

Internet Engineering Task Force
INTERNET-DRAFT
[draft-crawley-rsvp-over-atm-00](#)

M. Borden, E. Crawley, J. Krawczyk
Bay Networks
F. Baker
Cisco Systems
S. Berson
ISI
February 22, 1996

Issues for RSVP and Integrated Services over ATM

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. Internet Drafts 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."

Please check the I-D abstract listing contained in each Internet Draft directory to learn the current status of this or any other Internet Draft.

Abstract

The RSVP and Integrated Services working groups have been working towards the goal of an "Integrated Services Internet" where applications can request a Quality of Service (QoS) from the network in order to meet their end-to-end service requirements. Simultaneously, the IP-over-ATM working group has been specifying the requirements and interactions between IP and Asynchronous Transfer Mode (ATM) technology. The internals and interoperability of ATM devices are specified by the ATM Forum and its working groups. One of the features of ATM technology is the ability to request a virtual circuit (VC) with a specified QoS. Obviously, it would be useful for the work of the Integrated Services Internet to be able to take advantage of the QoS abilities of ATM.

This document outlines the issues related to running RSVP and Integrated Service models over ATM and presents an initial attempt at dividing up the issues between the various working groups and organizations. This is intended as a starting point for the discussions in the joint IP-ATM, RSVP, and Int-Serv meeting to be held at the Los Angeles IETF meeting in March, 1996.

1. Introduction

It is only natural that the RSVP protocol [[1](#)] and the Integrated Services (IntServ) models [[6,7](#)] would like to utilize the QoS properties of ATM as a subnetwork layer. [[2](#)] discusses providing RSVP services over ATM and brings out some of the open issues. Contributions to the ATM Forum MPOA Working Group [e.g. 10, 11, 12] have also investigated some of the issues related to running RSVP over ATM. However, no clear set of directions and plan of action has been determined. Further, there are several groups in the IETF and the ATM Forum that have expressed interest in the issues of RSVP and IntServ over ATM. While it is very important that the issues regarding RSVP and the IntServ models over ATM be addressed, it is important not to duplicate effort between groups. It is also important to avoid considering models and objectives that cannot be demonstrated to be solvable in other switched networks.

The purpose of this document is to briefly describe the issues as they are currently known and suggest which working groups should address them. It is not the purpose of this document to describe RSVP, the IntServ models, or IP over ATM. The reader should be familiar with the concepts of these protocols and models.

2. Issues Regarding the Operation of RSVP and IntServ over ATM

The issues related to RSVP and IntServ over ATM fall into several general classes:

- How to make RSVP run over ATM now and in the future
- When to set up a virtual circuit (VC) for a specific Quality of Service (QoS) related to RSVP
- How to map the IntServ models to ATM QoS models
- How to know that an ATM network is providing the QoS necessary for a flow
- How to handle the many-to-many connectionless features of IP multicast and RSVP in the one-to-many connection-oriented world of ATM

2.1 Modes/Models for RSVP and IntServ over ATM

[3] Discusses several different models for running IP over ATM networks. Any one of these models would work as long as the RSVP control packets (IP protocol 46) and data packets can follow the same IP path through the network. It is important that the RSVP PATH

messages follow the same IP path as the data such that appropriate PATH state may be installed in the routers along the path. For an ATM subnetwork, this means the ingress and egress points must be the same in both directions for the RSVP control and data messages. Within each of these models, there are decisions about using different types of data distribution in ATM as well as different connection initiation. The following sections look at some of the different ways QoS connections can be set up for RSVP.

Expires August 22, 1996

Page 2

INTERNET-DRAFT RSVP and INTSERV over ATM Issues February 22, 1996

2.1.1 UNI 3.x

In the User Network Interface (UNI) 3.0 and 3.1 specifications, [8,9] both permanent and switched virtual circuits (PVC and SVC) may be established with a specified service category (CBR, VBR, and UBR) and specific traffic descriptors in point-to-point and point-to-multipoint configurations. Additional QoS parameters are not available in UNI 3.x and those that are available are vendor-specific. Consequently, the level of QoS control available in standard UNI 3.x networks is somewhat limited. However, using these building blocks, it is possible to use RSVP and the IntServ models.

2.1.1.1 Permanent Virtual Circuits (PVCs)

PVCs emulate dedicated point-to-point lines in a network, so the operation of RSVP can be identical to the operation over any point-to-point network. The QoS of the PVC must be consistent and equivalent to the type of traffic and service model used. The devices on either end of the PVC have to provide traffic control services in order to multiplex multiple flows over the same PVC. With PVCs, there is no issue of when or how long it takes to set up VCs, since they are made in advance but the resources of the PVC are limited to what has been pre-allocated. PVCs that are not fully utilized can tie up ATM network resources that could be used for SVCs.

2.1.1.2 Switched Virtual Circuits (SVCs)

SVCs allow paths in the ATM network to be set up "on demand". This allows flexibility in the use of RSVP over ATM along with some complexity. Parallel VCs can be set up to allow best-effort and better service class paths through the network. The cost and time to set up SVCs can impact their use. For example, it may be better to initially route QoS traffic over existing VCs until an SVC with the desired QoS has been set up. Scaling issues can come into play if a single VC is used per RSVP flow. An RSVP flow is a data flow from one or more sources to one or more receivers as defined by an RSVP

filter specification. The number of VCs in any ATM device is limited so the number of RSVP flows that can be handled by a device would be strictly limited to the number of VCs available, if we assume one VC per flow.

2.1.1.3 Point to MultiPoint

In order to provide QoS for IP multicast, an important feature of RSVP, data flows must be distributed to multiple destinations from a given source. Point-to-multipoint VCs provide such a mechanism. It is important to map the actions of IP multicasting and RSVP (e.g. IGMP JOIN/LEAVE and RSVP RESV/RESV TEAR) to add party and drop party functions for ATM. Point-to-multipoint VCs as defined in UNI 3.x have a single service class for all destinations. This is

Expires August 22, 1996

Page 3

INTERNET-DRAFT RSVP and INTSERV over ATM Issues February 22, 1996

contrary to the RSVP "heterogeneous receiver" concept. It is possible to set up a different VC to each receiver requesting a different QoS, but this again could run into scaling problems.

RSVP sends messages both up and down the multicast distribution tree. In the case of a large ATM cloud, this could result in a RSVP message implosion at an RSVP traffic source with many receivers.

2.1.1.4 Multicast Servers

IP-over-ATM has the concept of a multicast server or reflector that can accept cells from multiple senders and send them via a point-to-multipoint VC to a set of receivers. This moves the VC scaling issues noted previously for point-to-multipoint VCs to the multicast server. Additionally, the multicast server will need to know how to interpret RSVP packets so it will be able to provide VCs of the appropriate QoS for the flows.

2.1.2 ATM Forum Release 4.0

Release 4.0 adds a Leaf Initiated Join (LIJ) capability and standardizes additional QoS parameters for finer control of the characteristics of a VC. LIJ allows an ATM end point to join into an existing point-to-multipoint VC without necessarily contacting the source of the VC. This will reduce the burden on the ATM source point for setting up new branches and more closely matches the receiver-based model of RSVP and IP multicast. However, many of the same scaling issues exist and the new branches added to a point-to-multipoint VC would use the same QoS as existing branches.

The ATM Forum Traffic Management (TM) specification 4.0 greatly enhances the flexibility of choosing QoS in an ATM network. Rather than QoS only being defined by a service category, it is possible to request individual traffic parameters. This will aid enormously in matching QoS models between ATM and IntServ. In addition, TM 4.0 defines new classes, such as VBR-rt and ABR that will be helpful in new service models.

2.1.3 Hop-by-Hop vs. Short Cut

If the ATM "cloud" is made up a number of logical IP subnets (LISs), then it is possible to use "short cuts" from a node on one LIS directly to a node on another LIS, avoiding router hops between the LISs. NHRP [4], is one mechanism for determining the ATM address of the egress point on the ATM network given a destination IP address. It is a topic for further study to determine if significant benefit is achieved from short cut routes vs. the extra state required.

However, if the PATH and RESV messages must follow the same set of IP hops, then short cut paths in both directions must be established.

Expires August 22, 1996

Page 4

INTERNET-DRAFT RSVP and INTSERV over ATM Issues February 22, 1996

This is not a problem for point-to-point VCs, but it is a problem for point-to-multipoint connections since they are unidirectional. This topic requires further study and guidelines.

2.1.4 Future Models

ATM is constantly evolving. If we assume that RSVP and IntServ applications are going to be wide-spread, it makes sense to consider changes to ATM that would improve the operation of RSVP and IntServ over ATM. Similarly, the RSVP protocol and IntServ models will continue to evolve and changes that affect them should also be considered.

2.1.4.1 Heterogeneous Point-to-MultiPoint

The IntServ models and RSVP support the idea of "heterogeneous receivers"; i.e., not all receivers of a particular multicast flow are required to ask for the same QoS from the network.

The most important case of this is where some receivers ask for a specific QoS while others receive the flow in a best-effort manner. In some cases where there are multiple senders on a shared-reservation flow (e.g., an audio conference), an individual receiver only needs to reserve enough resources to receive one sender

at a time. However, other receivers may elect to reserve more resources, perhaps to allow for some amount of "over- speaking" or in order to record the conference (post processing during playback can separate the senders by their source addresses).

In order to prevent denial-of-service attacks via reservations, the service models do not allow the service elements to simply drop non-conforming packets. Guaranteed Delay [6] service requires that the flow must be reshaped before running it through the policing function. Both Guaranteed and Controlled Load [7] assign non-conformant packets to best-effort status (which may result in packet drops if there is congestion).

Emulating these behaviors over an ATM network is problematic and needs to be studied. If a single maximum QoS is used over a point-to-multipoint VC, resources could be wasted if cells are sent over certain links where the reassembled packets will eventually be dropped. In addition, the "maximum QoS" may actually be a degradation in service to the best-effort branches.

The term "variegated VC" has been coined to describe a point-to-multipoint VC that allows a different QoS on each branch. This approach seems to match the spirit of the Integrated Service and RSVP models, but some thought has to be put into the cell drop strategy when traversing from a "bigger" branch to a "smaller" one. The "best-effort for non- conforming packets" behavior must be retained. Early Packet Discard (EPD) schemes must be used so that all the cells

Expires August 22, 1996

Page 5

INTERNET-DRAFT RSVP and INTSERV over ATM Issues February 22, 1996

for a given packet can be discarded at the same time rather than discarding only a few cells from several packets making all the packets useless to the receivers.

2.1.4.2 Lightweight Signalling

Q.2931 signalling is very complete and carries with it a significant burden for signalling in all possible public and private connections. It might be worth investigating a lighter weight signalling mechanism for faster connection setup in private networks.

2.1.4.3 QoS Renegotiation

Another change that would help RSVP over ATM is the ability to request a different QoS for an active VC. This would eliminate the need to setup and tear down VCs as the QoS changed. RSVP allows receivers to change their reservations and senders to change their traffic descriptors dynamically. This, along with the merging of

reservations can create a situation where the QoS needs of a VC can change. Allowing changes to the QoS of an existing VC would allow these features to work without creating a new VC.

2.1.4.4 Group Addressing

The model of one-to-many communications provided by point-to-multipoint VCs does not really match the many-to-many communications provided by IP multicasting. It would be nice to have a scaleable mapping from IP multicast addresses to an ATM "group address" to address this problem.

2.1.5 QoS Routing

RSVP is explicitly not a routing protocol. However, since it conveys QoS information, it may prove to be a valuable input to a routing protocol that could make path determinations based on QoS and network load information. In other words, instead of asking for just the IP next hop for a given destination address, it might be worthwhile for RSVP to provide information on the QoS needs of the flow if routing has the ability to use this information in order to determine a route. The work in this area is just starting and there are numerous issues to consider.

2.2 Reliance on Multicast Routing

RSVP was designed to support both unicast and IP multicast applications. This means that RSVP needs to work closely with multicast and unicast routing. Unicast routing over ATM has been addressed in [RFC1577](#)[13] and [RFC1755](#)[14]. MARS[5] provides multicast

Expires August 22, 1996

Page 6

INTERNET-DRAFT RSVP and INTSERV over ATM Issues February 22, 1996

address resolution for IP over ATM networks, an important part of the solution for multicast routing. Further work is needed for full interoperability of IP multicast routing with ATM.

2.3 Aggregation of Flows

Some of the scaling issues noted in previous sections can be addressed by aggregating several RSVP flows over a single VC if the destinations of the VC match for all the flows being aggregated. However, this causes some complexity in the management of VCs and in the scheduling of packets within each VC at the originating point of the VC. Note that the rescheduling of flows within a VC is not possible in the switches in the ATM network. Additionally, virtual

paths (VPs) could be used for aggregating multiple VCs. These topics need to be investigated and some guidelines need to be provided.

2.4 Mapping QoS Parameters

The mapping of QoS parameters from the IntServ models to the ATM service classes is an important issue in making RSVP and IntServ work over ATM. The issue is that while some guidelines can be developed for mapping the parameters of a given service model to the traffic descriptors of an ATM traffic class, implementation variables, policy, and cost factors can make strict standards problematic.

2.5 Directly Connected ATM Hosts

It is obvious that the needs of hosts that are directly connected to ATM networks must be considered for RSVP and IntServ over ATM. Functionality for RSVP over ATM must not assume that an ATM host has all the functionality of a router, but such things as MARS and NHRP clients might be worthwhile features.

2.6 API Issues

While it has not been much of a concern to the IETF, there are groups considering how applications can specify QoS parameters to the network in a manner that makes sense to application programmers. It makes sense that APIs specify parameters in a manner that is independent of the layer 2 network. It is also important for the application to be able to renegotiate and change its requested QoS in a manner that is consistent with the RSVP and IntServ models.

2.7 Accounting and Policy Issues

Since RSVP and IntServ create classes of preferential service, some form of administrative control or cost allocation is needed to control abuse. There are certain types of policies specific to ATM

and IP over ATM that need to be studied to determine how they interoperate with IP policies being developed. Typical IP policies would be that only certain users are allowed to make reservations. This policy would translate well to IP over ATM. There may be a need for policies specific to IP over ATM. For example, since signalling costs are high relative to IP, an IP over ATM specific policy might restrict the ability to change the prevailing QoS in a VC. If VCs are relatively scarce, there also might be specific accounting costs

in creating a new VC. The work so far has been preliminary, and much work remains to be done.

3. Responsibilities for Resolving Issues

The following sections list the working groups involved in some of the issues listed in the previous section. Along with each group is a list of issues that the authors believe each group should address for RSVP and IntServ over ATM. The goal is to not have multiple working groups working on the same problem(s). This list is a starting point for discussions. However, it is hoped that the discussions will be brief such that the working groups can get down to business quickly.

3.1 IETF Working Groups

The IETF has several working groups that can be involved in RSVP and IntServ over ATM. IETF working groups are chartered to solve specific problems and some of the issues noted in the following may not be specifically in their current charter so it may require amendment of the charter or creation of a new or sub working group.

3.1.1 IP over ATM (IPATM)

This group has been most involved in getting IP to work over ATM and is responsible for issues relating to both IP hosts and IP routers. They have spent time on encapsulations and multicast address resolution and have been the focal point for liaison activities between the ATM Forum and the IETF. They should continue to be the focal point for liaison activities as they relate to RSVP and IntServ over ATM and should address specific issues related to the operation of the IP protocols, including RSVP, over ATM.

As the focal point for liaison activities, the IPATM WG can encourage contributions for the ATM Forum as needed.

3.1.2 Integrated Services (INTSERV)

IntServ has been developing service models for the integrated services internet. These models turn into flow specification formats for RSVP. It is important for IntServ to examine methods by which these flow requirements can be achieved over ATM, using the ATM

notions of QoS. There is typically not a unique method of doing so,

and it is expected that the working group would be developing guidelines and not standardized rules.

In addition to describing the parameters of the flow there is the problem of managing the flows. This needs to be studied at the same time. The issues involve policies for aggregating flows over a single VC, when new VCs need to be established, and managing point-to-multipoint "variegated" VCs. Since most of these decisions are local, the need not be standardized but guidelines may be appropriate.

As new application requirements are generated, the IntServ WG will need to develop new service models, flow specification formats, and methods to achieve the services over ATM.

3.1.3 ReSerVation Protocol (RSVP)

The RSVP WG has been focused on developing version 1 of the RSVP protocol in a media and routing independent manner. They should continue addressing further issues with RSVP such as security, authentication, and access control. Related to ATM, the RSVP WG should consider the ramifications of RSVP in Non-Broadcast MultiAccess (NBMA) networks such as ATM and publish guidelines for RSVP over such nets.

3.1.4 InterDomain Multicast Routing (IDMR)

The IDMR WG has been concerned with scaleable multicast routing protocols. Obviously, there are going to be issues for multicast routing in NBMA networks and IDMR must consider the implications for the protocols being developed. QoS multicast routing is also something that IDMR should consider as it relates to RSVP and ATM.

3.1.5 Routing Over Large Clouds (ROLC)

As an NBMA network that can span a "large cloud," ATM can make use of the NHRP protocol being developed in the ROLC WG. The WG has been very aware of the needs of IP over ATM in the development of NHRP and should continue to do so. As mentioned earlier, the benefit of short cut routing needs to be weighed against the state and cost of establishing direct VCs between nodes on different ATM LISs. The benefit vs. the scaling issues must be considered for RSVP and IntServ.

3.2 ATM Forum

The ATM Forum exists for its members and all of its activities must be approved by the members. There are several working groups that

INTERNET-DRAFT RSVP and INTSERV over ATM Issues February 22, 1996

have interest in RSVP and IntServ over ATM. There was a well-attended Birds of a Feather (BOF) session for RSVP and IntServ over ATM as part of the MPOA WG at the February 1996 ATM Forum meeting in Beverly Hills, CA.

3.2.1 Signalling (SIG)

Obviously, any changes to ATM signalling to better support RSVP and IntServ will have to be approved by the Signalling WG. Such features include variegated point-to-multipoint VCs and renegotiated QoS for existing VCs. These are not trivial problems and are expected to take some time, but contributions from members are expected in this area soon.

The Signalling WG should also be involved in making a "group addressing" scheme to provide many-to-many communications in a scaleable manner, much like the current IP multicast model.

3.2.2 MultiProtocol Over ATM (MPOA)

The MPOA WG is addressing the issues of running multiple layer 3 protocols in an ATM environment. The features of ATM such as QoS and short cut routing are expected to be utilized. The features of RSVP that affect the MPOA environment should be considered by MPOA but there will be considerable overlap with the work of the IETF RSVP WG. In these cases, the RSVP WG should address the issues.

MPOA is also the focal point for the liaisons between the ATM Forum and the IETF and should continue to do so with feedback from other ATM Forum and IETF working groups feeding into the liaisons.

3.2.3 Traffic Management (TM)

The TM working group has recently been concerned with divorcing strict traffic classes from negotiated QoS parameters, and the introduction of the Available Bit Rate (ABR) class. The group is at the point where new work items are being defined, but it is plain that future work in the areas of clarification of QoS objectives and point-to- multipoint QoS is warranted. These are areas in which the developing IntServ ideas can have influence. In addition, it would be helpful for the TM group to provide examples of how IntServ QoS objectives can be obtained with TM models.

3.3 Other Organizations

3.3.1 IEEE

IEEE 802.1p and 802.9 have proposed various services that may be useful in running the IntServ models over IEEE- specified

Expires August 22, 1996

Page 10

INTERNET-DRAFT RSVP and INTSERV over ATM Issues February 22, 1996

subnetworks. To the extent that these have an impact on ATM, they should also be reflected in the layering of IntServ models on ATM. However, we recommend that IEEE considerations not further complicate these issues beyond liaison activities.

3.3.2 WinSock2 Forum

WinSock2 is an API for Windows platforms that can specify QoS parameters independent of the network transport. Obviously, the work of the groups above should be input to WinSock2 so the APIs match the services available.

4. Summary

In this paper, we have suggested a preliminary plan of action for addressing the issues in layering the Integrated Services models and RSVP onto an ATM subnetwork. This problem is complicated by the fact that multiple IETF working groups and multiple organizations are involved and by the fact that ATM technology, the Integrated Services models, and RSVP are evolving over time. Our objective is to see the work go forward openly and with a minimum of both duplicated effort and friction.

5. References

[1] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin. Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification. Internet Draft, [draft-ietf-rsvp-spec-09](#). February 1996.

[2] M. Borden, E. Crawley, B. Davie, S. Batsell. Integration of Real-time Services in an IP-ATM Network Architecture. Request for Comments (Informational) [RFC 1821](#), August 1995.

[3] R. Cole, D. Shur, C. Villamizar. IP over ATM: A Framework Document. Internet Draft, [draft-ietf-ipatm-framework-doc-07](#), February 1996.

[4] D. Katz, D. Piscitello, B. Cole, J. Luciani. NBMA Next Hop

Resolution Protocol (NHRP). Internet Draft, [draft-ietf-rolc-nhrp-07.txt](#), December 1995.

[5] G. Armitage, Support for Multicast over UNI 3.0/3.1 based ATM Networks. Internet Draft, [draft-ietf-ipatm-ipmc-11.txt](#). February 1996.

[6] S. Shenker, C. Partridge. Specification of Guaranteed Quality of Service. Internet Draft, [draft-ietf-intserv-guaranteed-svc-03.txt](#), November 1995.

Expires August 22, 1996

Page 11

INTERNET-DRAFT RSVP and INTSERV over ATM Issues February 22, 1996

[7] J. Wroclawski. Specification of the Controlled-Load Network Element Service. Internet Draft, [draft-ietf-intserv-ctrl-load-svc-01.txt](#), November 1995.

[8] ATM Forum. ATM User-Network Interface Specification Version 3.0. Prentice Hall, September 1993

[9] ATM Forum. ATM User Network Interface (UNI) Specification Version 3.1. Prentice Hall, June 1995.

[10] A. Demirtjis, S. Berson, B. Edwards, M. P. Maher, B. Braden, A. Mankin. RSVP and ATM Signalling, ATM Forum Contribution 96-0258, January 1996.

[11] M. Borden, K. Faulkner, E. Stern. Support for RSVP in ATM Networks, ATM Forum Contribution 96-0039, February 1996.

[12] F. Baker, A. Lin, Y. Rekhter, Support for RSVP and IP Integrated Services over ATM, ATM Forum Contribution 96-0258, January 1996.

[13] M. Laubach, Classical IP and ARP over ATM. Request for Comments (Proposed Standard) [RFC1577](#), January 1994.

[14] M. Perez, A. Mankin, E. Hoffman, G. Grossman, A. Malis, ATM Signalling Support for IP over ATM, Request for Comments (Proposed Standard) [RFC1755](#), February 1995.

6. Author's Address

Marty Borden
Eric S. Crawley
John J. Krawczyk
Bay Networks
3 Federal Street
Billerica, Ma 01821
+1 508 670-8888

mborden@baynetworks.com
esc@baynetworks.com
jj@baynetworks.com

Fred Baker
Cisco Systems
519 Lado Drive
Santa Barbara, California 93111
+1 805 681-0115
fred@cisco.com

Steve Berson
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
+1 310 822-1511
berson@isi.edu

Expires August 22, 1996

Page 12