

Network Working Group
Request for Comments: 4607
Category: Standards Track

H. Holbrook
Arastra, Inc.
B. Cain
Acopia Networks
August 2006

Source-Specific Multicast for IP

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

IP version 4 (IPv4) addresses in the 232/8 (232.0.0.0 to 232.255.255.255) range are designated as source-specific multicast (SSM) destination addresses and are reserved for use by source-specific applications and protocols. For IP version 6 (IPv6), the address prefix FF3x::/32 is reserved for source-specific multicast use. This document defines an extension to the Internet network service that applies to datagrams sent to SSM addresses and defines the host and router requirements to support this extension.

Table of Contents

| | |
|--|----|
| 1. Introduction | 3 |
| 2. Semantics of Source-Specific Multicast Addresses | 5 |
| 3. Terminology | 6 |
| 4. Host Requirements | 7 |
| 4.1. Extensions to the IP Module Interface | 7 |
| 4.2. Requirements on the Host IP Module | 8 |
| 4.3. Allocation of Source-Specific Multicast Addresses | 9 |
| 5. Router Requirements | 10 |
| 5.1. Packet Forwarding | 10 |
| 5.2. Protocols | 10 |
| 6. Link-Layer Transmission of Datagrams | 11 |
| 7. Security Considerations | 12 |
| 7.1. IPsec and SSM | 12 |
| 7.2. SSM and RFC 2401 IPsec Caveats | 12 |
| 7.3. Denial of Service | 13 |
| 7.4. Spoofed Source Addresses | 13 |
| 7.5. Administrative Scoping | 14 |
| 8. Transition Considerations | 14 |
| 9. IANA Considerations | 15 |
| 10. Acknowledgements | 15 |
| 11. Normative References | 16 |
| 12. Informative References | 17 |

1. Introduction

The Internet Protocol (IP) multicast service model is defined in RFC 1112 [RFC1112]. RFC 1112 specifies that a datagram sent to an IP multicast address (224.0.0.0 through 239.255.255.255) is delivered to each "upper-layer protocol module" that has requested reception of datagrams sent to address G. RFC 1112 calls the network service identified by a multicast destination address G a "host group". This model supports both one-to-many and many-to-many group communication. This document uses the term "Any-Source Multicast" (ASM) to refer to model of multicast defined in RFC 1112. RFC 3513 [RFC3513] specifies the form of IPv6 multicast addresses with ASM semantics.

IPv4 addresses in the 232/8 (232.0.0.0 to 232.255.255.255) range are currently designated as source-specific multicast (SSM) destination addresses and are reserved for use by source-specific applications and protocols [IANA-ALLOC].

For IPv6, the address prefix FF3x::/32 is reserved for source-specific multicast use, where 'x' is any valid scope identifier, by [IPv6-UBM]. Using the terminology of [IPv6-UBM], all SSM addresses must have P=1, T=1, and plen=0. [IPv6-MALLOC] mandates that the network prefix field of an SSM address also be set to zero, hence all SSM addresses fall in the FF3x::/96 range. Future documents may allow a non-zero network prefix field if, for instance, a new IP-address-to-MAC-address mapping is defined. Thus, address allocation should occur within the FF3x::/96 range, but a system should treat all of FF3x::/32 as SSM addresses, to allow for compatibility with possible future uses of the network prefix field.

Addresses in the range FF3x::4000:0001 through FF3x::7FFF:FFFF are reserved in [IPv6-MALLOC] for allocation by IANA. Addresses in the range FF3x::8000:0000 through FF3x::FFFF:FFFF are allowed for dynamic allocation by a host, as described in [IPv6-MALLOC]. Addresses in the range FF3x::0000:0000 through FF3x::3FFF:FFFF are invalid IPv6 SSM addresses. ([IPv6-MALLOC] indicates that FF3x::0000:0001 to FF3x::3FFF:FFFF must set P=0 and T=0, but for SSM, [IPv6-UBM] mandates that P=1 and T=1, hence their designation as invalid.) The treatment of a packet sent to such an invalid address is undefined -- a router or host MAY choose to drop such a packet.

Source-specific multicast delivery semantics are provided for a datagram sent to an SSM address. That is, a datagram with source IP address S and SSM destination address G is delivered to each upper-layer "socket" that has specifically requested the reception of datagrams sent to address G by source S, and only to those sockets. The network service identified by (S,G), for SSM address G and source

host address S, is referred to as a "channel". In contrast to the ASM model of RFC 1112, SSM provides network-layer support for one-to-many delivery only.

The benefits of source-specific multicast include:

Elimination of cross-delivery of traffic when two sources simultaneously use the same source-specific destination address. The simultaneous use of an SSM destination address by multiple sources and different applications is explicitly supported.

Avoidance of the need for inter-host coordination when choosing source-specific addresses, as a consequence of the above.

Avoidance of many of the router protocols and algorithms that are needed to provide the ASM service model. For instance, the "shared trees" and Rendezvous Points of the PIM - Sparse Mode (PIM-SM) protocol [PIM-SM] are not necessary to support the source-specific model. The router mechanisms required to support SSM are in fact largely a subset of those that are used to support ASM. For example, the shortest-path tree mechanism of the PIM-SM protocol can be adapted to provide SSM semantics.

Like ASM, the set of receivers is unknown to an SSM sender. An SSM source is provided with neither the identity of receivers nor their number.

SSM is particularly well-suited to dissemination-style applications with one or more senders whose identities are known before the application begins. For instance, a data dissemination application that desires to provide a secondary data source in case the primary source fails over might implement this by using one channel for each source and advertising both of them to receivers. SSM can be used to build multi-source applications where all participants' identities are not known in advance, but the multi-source "rendezvous" functionality does not occur in the network layer in this case. Just like in an application that uses unicast as the underlying transport, this functionality can be implemented by the application or by an application-layer library.

Multicast resource discovery of the form in which a client sends a multicast query directly to a "service location group" to which servers listen is not directly supported by SSM.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

This document defines the semantics of source-specific multicast addresses and specifies the policies governing their use. In particular, it defines an extension to the Internet network service that applies to datagrams sent to SSM addresses and defines host extensions to support the network service. Hosts, routers, applications, and protocols that use these addresses MUST comply with the policies outlined in this document. Failure of a host to comply may prevent that host or other hosts on the same LAN from receiving traffic sent to an SSM channel. Failure of a router to comply may cause SSM traffic to be delivered to parts of the network where it is unwanted, unnecessarily burdening the network.

2. Semantics of Source-Specific Multicast Addresses

The source-specific multicast service is defined as follows:

A datagram sent with source IP address *S* and destination IP address *G* in the SSM range is delivered to each host socket that has specifically requested delivery of datagrams sent by *S* to *G*, and only to those sockets.

Where, using the terminology of [IGMPv3],

"socket" is an implementation-specific parameter used to distinguish among different requesting entities (e.g., programs or processes or communication end-points within a program or process) within the requesting host; the socket parameter of BSD Unix system calls is a specific example.

Any host may send a datagram to any SSM address, and delivery is provided according to the above semantics.

The IP module interface to upper-layer protocols is extended to allow a socket to "Subscribe" to or "Unsubscribe" from a particular channel identified by an SSM destination address and a source IP address. The extended interface is defined in Section 4.1. It is meaningless for an application or host to request reception of datagrams sent to an SSM destination address *G*, as is supported in the any-source multicast model, without also specifying a corresponding source address, and routers MUST ignore any such request.

Multiple source applications on different hosts can use the same SSM destination address *G* without conflict because datagrams sent by each source host *S_i* are delivered only to those sockets that requested delivery of datagrams sent to *G* specifically by *S_i*.

The key distinguishing property of the model is that a channel is identified (addressed) by the combination of a unicast source address and a multicast destination address in the SSM range. So, for example, the channel

`S,G = (192.0.2.1, 232.7.8.9)`

differs from

`S,G = (192.0.2.2, 232.7.8.9),`

even though they have the same destination address portion. Similarly, for IPv6,

`S,G = (2001:3618::1, FF33::1234)`

and

`S,G = (2001:3618::2, FF33::1234)`

are different channels.

3. Terminology

To reduce confusion when talking about the any-source and source-specific multicast models, we use different terminology when discussing them.

We use the term "channel" to refer to the service associated with an SSM address. A channel is identified by the combination of an SSM destination address and a specific source, e.g., an (S,G) pair.

We use the term "host group" (used in RFC 1112) to refer to the service associated with "regular" ASM multicast addresses (excluding those in the SSM range). A host group is identified by a single multicast address.

Any host can send to a host group, and similarly, any host can send to an SSM destination address. A packet sent by a host S to an ASM destination address G is delivered to the host group identified by G. A packet sent by host S to an SSM destination address G is delivered to the channel identified by (S,G). The receiver operations allowed on a host group are called "join(G)" and "leave(G)" (as per RFC 1112). The receiver operations allowed on a channel are called "Subscribe(S,G)" and "Unsubscribe(S,G)".

The following table summarizes the terminology:

| | | |
|----------------------|-------------|------------------------|
| Service Model: | any-source | source-specific |
| Network Abstraction: | group | channel |
| Identifier: | G | S,G |
| Receiver Operations: | Join, Leave | Subscribe, Unsubscribe |

We note that, although this document specifies a new service model available to applications, the protocols and techniques necessary to support the service model are largely a subset of those used to support ASM.

4. Host Requirements

This section describes requirements on hosts that support source-specific multicast, including:

- Extensions to the IP Module Interface
- Extensions to the IP Module
- Allocation of SSM Addresses

4.1. Extensions to the IP Module Interface

The IP module interface to upper-layer protocols is extended to allow protocols to request reception of all datagrams sent to a particular channel.

Subscribe (socket, source-address, group-address, interface)

Unsubscribe (socket, source-address, group-address, interface)

where

"socket" is as previously defined in Section 2,

and, paraphrasing [IGMPv3],

"interface" is a local identifier of the network interface on which reception of the channel identified by the (source-address,group-address) pair is to be enabled or disabled. A special value may be used to indicate a "default" interface. If reception of the same channel is desired on multiple interfaces, Subscribe is invoked once for each.

The above are strictly abstract functional interfaces -- the functionality can be provided in an implementation-specific way. On a host that supports the multicast source filtering application programming interface of [MSFAPI], for instance, the Subscribe and Unsubscribe interfaces may be supported via that API. When a host has been configured to know the SSM address range (whether the configuration mechanism is manual or through a protocol), the host's operating system SHOULD return an error to an application that makes a non-source-specific request to receive multicast sent to an SSM destination address.

A host that does not support these IP module interfaces (e.g., ASM-only hosts) and their underlying protocols cannot expect to reliably receive traffic sent on an SSM channel. As specified below in Section 5.2, routers will not set up SSM forwarding state or forward datagrams in response to an ASM join request.

Widespread implementations of the IP packet reception interface (e.g., the `recvfrom()` system call in BSD Unix) do not allow a receiver to determine the destination address to which a datagram was sent. On a host with such an implementation, the destination address of a datagram cannot be inferred when the socket on which the datagram is received is Subscribed to multiple channels. Host operating systems SHOULD provide a way for a host to determine both the source and the destination address to which a datagram was sent. (As one example, the Linux operating system provides the destination of a packet as part of the response to the `recvmsg()` system call.) Until this capability is present, applications may be forced to use higher-layer mechanisms to identify the channel to which a datagram was sent.

4.2. Requirements on the Host IP Module

An incoming datagram destined to an SSM address MUST be delivered by the IP module to all sockets that have indicated (via Subscribe) a desire to receive data that matches the datagram's source address, destination address, and arriving interface. It MUST NOT be delivered to other sockets.

When the first socket on host H subscribes to a channel (S,G) on interface I, the host IP module on H sends a request on interface I to indicate to neighboring routers that the host wishes to receive traffic sent by source S to source-specific multicast destination G. Similarly, when the last socket on a host unsubscribes from a channel on interface I, the host IP module sends an unsubscription request for that channel to interface I.

These requests will typically be Internet Group Management Protocol version 3 (IGMPv3) messages for IPv4, or Multicast Listener Discovery Version 2 (MLDv2) messages for IPv6 [IGMPv3,MLDv2]. A host that supports the SSM service model MUST implement the host portion of [IGMPv3] for IPv4 and [MLDv2] for IPv6. It MUST also conform to the IGMPv3/MLDv2 behavior described in [GMP-SSM].

4.3. Allocation of Source-Specific Multicast Addresses

The SSM destination address 232.0.0.0 is reserved, and it must not be used as a destination address. Similarly, FF3x::4000:0000 is also reserved. The goal of reserving these two addresses is to preserve one invalid SSM destination for IPv4 and IPv6, which can be useful in an implementation as a null value. The address range 232.0.0.1 - 232.0.0.255 is currently reserved for allocation by IANA. SSM destination addresses in the range FF3x::4000:0001 through FF3x::7FFF:FFFF are similarly reserved for IANA allocation [IPv6-MALLOC]. The motivation to reserve these addresses is outlined below in Section 9, "IANA Considerations".

The policy for allocating the rest of the SSM addresses to sending applications is strictly locally determined by the sending host.

When allocating SSM addresses dynamically, a host or host operating system MUST NOT allocate sequentially starting at the first allowed address. It is RECOMMENDED to allocate SSM addresses to applications randomly, while ensuring that allocated addresses are not given simultaneously to multiple applications (and avoiding the reserved addresses). For IPv6, the randomization should apply to the lowest 31 bits of the address.

As described in Section 6, the mapping of an IP packet with SSM destination address onto a link-layer multicast address does not take into account the datagram's source IP address (on commonly-used link layers like Ethernet). If all hosts started at the first allowed address, then with high probability, many source-specific channels on shared-medium local area networks would use the same link-layer multicast address. As a result, traffic destined for one channel subscriber would be delivered to another's IP module, which would then have to discard the datagram.

A host operating system SHOULD provide an interface to allow an application to request a unique allocation of a channel destination address in advance of a session's commencement, and this allocation database SHOULD persist across host reboots. By providing persistent allocations, a host application can advertise the session in advance

of its start time on a web page or in another directory. (We note that this issue is not specific to SSM applications -- the same problem arises for ASM.)

This document neither defines the interfaces for requesting or returning addresses nor specifies the host algorithms for storing those allocations. One plausible abstract API is defined in RFC 2771 [RFC2771]. Note that RFC 2771 allows an application to request an address within a specific range of addresses. If this interface is used, the starting address of the range SHOULD be selected at random by the application.

For IPv6, administratively scoped SSM channel addresses are created by choosing an appropriate scope identifier for the SSM destination address. Normal IPv6 multicast scope boundaries [SCOPINGv6] are applied to traffic sent to an SSM destination address, including any relevant boundaries applied to both the source and destination address.

No globally agreed-upon administratively-scoped address range [ADMIN-SCOPE] is currently defined for IPv4 source-specific multicast. For IPv4, administrative scoping of SSM addresses can be implemented within an administrative domain by filtering outgoing SSM traffic sent to a scoped address at the domain's boundary routers.

5. Router Requirements

5.1. Packet Forwarding

A router that receives an IP datagram with a source-specific destination address MUST silently drop it unless a neighboring host or router has communicated a desire to receive packets sent from the source and to the destination address of the received packet.

5.2. Protocols

Certain IP multicast routing protocols already have the ability to communicate source-specific joins to neighboring routers (in particular, PIM-SM [PIM-SM]), and these protocols can, with slight modifications, be used to provide source-specific semantics. A router that supports the SSM service model MUST implement the PIM-SSM subset of the PIM-SM protocol from [PIM-SM] and MUST implement the router portion of [IGMPv3] for IPv4 and [MLDv2] for IPv6. An SSM router MUST also conform to the IGMPv3/MLDv2 behavior described in [GMP-SSM].

With PIM-SSM, successful establishment of an (S,G) forwarding path from the source S to any receiver depends on hop-by-hop forwarding of the explicit join request from the receiver toward the source. The protocol(s) and algorithms that are used to select the forwarding path for this explicit join must provide a loop-free path. When using PIM-SSM, the PIM-SSM implementation MUST (at least) support the ability to use the unicast topology database for this purpose.

A network can concurrently support SSM in the SSM address range and any-source multicast in the rest of the multicast address space, and it is expected that this will be commonplace. In such a network, a router may receive a non-source-specific, or "(*,G)" in conventional terminology, request for delivery of traffic in the SSM range from a neighbor that does not implement source-specific multicast in a manner compliant with this document. A router that receives such a non-source-specific request for data in the SSM range MUST NOT use the request to establish forwarding state and MUST NOT propagate the request to other neighboring routers. A router MAY log an error in such a case. This applies both to any request received from a host (e.g., an IGMPv1 or IGMPv2 [IGMPv2] host report) and to any request received from a routing protocol (e.g., a PIM-SM (*,G) join). The inter-router case is further discussed in Section 8, "Transition Considerations".

It is essential that all routers in the network give source-specific semantics to the same range of addresses in order to achieve the full benefit of SSM. To comply with this specification, a router MUST treat ALL IANA-allocated SSM addresses with source-specific semantics.

6. Link-Layer Transmission of Datagrams

Source-specific multicast packets are transmitted on link-layer networks as specified in RFC 1112 for IPv4 and as in [ETHERv6] for IPv6. On most shared-medium link-layer networks that support multicast (e.g., Ethernet), the IP source address is not used in the selection of the link-layer destination address. Consequently, on such a network, all packets sent to destination address G will be delivered to any host that has subscribed to any channel (S,G), regardless of S. Therefore, the IP module MUST filter packets it receives from the link layer before delivering them to the socket layer.

7. Security Considerations

This section outlines security issues pertaining to SSM. The following topics are addressed: IPsec, denial-of-service attacks, source spoofing, and security issues related to administrative scoping.

7.1. IPsec and SSM

The IPsec Authentication Header (AH) and Encapsulating Security Payload (ESP) can be used to secure SSM traffic, if a multicast-capable implementation of IPsec (as required in [RFC4301]) is used by the receivers.

7.2. SSM and RFC 2401 IPsec Caveats

For existing implementations of RFC 2401 IPsec (now superseded by [RFC4301]), there are a few caveats related to SSM. They are listed here. In RFC 2401 IPsec, the source address is not used as part of the key in the SAD lookup. As a result, two senders that happen to use the same SSM destination address and the same Security Parameter Index (SPI) will "collide" in the SAD at any host that is receiving both channels. Because the channel addresses and SPIs are both allocated autonomously by the senders, there is no reasonable means to ensure that each sender uses a unique destination address or SPI.

A problem arises if a receiver subscribes simultaneously to two unrelated channels using IPsec whose sources happen to be using the same IP destination address (IPDA) and the same IPsec SPI. Because the channel destination addresses are allocated autonomously by the senders, any two hosts can simultaneously use the same destination address, and there is no reasonable means to ensure that this does not happen. The <IPDA,SPI> tuple, however, consists of 56 bits that are generally randomly chosen (24 bits of the IP destination and 32 bits of the SPI), and a conflict is unlikely to occur through random chance.

If such a collision occurs, a receiver will not be able to simultaneously receive IPsec-protected traffic from the two colliding sources. A receiver can detect this condition by noticing that it is receiving traffic from two different sources with the same SPI and the same SSM destination address.

7.3. Denial of Service

A subscription request creates (S,G) state in a router to record the subscription, invokes processing on that router, and possibly causes processing at neighboring routers. A host can mount a denial-of-service attack by requesting a large number of subscriptions. Denial of service can result if:

- a large amount of traffic arrives when it was otherwise undesired, consuming network resources to deliver it and host resources to drop it;
- a large amount of source-specific multicast state is created in network routers, using router memory and CPU resources to store and process the state; or
- a large amount of control traffic is generated to manage the source-specific state, using router CPU and network bandwidth.

To reduce the damage from such an attack, a router MAY have configuration options to limit, for example, the following items:

- The total rate at which all hosts on any one interface are allowed to initiate subscriptions (to limit the damage caused by forged source-address attacks).
- The total number of subscriptions that can be initiated from any single interface or host.

Any decision by an implementor to artificially limit the rate or number of subscriptions should be taken carefully, however, as future applications may use large numbers of channels. Tight limits on the rate or number of channel subscriptions would inhibit the deployment of such applications.

A router SHOULD verify that the source of a subscription request is a valid address for the interface on which it was received. Failure to do so would exacerbate a spoofed-source address attack.

We note that these attacks are not unique to SSM -- they are also present for any-source multicast.

7.4. Spoofed Source Addresses

By forging the source address in a datagram, an attacker can potentially violate the SSM service model by transmitting datagrams on a channel belonging to another host. Thus, an application requiring strong authentication should not assume that all packets

that arrive on a channel were sent by the requested source without higher-layer authentication mechanisms. The IPSEC Authentication Header [RFC2401, RFC4301] may be used to authenticate the source of an SSM transmission, for instance.

Some degree of protection against spoofed source addresses in multicast is already fairly widespread, because the commonly deployed IP multicast routing protocols [PIM-DM, PIM-SM, DVMRP] incorporate a "reverse-path forwarding check" that validates that a multicast packet arrived on the expected interface for its source address. Routing protocols used for SSM SHOULD incorporate such a check.

Source Routing [RFC791] (both Loose and Strict) in combination with source address spoofing may be used to allow an impostor of the true channel source to inject packets onto an SSM channel. An SSM router SHOULD by default disallow source routing to an SSM destination address. A router MAY have a configuration option to allow source routing. Anti-source spoofing mechanisms, such as source address filtering at the edges of the network, are also strongly encouraged.

7.5. Administrative Scoping

Administrative scoping should not be relied upon as a security measure [ADMIN-SCOPE]; however, in some cases it is part of a security solution. It should be noted that no administrative scoping exists for IPv4 source-specific multicast. An alternative approach is to manually configure traffic filters to create such scoping if necessary.

Furthermore, for IPv6, neither source nor destination address scoping should be used as a security measure. In some currently-deployed IPv6 routers (those that do not conform to [SCOPINGv6]), scope boundaries are not always applied to all source address (for instance, an implementation may filter link-local addresses but nothing else). Such a router may incorrectly forward an SSM channel (S,G) through a scope boundary for S.

8. Transition Considerations

A host that complies with this document will send ONLY source-specific host reports for addresses in the SSM range. As stated above, a router that receives a non-source-specific (e.g., IGMPv1 or IGMPv2 or MLDv1 [RFC2710]) host report for a source-specific multicast destination address MUST ignore these reports. Failure to do so would violate the SSM service model promised to the sender: that a packet sent to (S,G) would only be delivered to hosts that specifically requested delivery of packets sent to G by S.

During a transition period, it would be possible to deliver SSM datagrams in a domain where the routers do not support SSM semantics by simply forwarding any packet destined to G to all hosts that have requested subscription of (S,G) for any S. However, this implementation risks unduly burdening the network infrastructure by delivering (S,G) datagrams to hosts that did not request them. Such an implementation for addresses in the SSM range is specifically not compliant with Section 5.2 of this document.

9. IANA Considerations

IANA allocates IPv4 addresses in the range 232.0.0.1 through 232.0.0.255 and IPv6 addresses in the range FF3x:4000:0001 to FF3x::7FFF:FFFF. These addresses are allocated according to IETF Consensus [IANA-CONSID]. These address ranges are reserved for services with wide applicability that either require that or would strongly benefit if all hosts use a well-known SSM destination address for that service. Any proposal for allocation must consider the fact that, on an Ethernet network, all datagrams sent to any SSM destination address will be transmitted with the same link-layer destination address, regardless of the source. Furthermore, the fact that SSM destinations in 232.0.0.0/24 and 232.128.0.0/24 use the same link-layer addresses as the reserved IP multicast group range 224.0.0.0/24 must also be considered. Similar consideration should be given to the IPv6 reserved multicast addresses. 232.0.0.0 and FF3x::4000:0000 should not be allocated, as suggested above.

Except for the aforementioned addresses, IANA SHALL NOT allocate any SSM destination address to a particular entity or application. To do so would compromise one of the important benefits of the source-specific model: the ability for a host to simply and autonomously allocate a source-specific multicast address from a large flat address space.

10. Acknowledgements

The SSM service model draws on a variety of prior work on alternative approaches to IP multicast, including the EXPRESS multicast model of Holbrook and Cheriton [EXPRESS], Green's [SMRP], and the Simple Multicast proposal of Perlman, et al. [SIMPLE]. We would also like to thank Jon Postel and David Cheriton for their support in reassigning the 232/8 address range to SSM. Brian Haberman contributed to the IPv6 portion of this document. Thanks to Pekka Savola for a careful review.

11. Normative References

- [ETHERv6] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [GMP-SSM] Holbrook, H. and B. Cain, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", RFC 4604, August 2006.
- [IGMPv3] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [IPv6-UBM] Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6 Multicast Addresses", RFC 3306, August 2002.
- [IPv6-MALLOC] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.
- [MLDv2] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [PIM-SM] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas. "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

12. Informative References

- [ADMIN-SCOPE] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, July 1998.
- [DVMRP] Waitzman, D., Partridge, C., and S. Deering, "Distance Vector Multicast Routing Protocol", RFC 1075, November 1988.
- [EXPRESS] Holbrook, H., and Cheriton, D. "Explicitly Requested Source-Specific Multicast: EXPRESS support for Large-scale Single-source Applications." Proceedings of ACM SIGCOMM '99, Cambridge, MA, September 1999.
- [IANA-ALLOC] Internet Assigned Numbers Authority, <http://www.iana.org/assignments/multicast-addresses>.
- [IANA-CONSID] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [IGMPv2] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [MSFAPI] Thaler, D., Fenner, B., and B. Quinn, "Socket Interface Extensions for Multicast Source Filters", RFC 3678, January 2004.
- [PIM-DM] Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", RFC 3973, January 2005.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.
- [RFC2771] Finlayson, R., "An Abstract API for Multicast Address Allocation", RFC 2771, February 2000.
- [SCOPINGv6] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, March 2005.

- [SIMPLE] R. Perlman, C-Y. Lee, A. Ballardie, J. Crowcroft, Z. Wang, T. Maufer, C. Diot, and M. Green, "Simple Multicast: A Design for Simple, Low-Overhead Multicast", Work in Progress, October 1999.
- [SMRP] Green, M. "Method and System of Multicast Routing for Groups with a Single Transmitter." United States Patent Number 5,517,494.

Authors' Addresses

Brad Cain
Acopia Networks

EMail: bcain99@gmail.com

Hugh Holbrook
Arastra, Inc.
P.O. Box 10905
Palo Alto, CA 94303

Phone: +1 650 331-1620
EMail: holbrook@arastra.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

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

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

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

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

