

Network Working Group
Request for Comments: 2357
Category: Informational

A. Mankin
USC/ISI
A. Romanow
MCI
S. Bradner
Harvard University
V. Paxson
LBL
With the TSV
Area Directorate
June 1998

IETF Criteria for Evaluating Reliable Multicast Transport
and Application Protocols

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 (1998). All Rights Reserved.

Abstract

This memo describes the procedures and criteria for reviewing reliable multicast protocols within the Transport Area (TSV) of the IETF. Within today's Internet, important applications exist for a reliable multicast service. Some examples that are driving reliable multicast technology are collaborative workspaces (such as whiteboard), data and software distribution, and (more speculatively) web caching protocols. Due to the nature of the technical issues, a single commonly accepted technical solution that solves all the demands for reliable multicast is likely to be infeasible [RMMinutes 1997].

A number of reliable multicast protocols have already been developed to solve a variety of problems for various types of applications. [Floyd97] describes one widely deployed example. How should these protocols be treated within the IETF and how should the IETF guide the development of reliable multicast in a direction beneficial for the general Internet?

The TSV Area Directors and their Directorate have outlined a set of review procedures that address these questions and set criteria and processes for the publication as RFCs of Internet-Drafts on reliable multicast transport protocols.

1.0 Background on IETF Processes and Procedures

In the IETF, work in an area is directed and managed by the Area Directors (ADs), who have authority over the chartering of working groups (WGs).

In addition, ADs review individually submitted (not by WGs) Internet-Drafts about work that is relevant to their areas prior to publication as RFCs (Experimental, Informational or, in rare cases, Standards Track). The review is done according to the guidelines set out in the Internet Standards Process, RFC 2026 [InetStdProc96].

The purpose of this document is to present the criteria that will be used by the TSV ADs in reviewing reliable multicast Internet-Drafts for any form of RFC publication.

For I-Ds submitted for Standards Track publication, these criteria must be met or else the ADs will decline to support publication of the document, which suffices to prevent publication. For I-Ds submitted as Experimental or Informational, these criteria must be met or else, at a minimum, the ADs will recommend publishing the I-D with an IESG note prepended stating that the protocol fails to comply with these criteria.

2.0 Introduction

There is a strong application demand for reliable multicast. Widespread use of the Internet makes the economy of multicast transport attractive. The current Internet multicast model offers best-effort many-to-many delivery service and offers no guarantees. One-to-many and few-to-few services may become more important in the future. Reliable multicast transports add delivery guarantees, not necessarily like those of reliable unicast TCP, to the group-delivery model of multicast. A panel of some major users of the Internet, convened at the 38th IETF, articulated reliable bulk transfer multicast as one of their most critical requirements [DiffServBOF97]. Examples of applications that could use reliable bulk multicast transfer include collaborative tools, distributed virtual reality, and software upgrade services.

To meet the growing demand for reliable multicast, there is a large number of protocol proposals. A few were published as RFCs before the impact of congestion from reliable multicast was fully

appreciated, and these should be deprecated [DeprRFCs]. Two surveys of other publications are [DiotCrow97], [Obraczka98].

As we discuss in Section 3, the issues raised by reliable multicast are considerably more complex than those related to reliable unicast. In particular, in today's Internet, reliable multicast protocols could do great damage through causing congestion disasters if they are widely used and do not provide adequate congestion control.

Because of the complexity of the technical issues, and the abundance of proposed solutions, we are putting in place review procedures that are more explicit than usual. We compare this action with an IESG action taken in 1991, RFC 1264 [Routing91], when community experience with standard Internet dynamic routing protocols was still limited, and extra review was deemed necessary to assure that the protocols introduced would be effective, correct and robust.

Section 3 describes in detail the nature of the particular challenges posed by reliable multicast. Section 4 describes the process for considering reliable multicast solutions. Section 5 details the additional requirements that need to be met by proposals to be published as Standards Track RFCs.

3.0 Issues in Reliable Multicast

Two aspects of reliable multicast make standardization particularly challenging. First, the meaning of reliability varies in the context of different applications. Secondly, if special care is not taken, reliable multicast protocols can cause a particular threat to the operation of today's global Internet. These issues are discussed in detail in this section.

3.1 One or Many Reliable Multicast Protocols or Frameworks?

Unlike reliable unicast, where a single transport protocol (TCP) is currently used to meet the reliable delivery needs of a wide range of applications, reliable multicast does not necessarily lend itself to a single application interface or to a single underlying set of mechanisms. For unicast transport, the requirements for reliable, sequenced data delivery are fairly general. TCP, the primary transport protocol for reliable unicast, is a mature protocol with delivery semantics that suit a wide range of applications.

In contrast, different multicast applications have widely different requirements for reliability. For example, some applications require that message delivery obey a total ordering while others do not. Some applications have many or all the members sending data while others have only one data source. Some applications have replicated

data, for example in an n-redundant file store, so that several members are capable of transmitting a data item, while for others all data originates at a single source. Some applications are restricted to small fixed-membership multicast groups, while other applications need to scale dynamically to thousands or tens of thousands of members (or possibly more). Some applications have stringent delay requirements, while others do not. Some applications such as file-transfer are high-bandwidth, while other applications such as interactive collaboration tools are more likely to be bursty but use low bandwidth overall. Some applications will sometimes trade off less than complete reliability for more timely delivery. These requirements each impact the design of reliable multicast protocols in a different way.

In addition, even for a specific application where the application's requirements for reliable multicast are well understood, there are many open questions about the underlying mechanisms for providing reliable multicast. A key question concerns the robustness of the underlying reliable multicast mechanisms as the number of senders or the membership of the multicast group grows.

One challenge to the IETF is to end up with the right match between applications' requirements and reliable multicast mechanisms. While there is general agreement that a single reliable multicast protocol or framework is not likely to meet the needs of all Internet applications, there is less understanding and agreement about the exact relationship between application-specific requirements and more generic underlying reliable multicast protocols or mechanisms. There are also open questions about the appropriate integration between an application and an underlying reliable multicast framework, and the potential generality of a single applications interface for that framework.

3.2 Congestion Control

A particular concern for the IETF is the impact of reliable multicast traffic on other traffic in the Internet in times of congestion, in particular the effect of reliable multicast traffic on competing TCP traffic. The success of the Internet relies on the fact that best-effort traffic responds to congestion on a link (currently as indicated by packet drops) by reducing the load presented to the network. Congestion collapse in today's Internet is prevented only by the congestion control mechanisms in TCP, standardized by RFC 2001 [CongAvoid97, Jacobson88].

There are a number of reasons to be particularly attentive to the congestion-related issues raised by reliable multicast proposals. Multicast applications in general have the potential to do more

congestion-related damage to the Internet than do unicast applications. One factor is that a single multicast flow can be distributed along a large, global multicast tree reaching throughout the entire Internet.

Unreliable multicast applications such as audio and video are, at the moment, usually accompanied by a person at the receiving end, and people typically unsubscribe from a multicast group if congestion is so heavy that the audio or video stream is unintelligible. Reliable multicast applications such as group file transfer applications, on the other hand, are likely to be between computers, with no humans in attendance monitoring congestion levels.

In addition, reliable multicast applications do not necessarily have the natural time limitations typical of current unreliable multicast applications. For a file transfer application, for example, the data transfer might continue until all of the data is transferred to all of the intended receivers, resulting in a potentially-unlimited duration for an individual flow. Reliable multicast applications also have to contend with a potential explosion of complex patterns of control traffic (e.g., ACKs, NACKs, status messages). The design of congestion control mechanisms for reliable multicast for large multicast groups is currently an area of active research.

The challenge to the IETF is to encourage research and implementations of reliable multicast, and to enable the needs of applications for reliable multicast to be met as expeditiously as possible, while at the same time protecting the Internet from the congestion disaster or collapse that could result from the widespread use of applications with inappropriate reliable multicast mechanisms. Because of the setbacks and costs that could result from the widespread deployment of reliable multicast with inadequate congestion control, the IETF must exercise care in the standardization of a reliable multicast protocol that might see widespread use.

The careful review and cautious acceptance procedures for proposals submitted as Internet-Drafts reflects our concern to meet the challenges described here.

4. IETF Process for Review and Publication of Reliable Multicast Protocol Specifications

In the general case of individually submitted Internet-Drafts (proposals not produced by an IETF WG), the process of publication as some type of RFC is described in RFC 2026 (4.2.3) [InetStdProc96]. This specifies that if the submitted Internet-Draft is closely related to work being done or expected to be done in the IETF, the

ADs may recommend that the document be brought within the IETF and progressed in the IETF context. Otherwise, the ADs may recommend that the Internet-Draft be published as an Experimental or Informational RFC, with or without an IESG annotation of its relationship to the IETF context.

The procedure for Reliable Multicast proposal publication will have as its default RFC status Experimental, when the technical criteria listed in Section 5 are deemed to be fulfilled. Both the criteria and the procedure reflect the AD's technical assessment of the current state of reliable multicast technology. It does not reflect the origins of the proposals, which we expect will be equally from commercial vendors with initial products and from researchers.

Work on the development and engineering of protocols that may eventually meet the review criteria could take place either in the IRTF Reliable Multicast Research Group (<http://www.irtf.org>) or a focused short IETF WG with an Experimental product.

When the work in reliable multicast technology has matured enough to be considered for standardization within the IETF, the TSV Area may charter appropriate working groups to develop standards track documents. The criteria for evaluation of standards track technology will be at least as stringent as those described herein (next section).

5. Technical Criteria for Reliable Multicast

The Internet-Draft must (in itself or a companion draft):

- a. Analyze the behavior of the protocol.
The vulnerabilities and performance problems must be shown through analysis. Especially the protocol behavior must be explained in detail with respect to scalability, congestion control, error recovery, and robustness.

For example the following questions should be answered:

How scalable is the protocol to the number of senders or receivers in a group, the number of groups, and wide dispersion of group members?

Identify the mechanisms which limit scalability and estimate those limits.

How does the protocol protect the Internet from congestion? How well does it perform? When does it fail?

Under what circumstances will the protocol fail to perform the functions needed by the applications it serves?

Is there a congestion control mechanism? How well does it perform? When does it fail? Note that congestion control mechanisms that operate on the network more aggressively than TCP will face a great burden of proof that they don't threaten network stability.

- b. Include a description of trials and/or simulations which support the development of the protocol and the answers to the above questions.
- c. Include an analysis of whether the protocol has congestion avoidance mechanisms strong enough to cope with deployment in the Global Internet, and if not, clearly document the circumstances in which congestion harm can occur. How are these circumstances to be prevented?
- d. Include a description of any mechanisms which contain the traffic within limited network environments. If the analysis in a or c shows that the protocol has potential to damage the Internet, then the analysis must include a discussion of ways to limit the scope or otherwise contain the protocol. We recognize that the confinement of Internet applications is an open research area.
- e. Reliable multicast protocols must include an analysis of how they address a number of security and privacy concerns. If the protocol can be used in different modes of secure operation, then each mode must be analyzed.

The analysis must document which of the various parties -- senders, routers (more generally, data forwarders), receivers, retransmission sources -- must be trusted in order to ensure secure operation and privacy of the transmitted data, to what degree, and why. (One issue to address here are "man-in-the-middle" attacks.)

To what degree can data be manipulated so that at least a subset of the receivers receive different copies? Does the protocol allow a group of receivers to determine whether they all received the same data?

What limitations are placed on the retransmission mechanism to prevent it from being abused to flood network links with excessive traffic? Which parties must be trusted to ensure this, and to what degree, and why? The presumption will be that either a congestion control mechanism will inherently limit the volume of retransmission traffic, and that this limiting

influence is robust under concerted attack; or that retransmission requests will be signed in a cryptographically strong manner so that abuses of the mechanism can be traced back to their source. Protocols that do not provide either of these forms of protection face a great burden of proof that they don't threaten network stability.

What sort of key management does the protocol require, and provide for?

6. Security Considerations

This memo specifies in Section 5.e. that reliable multicast Internet-Drafts reviewed by the Transport Area Directors must explicitly explore the security aspects of the proposed design.

7. Acknowledgments

Sally Floyd, Steve McCanne, Mark Handley, Steve Bellovin and Mike Reiter gave especially helpful comments on drafts of this document.

8. References

[RMMminutes 1997] Minutes the Second Reliable Multicast Research Group Meeting. September 1997. <http://www.east.isi.edu/rm>

[Floyd97] Floyd, S., Jacobson, V., Liu, C., McCanne, S., and Zhang, L., A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing. IEEE/ACM Transactions on Networking, December 1997. An online version of the paper is at <http://ee.lbl.gov/floyd/srm-paper.html>.

[InetStdProc96] Bradner, S., "The Internet Standards Process -- Revision 3", RFC 2026, October 1996.

[DiffServBOF97] [6] <http://www.ietf.org/proceedings/97apr> - Transport Area - FDDIFS BOF, April 1997.

[DeprRFCs] Freier, A., "Multicast Transport Protocol", RFC 1301, February 1992. and Braudes, R., and S. Zabele, "Requirements for Multicast Protocols", RFC 1458, May 1993.

[DiotCrow97] Diot, C., Crowcroft, J., Multicast Transport Survey. Journal of Selected Areas in Communications, 1997.

[Obraczka98] Obraczka, K., Multicast Transport Mechanisms: A Survey and Taxonomy. To appear in IEEE Communications, 1998.

[Routing91] Hinden, R., and Internet Engineering Task Force, "Internet Routing Protocol Standardization Criteria", RFC 1264, October 1991.

[CongAvoid97] Stevens, W., "TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms", RFC 2001, January 1997.

[Jacobson 1988] Jacobson, V., Congestion Avoidance and Control, Proceedings of SIGCOMM '88, August 1988, pp. 314-329. An updated version of this paper is available at "<ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z>".

9. Authors' Addresses

Allison Mankin - Past TSV Area Director
USC/ISI East
4350 N. Fairfax Dr., Suite 620
Arlington VA 22203
USA

Phone: 703 812 3706
EMail: mankin@east.isi.edu

Allyn Romanow - Past TSV Area Director
MCI Corporation
2560 North First Street
San Jose, CA 9531
USA

Phone: 408 922 7143
EMail: allyn@mci.net

Scott Bradner - TSV Co-Area Director
Harvard University
1350 Mass. Ave., Rm. 876
Cambridge MA 02138
USA

Phone: 617 495 3864
EMail: sob@harvard.edu

Vern Paxson - TSV Co-Area Director
MS 50B/2239
Lawrence Berkeley National Laboratory
University of California
Berkeley, CA 94720
USA

Phone: 510-486-7504
EMail: vern@ee.lbl.gov

10. Full Copyright Statement

Copyright (C) The Internet Society (1998). 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 assigns.

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.