

Network Working Group
Request for Comments: 4608
BCP: 120
Category: Best Current Practice

D. Meyer
R. Rockell
G. Shepherd
August 2006

Source-Specific Protocol Independent Multicast in 232/8

Status of This Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

IP Multicast group addresses in the 232/8 (232.0.0.0 to 232.255.255.255) range are designated as source-specific multicast destination addresses and are reserved for use by source-specific multicast applications and protocols. This document defines operational recommendations to ensure source-specific behavior within the 232/8 range.

Table of Contents

1. Introduction	2
1.1. BCP, Experimental Protocols, and Normative References	2
2. Operational practices in 232/8	4
2.1. Preventing Local Sources from Sending to Shared Tree	4
2.2. Preventing Remote Sources from Being Learned/Joined via MSDP	4
2.3. Preventing Receivers from Joining the Shared Tree	4
2.4. Preventing RPs as Candidates for 232/8	5
3. Acknowledgements	5
4. Security Considerations	5
5. References	6
5.1. Normative References	6
5.2. Informative References	6

1. Introduction

Current Protocol Independent Multicast - Sparse Mode (PIM-SM) [RFC4601] relies on the shared Rendezvous Point (RP) tree to learn about active sources for a group and to support group-generic (Any Source Multicast or ASM) data distribution. The IP Multicast group address range 232/8 has been designated for Source-Specific Multicast [RFC3569] applications and protocols [IANA] and SHOULD support source-only trees only, precluding the requirement of an RP and a shared tree; active sources in the 232/8 range will be discovered out of band. PIM Sparse Mode Designated Routers (DR) with local membership are capable of joining the shortest path tree for the source directly using SSM functionality of PIM-SM.

Operational best common practices in the 232/8 group address range are necessary to ensure shortest path source-only trees across multiple domains in the Internet [RFC3569], and to prevent data from sources sending to groups in the 232/8 range from arriving via shared trees. This avoids unwanted data arrival and allows several sources to use the same group address without conflict at the receivers.

The operational practices SHOULD:

- o Prevent local sources from sending to shared tree
- o Prevent receivers from joining the shared tree
- o Prevent RPs as candidates for 232/8
- o Prevent remote sources from being learned/joined via Multicast Source Discovery Protocol (MSDP) [RFC3618]

1.1. BCP, Experimental Protocols, and Normative References

This document describes the best current practice for a widely deployed Experimental protocol, MSDP. There is no plan to advance MSDP's status (for example, to Proposed Standard). The reasons for this include:

- o MSDP was originally envisioned as a temporary protocol to be supplanted by whatever the Inter-Domain Multicast Routing (IDMR) working group produced as an inter-domain protocol. However, the IDMR WG (or subsequently, the Border Gateway Multicast Protocol (BGMP) WG) never produced a protocol that could be deployed to replace MSDP.

- o One of the primary reasons given for MSDP to be classified as Experimental was that the MSDP Working Group came up with modifications to the protocol that the WG thought made it better but that implementors didn't see any reasons to deploy. Without these modifications (e.g., UDP or GRE encapsulation), MSDP can have negative consequences to initial packets in datagram streams.
- o Scalability: Although we don't know what the hard limits might be, readvertising everything you know every 60 seconds clearly limits the amount of state you can advertise.
- o MSDP reached nearly ubiquitous deployment as the de facto standard inter-domain multicast protocol in the IPv4 Internet.
- o No consensus could be reached regarding the reworking of MSDP to address the many concerns of various constituencies within the IETF. As a result, a decision was taken to document what is (ubiquitously) deployed and to move that document to Experimental. Although advancement of MSDP to Proposed Standard was considered, for the reasons mentioned above, it was immediately discarded.
- o The advent of protocols such as source-specific multicast and bi-directional PIM, as well as embedded RP techniques for IPv6, have further reduced consensus that a replacement protocol for MSDP for the IPv4 Internet is required.

The RFC Editor's policy regarding references is that they be split into two categories known as "normative" and "informative". Normative references specify those documents that must be read for one to understand or implement the technology in an RFC (or whose technology must be present for the technology in the new RFC to work) [RFCED]. In order to understand this document, one must also understand both the PIM [RFC4601] and MSDP [RFC3618] documents. As a result, references to these documents are normative.

The IETF has adopted the policy that BCPs must not have normative references to Experimental protocols. However, this document is a special case in that the underlying Experimental document (MSDP) is not planned to be advanced to Proposed Standard.

The MBONED Working Group requests approval under the Variance Procedure as documented in RFC 2026 [RFC2026]. The IESG followed the Variance Procedure and, after an additional 4-week IETF Last Call, evaluated the comments and status and has approved the document.

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].

2. Operational practices in 232/8

2.1. Preventing Local Sources from Sending to Shared Tree

In order to eliminate the use of shared trees for groups in 232/8, while maintaining coexistence with ASM in PIM-SM, the behavior of the RP and/or the DR needs to be modified. This can be accomplished by

- preventing data for 232/8 groups from being sent encapsulated to the RP by the DR,
- preventing the RP from accepting registers for 232/8 groups from the DR, and
- preventing the RP from forwarding accepted data down (*,G) tree for 232/8 groups.

2.2. Preventing Remote Sources from Being Learned/Joined via MSDP

SSM does not require active source announcements via MSDP. All source announcements are received out of band, and the last hop router is responsible for sending (S,G) joins directly to the source. To prevent propagation of SAs in the 232/8 range, an RP SHOULD

- never originate an SA for any 232/8 groups, and
- never accept or forward an SA for any 232/8 groups.

2.3. Preventing Receivers from Joining the Shared Tree

Local PIM domain practices need to be enforced to prevent local receivers from joining the shared tree for 232/8 groups. This can be accomplished by

- preventing DR from sending (*,G) joins for 232/8 groups, and
- preventing RP from accepting (*,G) join for 232/8 groups.

However, within a local PIM domain, any last-hop router NOT preventing (*,G) joins may trigger unwanted (*,G) state toward the RP that intersects an existing (S,G) tree, allowing the receiver on the shared tree to receive the data, which breaks the source-specific

[RFC3569] service model. It is therefore recommended that ALL routers in the domain MUST reject AND never originate (*,G) joins for 232/8 groups.

In those cases in which an ISP is offering its customers (or others) the use of the ISP's RP, the ISP SHOULD NOT allow (*,G) joins in the 232/8 range.

2.4. Preventing RPs as Candidates for 232/8

Because SSM does not require an RP, all RPs SHOULD NOT offer themselves as candidates in the 232/8 range. This can be accomplished by

- preventing RP/BSR from announcing in the 232/8 range,
- preventing ALL routers from accepting RP delegations in the 232/8 range, and
- precluding RP functionality on RP for the 232/8 range.

Note that in typical practice, RPs announce themselves as candidates for the 224/4 (which obviously includes 232/8). It is still acceptable to allow the advertisement of 224/4 (or any other superset of 232/8); however, this approach relies on the second point, above; namely, that routers silently ignore the RP delegation in the 232/8 range and prevent sending or receiving using the shared tree, as described previously. Finally, an RP SHOULD NOT be configured as a candidate RP for 232/8 (or for a more specific range).

3. Acknowledgements

This document is the work of many people in the multicast community, including (but not limited to) Dino Farinacci, John Meylor, John Zwiebel, Tom Pusateri, Dave Thaler, Toerless Eckert, Leonard Giuliano, Mike McBride, and Pekka Savola.

4. Security Considerations

This document describes operational practices that introduce no new security issues to PIM-SM [RFC4601] in either or SSM [RFC3569] or ASM operation.

However, in the event that the operational practices described in this document are not adhered to, some problems may surface. In particular, Section 2.3 describes the effects of non-compliance of last-hop routers (or, to some degree, rogue hosts sending PIM messages themselves) on the source-specific service model. Creating

the (*,G) state for source-specific (S,G) could enable a receiver to receive data it should not get. This can be mitigated by host-side multicast source filtering.

5. References

5.1. Normative References

- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [RFC3569] Bhattacharyya, S., "An Overview of Source-Specific Multicast (SSM)", RFC 3569, July 2003.
- [RFC3618] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.

5.2. Informative References

- [IANA] <http://www.iana.org>
- [RFCED] <http://www.rfc-editor.org/policy.html>

Authors' Addresses

David Meyer

EMail: dmm@1-4-5.net

Robert Rockell
Sprint

EMail: rrockell@sprint.net

Greg Shepherd
Cisco

EMail: gjshep@gmail.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).

