

## An Overview of Source-Specific Multicast (SSM)

### Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

### Abstract

The purpose of this document is to provide an overview of Source-Specific Multicast (SSM) and issues related to its deployment. It discusses how the SSM service model addresses the challenges faced in inter-domain multicast deployment, changes needed to routing protocols and applications to deploy SSM and interoperability issues with current multicast service models.

### 1. Introduction

This document provides an overview of the Source-Specific Multicast (SSM) service and its deployment using the PIM-SM and IGMP/MLD protocols. The network layer service provided by SSM is a "channel", identified by an SSM destination IP address (G) and a source IP address S. An IPv4 address range has been reserved by IANA for use by the SSM service. An SSM destination address range already exists for IPv6. A source S transmits IP datagrams to an SSM destination address G. A receiver can receive these datagrams by subscribing to the channel (Source, Group) or (S,G). Channel subscription is supported by version 3 of the IGMP protocol for IPv4 and version2 of the MLD protocol for IPv6. The interdomain tree for forwarding IP multicast datagrams is rooted at the source S, and is constructed using the PIM Sparse Mode [9] protocol.

This document is not intended to be a standard for Source-Specific Multicast (SSM). Instead, its goal is to serve as an introduction to SSM and its benefits for anyone interested in deploying SSM services. It provides an overview of SSM and how it solves a number of problems faced in the deployment of inter-domain multicast. It outlines changes to protocols and applications both at end-hosts and routers

for supporting SSM, with pointers to more detailed documents where appropriate. Issues of interoperability with the multicast service model defined by RFC 1112 are also discussed.

This memo is a product of the Source-Specific Multicast (SSM) Working Group of the Internet Engineering Task Force.

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as defined in BCP 14, RFC 2119 [28].

## 2. Terminology

This section defines some terms that are used in the rest of this document:

**Any-Source Multicast (ASM):** This is the IP multicast service model defined in RFC 1112 [25]. An IP datagram is transmitted to a "host group", a set of zero or more end-hosts (or routers) identified by a single IP destination address (224.0.0.0 through 239.255.255.255 for IPv4). End-hosts may join and leave the group any time, and there is no restriction on their location or number. Moreover, this model supports multicast groups with arbitrarily many senders - any end-host (or router) may transmit to a host group, even if it is not a member of that group.

**Source-Specific Multicast (SSM):** This is the multicast service model defined in [5]. An IP datagram is transmitted by a source S to an SSM destination address G, and receivers can receive this datagram by subscribing to channel (S,G). SSM provides host applications with a "channel" abstraction, in which each channel has exactly one source and any number of receivers. SSM is derived from earlier work in EXPRESS [1]. The address range 232/8 has been assigned by IANA for SSM service in IPv4. For IPv6, the range FF3x::/96 is defined for SSM services [21].

**Source-Filtered Multicast (SFM):** This is a variant of the ASM service model, and uses the same address range as ASM (224.0.0.0-239.255.255.255). It extends the ASM service model as follows. Each "upper layer protocol module" can now request data sent to a host group G by only a specific set of sources, or can request data sent to host group G from all BUT a specific set of sources. Support for source filtering is provided by version 3 of the Internet Group Management Protocol (or IGMPv3) [3] for IPv4, and version 2 of the Multicast Listener Discovery (or MLDv2) [22] protocol for IPv6. We shall henceforth refer to these two protocols as "SFM-capable". Earlier versions of these protocols - IGMPv1/IGMPv2 and MLDv1 - do not provide support for

source-filtering, and are referred to as "non-SFM-capable". Note that while SFM is a different model than ASM from a receiver standpoint, there is no distinction between the two for a sender.

For the purpose of this document, we treat the scoped multicast model of [12] to be a variant of ASM since it does not explicitly restrict the number of sources, but only requires that they be located within the scope zone of the group.

### 3. The IGMP/PIM-SM/MSDP/MBGP Protocol Suite for ASM

As of this writing, all multicast-capable networks support the ASM service model. One of the most common multicast protocol suites for supporting ASM consists of IGMP version 2 [2], PIM-SM [8,9], MSDP [13] and MBGP [26]. IGMPv2 is the most commonly used protocol for hosts to specify membership in a multicast group, and nearly all multicast routers support (at least) IGMPv2. In case of IPv6, MLDv1 [21] is the commonly used protocol.

Although a number of protocols such as PIM-DM [10], CBT [24,11], DVMRP [6], etc. exist for building multicast tree among all receivers and sources in the same administrative domain, PIM-SM [8,9] is the most widely used protocol. PIM-SM builds a spanning multicast tree rooted at a core rendezvous point or RP for all group members within a single administrative domain. A 'first-hop' router adjacent to a multicast source sends the source's traffic to the RP for its domain. The RP forwards the data down the shared spanning tree to all interested receivers within the domain. PIM-SM also allows receivers to switch to a source-based shortest path tree.

As of this writing, multicast end-hosts with SFM capabilities are not widely available. Hence a client can only specify interest in an entire host group and receives data sent from any source to this group.

Inter-domain multicast service (i.e., where sources and receivers are located in different domains) requires additional protocols - MSDP [13] and MBGP [26] are the most commonly used ones. An RP uses the MSDP protocol to announce multicast sources to RPs in other domains. When an RP discovers a source in a different domain transmitting data to a multicast group for which there are interested receivers in its own domain, it joins the shortest-path source based tree rooted at that source. It then redistributes the data received to all interested receivers via the intra-domain shared tree rooted at itself.

MBGP defines extensions to the BGP protocol to support the advertisement of reachability information for multicast routes. This allows an autonomous system (AS) to support incongruent unicast and multicast routing topologies, and thus implement separate routing policies for each.

However, the last-hop routers of interested receivers may eventually switch to a shortest-path tree rooted at the source that is transmitting the data.

#### 4. Problems with Current Architecture

There are several deployment problems associated with current multicast architecture:

##### A) Address Allocation:

Address allocation is one of core deployment challenges posed by the ASM service model. The current multicast architecture does not provide a deployable solution to prevent address collisions among multiple applications. The problem is much less serious for IPv6 than for IPv4 since the size of the multicast address space is much larger. A static address allocation scheme, GLOP [17] has been proposed as an interim solution for IPv4; however, GLOP addresses are allocated per registered AS, which is inadequate in cases where the number of sources exceeds the AS numbers available for mapping. RFC 3138 expands on RFC 2770 to allow routing registries to assign multicast addresses from the GLOP space corresponding to the RFC 1930 private AS space [27]. This space is referred to as the EGLOP (Extended GLOP) address space. Proposed longer-term solutions such as the Multicast Address Allocation Architecture [14] are generally perceived as being too complex (with respect to the dynamic nature of multicast address allocation) for widespread deployment.

##### B) Lack of Access control:

In the ASM service model, a receiver cannot specify which specific sources it would like to receive when it joins a given group. A receiver will be forwarded data sent to a host group by any source. Moreover, even when a source is allocated a multicast group address to transmit on, it has no way of enforcing that no other source will use the same address. This is true even in the case of IPv6, where address collisions are less likely due to the much larger size of the address space.

C) Inefficient handling of well-known sources:

In cases where the address of the source is well known in advance of the receiver joining the group, and when the shortest forwarding path is the preferred forwarding mode, then shared tree mechanisms are not necessary.

5. Source Specific Multicast (SSM): Benefits and Requirements

As mentioned before, the Source Specific Multicast (SSM) service model defines a "channel" identified by an (S,G) pair, where S is a source address and G is an SSM destination address. Channel subscriptions are described using an SFM-capable group management protocol such as IGMPv3 or MLDv2. Only source-based forwarding trees are needed to implement this model.

The SSM service model alleviates all of the deployment problems described earlier:

- A) Address Allocation: SSM defines channels on a per-source basis, i.e., the channel (S1,G) is distinct from the channel (S2,G), where S1 and S2 are source addresses, and G is an SSM destination address. This averts the problem of global allocation of SSM destination addresses, and makes each source independently responsible for resolving address collisions for the various channels that it creates.
- B) Access Control: SSM lends itself to an elegant solution to the access control problem. When a receiver subscribes to an (S,G) channel, it receives data sent only by the source S. In contrast, any host can transmit to an ASM host group. At the same time, when a sender picks a channel (S,G) to transmit on, it is automatically ensured that no other sender will be transmitting on the same channel (except in the case of malicious acts such as address spoofing). This makes it much harder to "spam" an SSM channel than an ASM multicast group.
- C) Handling of well-known sources: SSM requires only source-based forwarding trees; this eliminates the need for a shared tree infrastructure. This implies that neither the RP-based shared tree infrastructure of PIM-SM nor the MSDP protocol is required. Thus the complexity of the multicast routing infrastructure for SSM is low, making it viable for immediate deployment. Note that there is no difference in how MBGP is used for ASM and SSM.

6. SSM Framework

Figure 1 illustrates the elements in an end-to-end implementation framework for SSM:

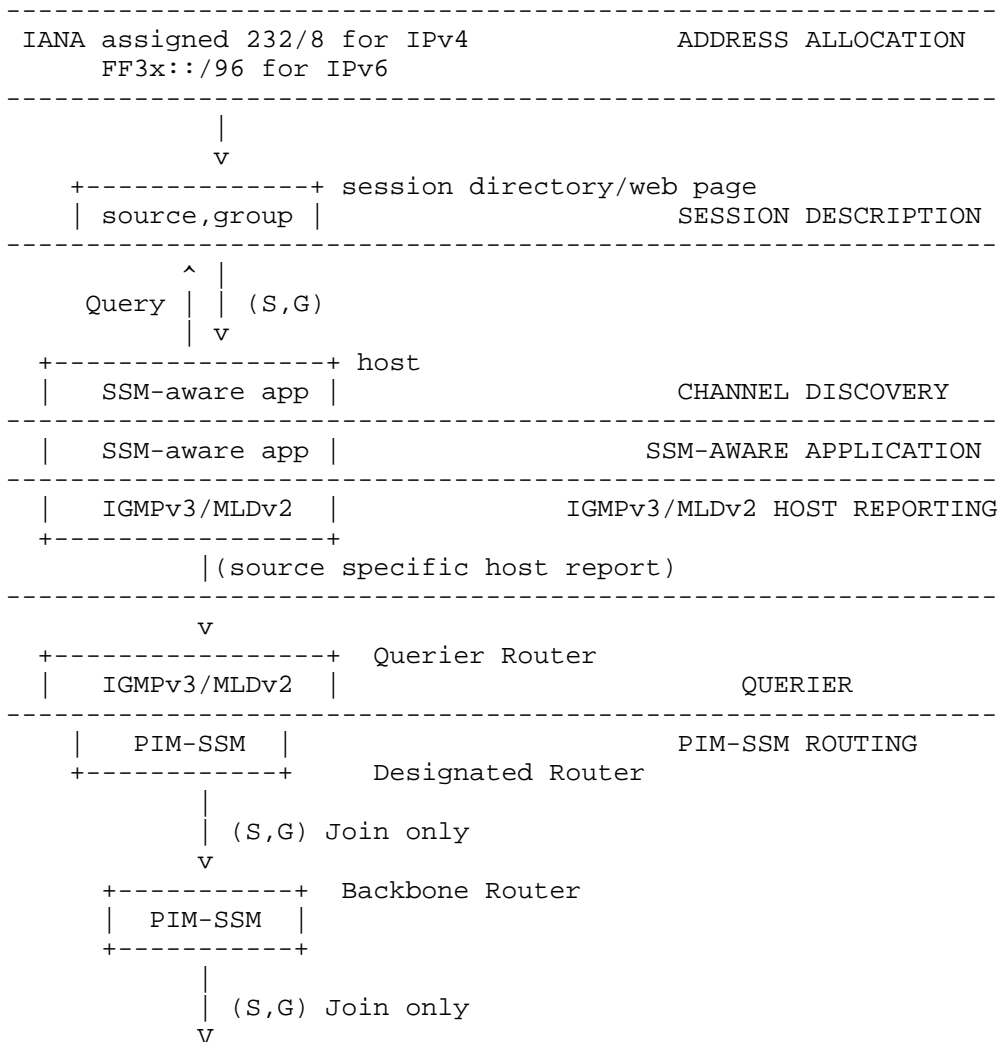


Figure 1: SSM Framework: elements in end-to-end model

We now discuss the framework elements in detail:

### 6.1. Address Allocation

For IPv4, the address range of 232/8 has been assigned by IANA for SSM. To ensure global SSM functionality in 232/8, including in networks where routers run non-SFM-capable protocols, operational policies are being proposed [9] which recommend that routers should not send SSM traffic to parts of the network that do not have channel subscribers.

Note that IGMPv3/MLDv2 does not limit (S,G) joins to only the 232/8 range. However, SSM service, as defined in [5], is available only in this address range for IPv4.

In case of IPv6, [23] has defined an extension to the addressing architecture to allow for unicast prefix-based multicast addresses. See RFC 3306 for details.

### 6.2. Session Description and Channel Discovery

An SSM receiver application must know both the SSM destination address G and the source address S before subscribing to a channel. Channel discovery is the responsibility of applications. This information can be made available in a number of ways, including via web pages, sessions announcement applications, etc. This is similar to what is used for ASM applications where a multicast session needs to be announced so that potential subscribers can know of the multicast group address, encoding schemes used, etc. In fact, the only additional piece of information that needs to be announced is the source address for the channel being advertised. However, the exact mechanisms for doing this is outside the scope of this framework document.

### 6.3. SSM-Aware Applications

There are two main issues in making multicast applications "SSM-aware":

- An application that wants to receive an SSM session must first discover the channel address in use.
- A receiving application must be able to specify both a source address and a destination address to the network layer protocol module on the end-host.

Specific API requirements are identified in [16]. [16] describes a recommended application programming interface for a host operating system to support the SFM service model. Although it is intended for SFM, a subset of this interface is sufficient for supporting SSM.

#### 6.4. IGMPv3/MLDv2 Host Reporting and Querier

In order to use SSM service, an end-host must be able to specify a channel address, consisting of a source's unicast address and an SSM destination address. IGMP version 2 [3] and MLD version 1 [19] allows an end-host to specify only a destination multicast address. The ability to specify an SSM channel address *c* is provided by IGMP version 3 [3] and MLD version 2 [20]. These protocols support "source filtering", i.e., the ability of an end-system to express interest in receiving data packets sent only by SPECIFIC sources, or from ALL BUT some specific sources. In fact, IGMPv3 provides a superset of the capabilities required to realize the SSM service model.

A detailed discussion of the use of IGMPv3 in the SSM destination address range is provided in [4].

The Multicast Listener Discovery (MLD) protocol used by an IPv6 router to discover the presence of multicast listeners on its directly attached links, and to discover the multicast addresses that are of interest to those neighboring nodes. MLD version 1 is derived from IGMPv2 and does not provide the source filtering capability required for the SSM service model. MLD version 2 is derived from, and provides the same support for source-filtering as, IGMPv3. Thus IGMPv3 (or MLDv2 for IPv6) provides a host with the ability to request the network for an SSM channel subscription.

#### 6.5. PIM-SSM Routing

[9] provides guidelines for how a PIM-SM implementation should handle source-specific host reports as required by SSM. Earlier versions of the PIM protocol specifications did not describe how to do this.

The router requirements for operation in the SSM range are detailed in [5]. These rules are primarily concerned with preventing ASM-style behaviour in the SSM address range. In order to comply with [5] several changes to the PIM-SM protocol are required, as described in [9]. The most important changes in PIM-SM required for compliance with [5] are:



- When a DR receives an (S,G) join request with the address G in the SSM address range, it MUST initiate a (S,G) join, and NEVER a (\*,G) join.
- Backbone routers (i.e., routers that do not have directly attached hosts) MUST NOT propagate (\*,G) joins for group addresses in the SSM address range.
- Rendezvous Points (RPs) MUST NOT accept PIM Register messages or (\*,G) Join messages in the SSM address range.

Note that only a small subset of the full PIM-SM protocol functionality is needed to support the SSM service model. This subset is explicitly documented in [9].

## 7. Interoperability with Existing Multicast Service Models

Interoperability with ASM is one of the most important issues in moving to SSM deployment, since both models are expected to be used at least in the foreseeable future. SSM is the ONLY service model for the SSM address range - the correct protocol behaviour for this range is specified in [5]. The ASM service model will be offered for the non-SSM address range, where receivers can issue (\*,G) join requests to receive multicast data. A receiver is also allowed to issue an (S,G) join request in the non-SSM address range; however, in that case there is no guarantee that it will receive service according to the SSM model.

Another interoperability issue concerns the MSDP protocol, which is used between PIM-SM rendezvous points (RPs) to discover multicast sources across multiple domains. MSDP is not needed for SSM, but is needed if ASM is supported. [9] specifies operational recommendations to help ensure that MSDP does not interfere with the ability of a network to support the SSM service model. Specifically, [9] states that RPs must not accept, originate or forward MSDP SA messages for the SSM address range.

## 8. Security Considerations

SSM does not introduce new security considerations for IP multicast. It can help in preventing denial-of-service attacks resulting from unwanted sources transmitting data to a multicast channel (S, G). However no guarantee is provided.

## 9. Acknowledgments

We would like to thank Gene Bowen, Ed Kress, Bryan Lyles, Timothy Roscoe, Hugh Holbrook, Isidor Kouvelas, Tony Speakman and Nidhi Bhaskar for participating in lengthy discussions and design work on SSM, and providing feedback on this document. Thanks are also due to Mujahid Khan, Ted Seely, Tom Pusateri, Bill Fenner, Kevin Almeroth, Brian Levine, Brad Cain, Hugh LaMaster and Pekka Savola for their valuable insights and continuing support.

## 10. References

### 10.1. Informative References

- [1] Holbrook, H. and D.R. Cheriton, "IP Multicast Channels: EXPRESS Support for Large-scale Single-Source Applications", In Proceedings of SIGCOMM 1999.
- [2] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [3] Cain, B., Deering, S., Kouvelas, I. and A. Thyagarajan, "Internet Group Management Protocol, Version 3.", RFC 3376, October 2002.
- [4] Holbrook, H. and B. Cain, "Using IGMPv3 and MLDv2 for Source-Specific Multicast", Work In Progress.
- [5] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", Work in Progress.
- [6] Deering, S. and D. Cheriton, "Multicast Routing in Datagram Networks and Extended LANs", ACM Transactions on Computer Systems, 8(2):85-110, May 1990.
- [7] Deering, S. et al., "PIM Architecture for Wide-Area Multicast Routing", IEEE/ACM Transaction on Networking, pages 153-162, April 1996.
- [8] 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.
- [9] Fenner, B., Handley, M., Holbrook, H. and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", Work In Progress.

- [10] Adams, A., Nicholas, J. and W. Siadek, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", Work In Progress.
- [11] Ballardie, A., "Core-Based Trees (CBT) Multicast Routing Architecture", RFC 2201, September 1997.
- [12] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, July 1998.
- [13] Farinacci, D. et al., "Multicast Source Discovery Protocol", Work In Progress.
- [14] Thaler, D., Handley, M. and D. Estrin, "The Internet Multicast Address Allocation Architecture", RFC 2908, September 2000.
- [15] Diot, C., Levine, B., Lyles, B., Kassem, H. and D. Balensiefen, "Deployment Issues for the IP Multicast Service and Architecture", In IEEE Networks Magazine's Special Issue on Multicast, January, 2000.
- [16] Thaler, D., Fenner B. and B. Quinn, "Socket Interface Extensions for Multicast Source Filters", Work in Progress.
- [17] Meyer, D. and P. Lothberg, "GLOP Addressing in 233/8", BCP 53, RFC 3180, September 2001.
- [18] Levine, B. et al., "Consideration of Receiver Interest for IP Multicast Delivery", In Proceedings of IEEE Infocom, March 2000.
- [19] Deering, S., Fenner, W. and B. Haberman, "Multicast Listener Discovery for IPv6", RFC 2710, October 1999.
- [20] Vida, R. et. al., "Multicast Listener Discovery Version 2(MLDv2) for IPv6", Work In Progress.
- [21] Haberman, B. and D. Thaler, "Unicast-Prefix-Based IPv6 Multicast Addresses", RFC 3306, August 1992.
- [22] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [23] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.

- [24] Ballardie, A., "Core-Based Trees (CBT Version 2) Multicast Routing -- Protocol Specification", RFC 2189, September 2001.
- [25] Deering, S., "Host Extensions for IP Multicasting", STD 5, RFC 1112, August 1989.
- [26] Bates, T., Rekhter, Y., Chandra, R. and D. Katz, "Multiprotocol Extensions for BGP-4", RFC 2858, June 2000.
- [27] Meyer, D., "Extended Assignments in 233/8", RFC 3138, June 2001.

## 10.2. Normative References

- [28] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

## 11. Contributors

Christophe Diot  
Intel  
EMail: christophe.diot@intel.com

Leonard Giuliano  
Juniper Networks  
EMail: lenny@juniper.net

Greg Shepherd  
Procket Networks  
EMail: shep@procket.com

Robert Rockell  
Sprint  
EMail: rrockell@sprint.net

David Meyer  
Sprint  
EMail: dmm@1-4-5.net

John Meylor  
Cisco Systems  
EMail: jmeylor@cisco.com

Brian Haberman  
Caspian Networks  
EMail: bkhabs@nc.rr.com

12. Editor's Address

Supratik Bhattacharyya  
Sprint

EMail: [supratik@sprintlabs.com](mailto:supratik@sprintlabs.com)

### 13. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

### Acknowledgement

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