

**Extensions to the MARS model for Integrated Services
draft-kandlur-issll-rsvp-mars-00.txt**

Status of This Memo

This document is an Internet-Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months, and may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material, or to cite them other than as a ``working draft'' or ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the internet-drafts Shadow Directories on ds.internic.net (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

Abstract

This memo describes an extension to the IP multicast architecture now being developed by the IP over ATM working group, also known as MARS (Multicast Address Resolution Server). This extension expands the responsibilities of the Multicast Server (MCS) to adapt it to the Integrated Services Internet Architecture and related signaling protocols such as RSVP.

1. Introduction

The MARS document [[Arm96](#)] describes a mechanism for handling IP multicast over ATM. Clusters of endpoints, currently defined by IP subnetwork boundaries, share a MARS and use it to track and disseminate information on membership in multicast groups. This allows endpoints to establish VCCs and to transmit to their multicast group over these VCCs. Multicasting is supported using either meshes of VCs or ATM-level multicast servers (MCSs). The choice is made on a per-group basis and is transparent to the endpoints.

1.1. The Problem

In a IP/ATM subnetwork using the MARS protocol for IP multicast address resolution, an Multicast Server (MCS) may be used to forward IP multicast datagrams. This Multicast Server is transparent to the end stations and its operation is further defined in [[RT96](#)]. An end station (the MARS client) requests the MARS server for resolution of IP multicast addresses and receives a list of ATM addresses. It then creates a virtual connection to members of this list, which may consist of the MCS or all of the local members of the multicast group, to deliver packets addressed to this multicast address. Currently, this connection is set up with best-effort service class. The purpose of this document is to investigate the operation of the MARS/MCS in the presence of IP Integrated Services and RSVP.

We consider the ``classical'' framework for handling RSVP and Integrated Services over ATM described in [[Ber96b](#), [Ber96a](#)]. In this model, an RSVP ``flow'' is mapped onto an ATM VCC with QoS attributes. The guidelines for this mapping are described in [[BG96](#)]. Specifically, for multicast flows, when a source receives a RESV message for a given multicast session, it is responsible for creating a virtual connection with the appropriate quality-of-service specification to match the service requested in the RESV message. However, when an MCS is used, this virtual connection only extends from the source to the MCS (see Fig. 1). The forwarding path from the MCS to the destination end-stations will continue to be best-effort since the MCS has no knowledge of the QoS information, and the delivery mechanisms currently defined for the MCS only specify UBR VCCs. Moreover, since the MCS does not process RSVP messages, the destination stations would not be able to detect this breach of service.

Figure 1: Data flow with MCS

2. Possible Solutions

One solution is to place IP multicast group information in the ATM setup message when the sender creates an ATM VC (to the MCS). The receiver, in this case the MCS, can then use this information to create a QoS VC towards the receivers of the multicast group. It would derive the QoS parameters from the QoS and traffic information elements received in the ATM SETUP message, and IP multicast group information from the information elements reserved for protocol information. Unfortunately, this approach has several deficiencies:

1. the MCS does not have access to RESV messages so it is not able to distinguish between receivers which have made reservations from those that have not made any reservation
2. the MCS would require additional information to handle shared reservations

A second approach would be to apply the SBM mechanism [[RY96](#)] to ATM. In this approach, the SBM would be co-located with the MCS and

receivers would send RESV messages to the SBM/MCS. However, it has some disadvantages in the ATM context:

1. unlike the Ethernet environment, the SBM/MCS is used only for reservations for multicast sessions. This requires special mechanism for differentiating between unicast and multicast sessions.
2. it requires either additional configuration of end-stations (for the SBM address) or a cheap multicast mechanism.

A third, and we believe more effective, solution would be to make the MCS participate in RSVP signaling by processing IP/RSVP packets. The MCS would intercept PATH messages and change the RSVP_PHOP field to the address of the MCS. This change would ensure that the RSVP agents on the receivers would direct their RESV responses to the MCS. The MCS would process these RESV responses, perform merging as applicable, and then forward them to the appropriate previous hops. The processing done by the MCS is similar to that performed by an RSVP router, or a subnetwork bandwidth manager (SBM). This approach differs from the SBM approach in that the MCS also processes RSVP PATH messages. There are several advantages to this approach:

1. it is transparent to end-stations
2. it does not require any protocol modifications to RSVP
3. it places responsibility on the MCS, which is part of the data forwarding path and, hence, necessarily involved with the setup of QoS connections.
4. it requires no additional configuration to end-stations or a multicast mechanism for discovery of the MCS.

3. Summary

This draft described several approaches to handling IP Integrated Services flows in the presence of an MCS. From the discussion, it is apparent that the preferred solution is one in which the MCS is RSVP aware and participates in the signaling process in a manner that is transparent to end stations. Further specification of this approach will be pursued based on the recommendation of the working group.

References

- [Arm96] G. Armitage. Support for Multicast over UNI 3.1 based ATM Networks. Internet Draft, [draft-ietf-ipatm-ipmc-12.txt](#), February 1996.
- [Ber96a] L. Berger. RSVP over ATM: Framework and UNI 3.0/3.1 Method. Internet Draft, [draft-berger-rsvp-over-atm-00](#).ps, June 1996.
- [Ber96b] S. Berson. IP Integrated Services Support in ATM. Internet Draft, [draft-ietf-issll-atm-support-00](#).ps, June 1996.
- [BG96] M. Borden and M. Garrett. Interoperation of Controlled-Load and Guaranteed-Service with ATM. Internet Draft, [draft-borden-intserv-atm-mapping-00.txt](#), June 1996.
- [RT96] M. Ammar R. Talpade. Multicast Server Architectures for MARS-based ATM multicasting. Internet Draft, [draft-ietf-ion-marsmcs-00.txt](#), June 1996.
- [RY96] D. Hoffman R. Yavatkar, E. Patki. SBM (Subnet Bandwidth Manager): A Proposal for Admission Control over Ethernet. Internet draft, [draft-yavatkar-sbm-ethernet-00.txt](#), June 1996.

4. Authors' Address

Roch Guerin
IBM T.J. Watson research Center
P.O. Box 704
Yorktown Heights, NY 10598
guerin@watson.ibm.com
VOICE +1 914 784-7038
FAX +1 914 784-6318

Dilip Kandlur
IBM T.J. Watson research Center
P.O. Box 704
Yorktown Heights, NY 10598
kandlur@watson.ibm.com
VOICE +1 914 784-7722
FAX +1 914 784-6205

Doug Williams
IBM T.J. Watson research Center
P.O. Box 704

Yorktown Heights, NY 10598
dougw@watson.ibm.com
VOICE +1 914 784-5047
FAX +1 914 784-6318