Expires May 3, 2003

H. Holbrook
Cisco Systems
B. Cain
Cereva Networks

3 Nov 2002

Source-Specific Multicast for IP
<draft-ietf-ssm-arch-01.txt>

Status of this Memo

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 [RFC 2119].

Abstract

IP 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 [IANA-ALLOCATION]. For IP version 6 (IPv6), the address prefix FF3x::/32 is reserved for Source-Specific Multicast use, where 'x' is any valid scope identifier [IPV6-UBM]. This document defines the semantics of source-specific multicast addresses and specifies the policies governing their use. It 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.

A companion document will describe how the Internet Group Management Protocol Version 3 [IGMPv3] and the Multicast Listener Discovery Protocol Version 2 [MLDv2] are adapted to support source-specific multicast.

1. Overview and Rationale

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) G 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 the RFC 1112 model of multicast. RFC 2373 [RFC2373] specifies the form of IPv6 multicast addresses with ASM semantics.

IP 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-ALLOCATION].

For IPv6, the address prefix FF3x::/32 is reserved for Source-Specific Multicast use, where 'x' is any valid scope identifier [IPV6-UBM]. Addresses in the range FF3x::4000:0000 through FF3x::7fff:ffff are reserved for allocation by IANA, and 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, per [IPV6-UBM]. 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 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.

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 from a host.

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

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 = (36.18.0.1, 232.7.8.9)
```

differs from

```
    S,G = (36.18.0.2, 232.7.8.9),
```

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

```
   S,G = (2001:3618::1, FF23::1234)
```

and

S,G = (2001:3618::2, FF23::1234)

are different channels.

3. Terminology

To avoid 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], the Subscribe and Unsubscribe interfaces may be supported via that API.

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 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 out interface I.

These requests will typically be IGMPv3 messages for IPv4, or MLDv2 messages for IPv6. The exact rules for sending source-specific subscription and unsubscription requests and the algorithms used to maintain subscriptions are defined in other documents.

4.3. Allocation of Source-Specific Multicast Addresses

The SSM destination address 232.0.0.0 is reserved, and hosts MUST NOT send datagrams with destination address of 232.0.0.0. The address range 232.0.0.1-232.0.0.255 is currently reserved for allocation by IANA. The IPv6 SSM address range FF2x:: is reserved for IANA allocation.

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 address range for IPv4). For IPv6, the randomization should apply to the lower 32 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 collide on 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 reject 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.

No globally agreed-upon administratively-scoped address range [ADMIN-SCOPE] is currently defined for source-specific multicast. Note that there is no possibility of address conflict between hosts in different administrative domains (or between two hosts of any kind). Administrative scoping of SSM addresses can be implemented within an administrative domain by filtering at domain 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), and these protocols can, with slight modifications, be used to provide source-specific semantics. Companion documents will specify the required modifications to those protocols to support the source-specific address range.

A network can concurrently support SSM semantics 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. This applies both to any request received from a host, e.g., an IGMPv1 or IGMPv2 host report, and to any request received from a routing protocol, e.g., a PIM-SM (*,G) join [PIM-SM]. 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 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. And therefore, the IP module MUST filter packets it receives from the link layer before delivering them to the socket layer. A socket on which an (S,G) subscription has been requested MUST NOT receive packets whose source and destination address do not match the requested subscription(s) for that socket.

7. Security Considerations

7.1. Denial-of-Service

A subscription request creates (S,G) state in a router to record the subscription and 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. A 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

- 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 a configuration option to limit 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 initated 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.

7.2. 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 [IPSEC] 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 MUST have a configuration option to disable source routing to an SSM destination addresses, and the default value SHOULD be to disable Source Routing to an SSM destination address. Anti-source spoofing mechanisms like source address filtering at the edges of the network are also strongly encouraged.

8. Transition Considerations

A host that complies with this document will send ONLY source-specific host reports for addresses in the SSM range. A router that receives a non-source-specific (IGMPv1 or IGMPv2) host report for a source-specific destination addresses SHOULD 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 deliver (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

Addresses in the range 232.0.0.1 through 232.0.0.255 and IPv6 addresses with prefix FF2x:: are reserved for services with wide applicability that either require or would strongly benefit if all hosts used a well-known SSM destination address for that service. IANA shall allocate addresses in this range according to IETF Consensus [IANA-CONSIDERATIONS]. 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.

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 address from a large flat address space.

10. Acknowledgments

The SSM service model draws on a variety of prior work on alternative aproaches 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.

11. References

11.1. Normative

[RFC791] Postel, J., ed., "Internet Protocol, Darpa Internet Program Protocol Specification," September 1981.

[IPV6-UBM] B. Haberman, D. Thaler, "Unicast-Prefix-based IPv6 Multicast Addresses.", RFC 3306, August 2002.

[IPV6-MALLOC] B. Haberman, "Dynamic Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.

[ETHERv6] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC2464, Dec 1998.

[RFC1112] Deering, S., "Host Extensions for IP Multicasting," RFC 1112, August 1989.

[RFC2373] Hinden, R. and Deering, S. "IP Version 6 Addressing Architecture." RFC 2373, July 1998.

11.2. Non-normative

[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, Nov 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-ALLOCATION] Internet Assigned Numbers Authority, http://www.isi.edu/in-notes/iana/assignments/multicast-addresses.

[IANA-CONSIDERATIONS] Narten, T., and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs," RFC 2434, October 1998.

[IGMPv2] Fenner, W., "Internet Group Management Protocol, Version 2," RFC 2236, November 1997.

[IGMPv3] Cain, B., Deering, S., and A. Thyagarajan, "Internet Group Management Protocol, Version 3," RFC 3376, October 2002.

[IPSEC] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol.", RFC 2401.

[MLDv2] R. Vida, L. Costa, R. Zara, S. Fdida, S. Deering, B. Fenner, I. Kouvelas, B. Haberman. "Multicast Listener Discovery Version 2 (MLDv2) for IPv6." Work in Progress.

[MSFAPI] Thaler, D., Fenner, B., and Quinn, B. "Socket Interface Extensions for Multicast Source Filters." Work in Progress.

[PIM-SM] Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., Jacobson, V., Liu, C., Sharma, P. and L. Wei, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification," RFC 2362, June 1998.

[PIM-DM] Deering, S., Estrin, D., Farinacci, D., Jacobson, V., Helmy, A., Meyer, D., and L. Wei, "Protocol Independent Multicast Version 2 Dense Mode Specification," Work in Progress.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," RFC 2119, March 1997.

[RFC2710] S. Deering, W. Fenner, 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.

[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.

[SMRP] Green, M. "Method and System of Multicast Routing for Groups with a Single Transmitter." United States Patent Number 5,517,494.

12. Author's Address

Brad Cain
Cereva Networks
3 Network Drive
Marlborough, MA 01752
bcain@cereva.com

Hugh Holbrook
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
holbrook@cisco.com

This document expires May 3, 2003.