Internet Engineering Task Force
Internet-Draft
Expires: August 15, 2005

R. Lehtonen
TeliaSonera
S. Venaas
University of Southampton
M. Hoerdt
University Louis Pasteur - LSIIT

Feb 11, 2005

Requirements for discovery of dynamic SSM sources
draft-lehtonen-mboned-dynssm-req-00.txt

Status of this Memo

Copyright Notice

Abstract

This draft identifies the need for discovering new SSM sources in a multicast session. It also defines the basic requirements for such functionality.

Table of Contents

1. Introduction

Many multi-party applications make use of multicast. Multicast offers obvious benefits when one party is sending the same content to multiple receivers. Also traditional multicast [3] allows for a multi-party session to be identified by a single group address.

Any participant can send to the group without knowing who the receivers are, and all receivers can join the group without knowing who will send. This means for e.g. a video conference, one only needs to give the multicast group address to potential participants, possibly making a public announcement. Participants can then come and go as they like, and those that happen to be sending to or receiving from the group address simultaneously will be able to reach each other.

There are several problems with traditional multicast [6] and it's widely believed that Source-Specific Multicast (SSM) [1] is a more scalable and easier to deploy interdomain multicast technology in the long term. One of the simplifications of SSM is that source discovery is not done in the network. It's however precisely the network source discovery (typically using MSDP and Rendezvous-points) that allows anyone to start sending at any point and the receivers to get the data without knowing who the sources are.

SSM requires the application to specify exactly which sources it will receive data from. The source addresses must somehow be learned by the receivers out-of-band. With traditional multicast the multicast group to use must be announced out-of-band before the session is to take place. It may be announced using e.g. SDP over HTTP or SAP. For SSM one could also list the addresses of the sources. This works well where all the sources (participants) are known in advance, but this will often not be the case. Also when announcing a multi-party session publicly and allowing anyone to join, there should be a simple mechanism for registering a participant and getting it announced to the others. This draft defines the basic requirements for such functionality.

2. Problem Statement

As illustrated in the table 1, various multicast applications requirements may require different degree of dynamics in the source discovery process. Existing and future applications require an out of band multicast source discovery mechanism which ideally offers the same level of performances than the current ASM architecture is offering now by in-band means. Distributed Interactive Simulation (D.I.S) [4] is a good example of applications which require a very high level of performance from multicast source discovery.

Applications type	Potential transient degree of the sources	Current solution for for source discovery
Games, D.I.S	High (measured in seconds)	ASM or client-server
Distributed File Sharing	High	client-server
Conferencing	Medium (measured in minutes)	ASM or client-server
TV/Radio	Low (measured in days)	Static publishing

Table 1 : Transient degree of various potential multicast applications.

Today most of the multi-party applications containing transient source of data are client-server based and do not take advantage of IP multicast for source discovery. This limits their efficiency and scalability.

Operational experience and analysis [2] have shown that current ASM model implemented with MSDP [5] for source discovery has severe scaling and security problems in inter-domain scale. In addition to MSDP there are other problems with the RP based source discovery mechanisms like deployment of new mechanisms into routers, RP as the traffic congestion point and insufficient support for both IP versions.

SSM model provides good scaling and security properties and works for both IP versions, but does not provide direct support for source discovery. a typical application that requires discovery of sources during the session is video conferencing. The solution for discovery of new sources during the ongoing session should be standardized for several reasons:

- Clients from different origins/vendors may participate in the same multicast session.

- Without a standardized solution application writers may decide to solve this problem in different non-compatible ways.

- Source discovery must be manageable, so we need a standard that is stable and can be managed/monitored (e.g. to prevent DoS attacks against the multicast infrastructure).

In any cases this source discovery standard will facilitate the

multi-source multicast applications writers to produce new applications with less cost and wider compatibility across the Internet. The multi-source multicast applications deployment effort can be improved by such a solution.

3. Applicability Statement

The source discovery mechanism and its requirements only need that the underlying network supports SSM natively end-to-end. The source discovery mechanism is intended to work in both inter-domain and intra-domain cases. The source discovery mechanism should provide required SSM channel information to receivers. Other application specific discovery requirements are out-of-scope (e.g. discovery of source bandwidth, supported codecs, identification, etc.).

4. Requirements

This section lists the requirements for discovery of dynamic SSM sources. The requirements are separated into general, host and signalling parts.

4.1 General requirements

1. Solution must provide discovery of dynamic SSM sources during the session, offering a comparable level of performance to the current ASM architecture.

2. Solution shouldn't introduce additional requirements on the network (in addition to SSM support).

3. Solution must work in SSM address space.

4. Solution should be easily manageable and provide good security and control properties.

5. Solution should allow co-existence with other source discovery mechanisms.

6. Gradual deployment must be possible without affecting the operation of other SSM hosts.

7. Adding AAA and related functionalities (e.g. source access control) must be possible.

4.2 Host requirements

1. Source discovery functionality must have at least three different

separatable elements; source, receiver and rendezvous elements. Sources are required to register themselves to discovery process. Receivers are required to understand source discovery signalling. Rendezvous function is needed between sources and receivers for matchmaking and control purposes, a common point source discovery signalling.

2. Multiple applications and users on the same host must be able to use the source discovery functionality.

3. A group of users must be able to set up their own rendezvous function.

4. Rendezvous functionality must be able to work in routers and/or in specific hosts if needed for redundancy, availability and control purposes.

5. Rendezvous function should be possible to implement in the SSM application itself.

6. Overhead of being rendezvous must not be too big in terms of processing power, memory or signalling traffic consumption.

7. It must be possible to add rendezvous fallback and load-sharing properties (these functions are not part of the basic requirement set).

8. Registering sources to discovery process must be as simple as possible.

9. There shouldn't be additional hypothesis on the receivers than SSM already brings.

10. Discovery mechanism should provide enough information on the sources that non-active sources and respective SSM channels can be teared down by the receivers.

4.3 Signalling requirements

1. Source discovery signalling must be separated from the actual multicast traffic to achieve the following advantages:

   * Allow path setup before actual traffic is sent (also initial multicast packet gets delivered to receivers).

   * Allow multicast traffic patterns to be different from source discovery. Sources might send at low or high rates

independent of signalling rate. Also the source discovery signalling might be periodical but not necessarily the traffic itself.

* While signalling may go via a central control point the multicast traffic should always take the optimal path (no traffic congestion point).

2. Signalling architecture must be robust. Information about new sources must be distributed to receivers in less than 1 second. New receivers must get the complete source information in less than 15 seconds (worst case).

3. Discovery mechanism must support both IP versions.

4. Minimum amount of extra state in routers for source discovery.

5. Discovery should use multicast where possible, to reduce the overhead of hosting rendezvous function.

6. Standardized and available mechanisms and protocols should be used where appropriate.

7. Source discovery signalling should not necessarily be centralized, the rendezvous function may be distributed across the network to improve the robustness of the mechanism.

5. Security Considerations

This document specifies requirements for source discovery and introduces no new security threats. Security is an important aspect when deciding on a solution.

6. IANA Considerations

This document has no actions for IANA.

7. Acknowledgements

Thanks to Jean-Jacques Pansiot and Pekka Savola for review and constructive comments.

8. References

8.1 Normative References

[1] Bhattacharyya, S., "An Overview of Source-Specific Multicast

(SSM)", RFC 3569, July 2003.

8.2 Informative References

[2] Rajvaidya, P., Ramachandran, K. and K. Almeroth, "Detection and Deflection of DoS Attacks Against the Multicast Source Discovery Protocol", IEEE Infocom 2003.

[3] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.

[4] Pullen, J., Myjak, M. and C. Bouwens, "Limitations of Internet Protocol Suite for Distributed Simulation the Large Multicast Environment", RFC 2502, February 1999.

[5] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.

[6] Savola, P., "IPv6 Multicast Deployment Issues", Internet-Draft draft-ietf-mboned-ipv6-multicast-issues-01, September 2004.

Authors' Addresses

Rami Lehtonen
TeliaSonera
Hataanpaan valtatie 20
Tampere 33100
Finland

Email: rami.lehtonen@teliasonera.com

Stig Venaas
University of Southampton
School of Electronics and Computer Science
Southampton, Hampshire SO17 1BJ
United Kingdom

Email: stig.venaas@uninett.no

Mickael Hoerdt
University Louis Pasteur - LSIIT
C422 - Pole API - Boulevard Sebastien Brant
67400 ILLKIRCH Cedex
France

Email: hoerdt@clarinet.u-strasbg.fr

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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