)

added request to replace rfc1112 with thisRFC for the new RFC9542 registry for MAC addresses.

12.3 security considerations

clarified/refined/expanded sender control text (Eric, Brian).

A.2 appendix for discussion about transient IP multicast addresses

moved existing text from Steve Deering (mostly never realized ideas) to historic subsection, added subsection for best decade-long solution SAP/SDP and paragraph about new evolving solution GAAP.

Various textual nits.

Added text/reference for recent new IANA Multicast / 802x address registries

Added explanation for ND in Figure 1.

Other:

Changed TO_BE_REMOVED_SECTION to indicate keeping it also for IETF/ IESG review as it seems usefulto help such further broader reviewers.

B.4.2. draft-ietf-pim-rfc1112bis-07

Revision for early reviews from directorates. Added to-be-removed contextual explanations for those reviews.

B.4.3. draft-ietf-pim-rfc1112bis-06

Added To-Be-Removed note for reviewers to compare with rfc1112 to find pre-existing sections.

Removed erroneous reference to UDP in 7.1 (socket calls in referenced docs are not specific to UDP).

Changed order of authors.

Included fixes from Stig Veenas' review:

Variety of typos.

Expanded "protocol field in IP header" to be explicit about the complex IPv6 options.

Clarified that "IP multicast address" covers host group and SSM channel destination addresses and fixed text that applies to both ASM and SSM touse "IP multicast address" instead of host group (address).

removed IGMPv3lite term

Added 6 pages of Security Considerations and two pages of Appendix for application socket security considerations.

B.4.4. draft-ietf-pim-rfc1112bis-05

Brian pointing to the requirement to support link-local IPv6 multicast in RFC4291, section 2.8, accordingly changed the requirement to MUST for Level 2L and explanation about that.

B.4.5. draft-ietf-pim-rfc1112bis-04

1. Some textual nit improvements - introduced "all-nodes also for IPv6 (but be careful to only call it Link-Local, as there are scope relative ones too), adding references to RFC8504, referring to "host-side" impleemntation of IGMP/MLD. Shoveling sentence in 4. to make reading more logical.

2. "Levels of Conformance": Made support for IP multicast (Level 2 = sending/receiving) RECOMMENDED for all IPv4 / IPv6 host stack. For the past 36 years, there was only the RFC1122 requirement

(see below) for IPv4. For IPv6 there was no requirement to support IPv6 multicast at all. Instead, there was only a dependency to support it when implementing widespread IPv6 protocols (SLAAC, ND).

3. Section 3.4: Introduction of conformance Level 2L to describe IPv4 multicast with link-local only sending/receiving. Primarily because RFC1122 specified it, but also because there are sufficiently many devices that do implement this at their core - e.g.: router operating systems in suport of OSPF etc (most have been updated to also support IGMP.

4. Section 7.2: (re-)introduced permanent joining of all-groups as a SHOULD requirement.

5. Section 9.4 and header: Defining this doc as update to RFC1122 to override the 36 year long recommendation of only implementing IP multicast without IGMP.

6. New sections 10.7 to explain RFC1122 and Level 2L

7. New section 10.8 to explain/justify recommendation to SHOULD support IP multicast on all hosts.

8. Rewrote Section 10.10 for permanently join all-nodes group.

B.4.6. draft-ietf-pim-rfc1112bis-03

1. Changed document text to make the term "ASM" apply only to the IP service interface (extensions) specified by the document (and shown and explained in existing text), instead of the whole host extensions specified in this document (as it was written up to up to -02). This is the only correct semantic, given how all the host exensions specified in this document are shared by SSM, only the IP service interface is changed/amended by SSM.

2. Subdivided section 2 (INTRODUCTION) into sections 2.1 (Summary), which contains new text from this spec, and 2.2 (Overview), which is unchanged RFC1112 text. Newly written section 2.1 to summarize the key content of this document. This was so far only explained in the much later changes from rfc1112 section. Includes IPv4/IPv6 applicability, ASM/SSM naming and maintaining most of RFC1112 text as a goal.

3. Introduced text to define and explain link local IPv4 host group addresses 224.0.0.0 - 224.0.0.255. This was triggered by trying to fix the rfc1112 text sections that Brian Haberman was concerned about, which did cover behavior for 224.0.0.1.

As it turns out, the behavior for 224.0.0.1 was quickly adopted by other protocols getting 224.0.0.0/24 addresses and there has been no functional specification to explain the non-forwarding behavior for these link-local addresses. Instead, only IANA allocation guideline RFCs where introducing them. This is now rectified with new explanatory text in this spec. and a new MAY requirement to permit non-use of IGMP for those groups. See Section 7.2.

1. Changed references to IGMPv3 and MLDv2 to the -bis drafts currently in RFC-editor queue. Also triggered by Brian Haberman mentioning them.

2. Improved wording in "(Normative) Status Change" section 9.

5.1 Removed "Update to rfc791" as an open issue and instead claimed it as fact in section 9.3. Added discussion about this point to the discussion appendix that is to be removed by RFC-editor.

5.1 Also added subsection to declare that this document replaces RFC1112 in STD5.

1. Enhanced/New text in section 10., "changes from RFC1112"

Especially explaining the changes in the normative section explained above and below, triggered by Brian’s review.

1. Applying changes proposed by Brian Haberman during WGLC.

7.1 Changed meaning of IP from "IPv4" to "IPv4 and IPv6", accordingly updated all text. Makes a lot of sense given the goal of showing how most of the IP multicast host stack operates the same for IPv4 and IPv6.

7.2 Re-added requirement for routers not to forward link-local multicast

7.3 adding MAY requirement to allow non-signaling of Link-Local scope IPv4 multicast and IPv6 all-nodes group, and explanations how this is better than the prior definitions from rfc1112. Also includes new (length) Appendix A.3 to justify this for IPv4.

7.4 text nits (thanks, Brian).

B.4.7. draft-ietf-pim-rfc1112bis-02

Removed unused references, fefresh - waiting for more reviews. Added IANA section for updates from RFC1112 to RFC1112bis. Added references to RFC5771 and RFC6034 because they actually are the references for the IANA 224.0.0.0/4 registrations, which seems a bit undocumented given how RFC1112 did introduce the definition (before IANA).

B.4.8. draft-ietf-pim-rfc1112bis-01

Fix up reference for IGMPv3. Refined candidate open issues. Removed author discussion.

B.4.9. 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 Dino’s 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.4.10. draft-ietf-pim-rfc1112bis-00

Just changed title, added github pointer.

B.4.11. 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.4.12. draft-eckert-pim-rfc1112bis-00

Initial version based on RFC1112 text version, edited.

Authors' Addresses

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

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