

**ATM Virtual Circuit Identification Support for an RSVP-based Service
draft-williams-issll-vcuse-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

In this document we present a specification for associating RSVP flows with Virtual Connections on an ATM network. The text in this document is meant as an addendum to the text in [Subsection 3.2](#) ``Sending RSVP Messages'' of the RSVP specification [Bzb+96]. This document specifies a usage of the Logical Interface Handle (LIH) in the RSVP_HOP Class Object when either a host or a router is attached to an ATM network. This document also specifies ATM UNI [[For94](#)] signaling semantics for associating a switched ATM connection with RSVP flows. These ATM UNI semantics use the Broadband High Layer Information (BHLI) Information Element.

1. Motivation

The purpose for this document is to propose a standard method for hosts and routers using the RSVP Protocol [Bzb+96] over an ATM network to associate RSVP flows with a unique Virtual Connection (permanent or switched). Specifically this document presents a specification that is to be used as addendum text to [Subsection 3.2](#) ``Sending RSVP Messages'' of the RSVP Specification. This document does not present any new RSVP Objects, or in fact, any particular change to the basic framework of either RSVP or ATM UNI signaling, rather, this document presents a usage of existing protocol mechanisms.

The motivation for this work is to allow hosts and routers to efficiently and consistently use the underlying ATM link layer technology (i.e. manage VCs). While there are several alternatives to mapping RSVP sessions to a data-link connection, this document specifies a "one connection per session" model, in which a new connection is created for each RSVP session. This model has the following advantages:

- It exploits the traffic control and scheduling capabilities of the data link, which are typically supported with hardware mechanisms.
- It reduces the network level processing load by removing the need for multiplexing of RSVP sessions onto a connection.

The advantages of this model can be extended further when the receiving end-point of the connection is made aware of the association between the RSVP session and the data-link connection. For example, when the receiver is a network device, the data-link connection can be used as a mechanism for pre-sorting packets to their network level session. This can assist the network device in performing appropriate forwarding and scheduling actions at very high speeds. Similarly, when the receiver is an end-station, the data link connection can be exploited to perform an early de-multiplexing of arriving packets to their appropriate application endpoint. This early de-multiplexing has been shown to reduce latency and processing overheads within the end-station, thereby improving the performance of multimedia applications. Moreover, when the consuming application endpoint is a device within the end-station (such as a graphics display adapter) the data link connection can be used to direct packets to the device.

Another reason for keeping track of RSVP-related connections at the data link layer is related to connection management. The IP over ATM signaling standard ([RFC 1755](#)) specifies that end systems must support the use of inactivity timers for connections. However, this

provision should not apply to connections that were created for RSVP. The presence of RSVP signaling messages is sufficient to indicate that the application desires a reservation, even when there is no data traffic on the session. Since these data link connections are established and maintained at the behest of RSVP, it is logical to leverage RSVP to convey data link connection information between the endpoints. Moreover, RSVP messages contain all the elements necessary to identify the data flows that are carried on the connection.

The advantages of data link connections for supporting QoS have been recognized in many different forms, and there have been proposals put forth to define protocols to exploit them. Among these proposals are SINP and the Ipsilon flow management protocol. Our approach differs from these efforts in that it is defined as a usage of existing protocols, namely RSVP and ATM UNI Signaling. We feel this to be of importance as while the functionality we seek to provide is useful in many different settings, its support should be through existing signaling protocols rather than by developing a new protocol for that purpose.

2. Specification

The following text is meant to be an addendum to the end of [Subsection 3.2](#) ``Sending RSVP Messages'' in the RSVP Specification.

When operating over and ATM link-layer, it can be desirable to associate each RSVP session with a separate Virtual Connection (either Permanent or Switched). ATM attached RSVP capable hosts and routers can support this capability in the following manner:

1. When a node N forwards a PATH message over an ATM network connection, this PATH message is sent on a default IP/Control Virtual Connection. The message includes an RSVP_HOP object class with one of the following codings for the LIH:


- In a Permanent Virtual Connection environment, the LIH contains the VPI:VCI of the ATM Virtual Connection as follows:

```

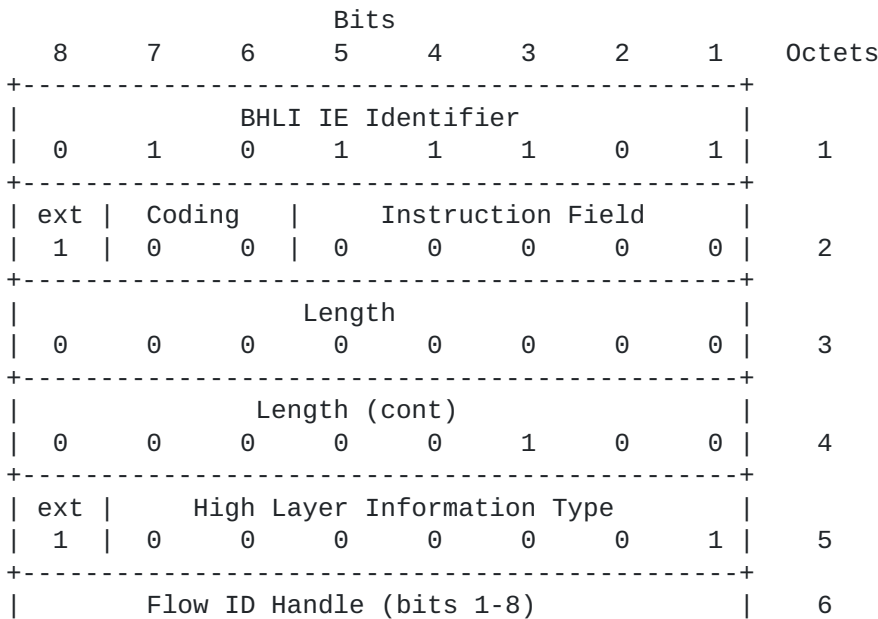
+-----+-----+-----+-----+
|   NULL       VPI       VCI       VCI   |
+-----+-----+-----+-----+

```

- In a Switched Virtual Connection environment, the Logical Interface Handle contains a Flow ID Handle as follows:




2. Both N and the next hop node N' store the LIH in their Path State Tables.
3. When N' sends a RESV message to N, it includes the LIH value from the Path State Table in an RSVP_HOP object.
4. When a RESV message(s) arrives at N for this session, the LIH information from the Path State Table is used to identify the Virtual Connection for this session's data packets. In the case of PVCs, this association is immediate, in the case of SVCs, the VPI:VCI is identified after node N performs the ATM UNI signaling described in Step 5 below.
5. In the case of a Switched Virtual Connection, when N receives the RESV message it performs ATM UNI signaling [For94] to establish the connection for this reservation. In the UNI SETUP Message, a Broadband High Layer Information Information Element is included with the 32-bit Flow ID Handle for the session that will be using this Virtual Connection (i.e the same Flow ID that was sent in the original PATH message for this session). When N' receives the ATM UNI SETUP message, it extracts the Flow ID Handle from the BHLI field and uses it to associate the particular RSVP session with the VPI:VCI being setup. The BHLI Information Element is coded as follows:



Flow ID Handle (bits 9-16)	7
Flow ID Handle (bits 17-24)	8
Flow ID Handle (bits 25-32)	9

3. Example

The figure below shows a simple example network. In which two IP hosts are connected to each other through an ATM switch and a router.



H1 is an RSVP sender and H2 is an RSVP receiver; assume two PVCs between H1 and R1 and R1 and H2 are connected by an Ethernet. Using the mechanisms outlined in this document, H1 sends an RSVP PATH message over one of the two PVCs (referred to as the default VC). This PATH message flows as normal through the network to the destination, in this case H2 is the destination. The PATH message contains the VPI:VCI of the second PVC (referred to as the reserved data PVC) in the LIH field of the RSVP_HOP Object as shown above. When R1 receives the PATH message, it creates a Path State Table entry that includes the VPI:VCI that is located in the LIH field. H2 determines that it would like to setup a Reservation for this particular session, that is, it sends a RESV message. When R1 receives this message, along with normal RSVP processing, it associates the PVC connection that corresponds to the VPI:VCI that was stored in the Path State Table with this session. R1 then sends the RESV message to H1 with the LIH again set to its VPI:VCI for the reserved data PVC. When H1 receives the RESV message it sets up its resources and sends the session's data on the reserved data PVC.

In the case of an SVC connection, the flow is the same as the PVC case, with the exception that H1 includes a 32-bit Flow ID handle in the LIH field instead of a VPI:VCI (note that in the SVC case, the reserved data SVC is not established until the sender receives the RESV message), R1 stores the Flow ID handle in the Path State Table. When the RESV message arrives at H1, it requests a reserved data VC to R1 using UNI signaling. The UNI signaling includes the Flow ID handle in the BHLI as shown above. When R1 receives the call request, it is responsible for associating this particular VC with the RSVP session that is identified in the Path State Table (i.e.

comparing Flow ID handles). When the ATM SVC is fully established, the reserved data flows over this SVC.

4. References

[Bzb+96] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1, Functional Specification", Internet Draft, [draft-ietf-rsvp-spec-12.txt](#), May 1996.

[For94] ATM Forum, "ATM User-Network Interface Specifications, Version 3.1", September 1994.

[For95] ATM Forum, "ATM User-Network Interface Specifications, Version 4.0", ATM Forum/94-1018.

5. Authors' Address

Roch Guerin
T. J. Watson Research Center
IBM Corporation
30 Saw Mill River Rd.
Hawthorne, NY 10532

Phone: +1-914-784-7038
E-mail: guerin@watson.ibm.com

Dilip Kandlur
T. J. Watson Research Center
IBM Corporation
30 Saw Mill River Rd.
Hawthorne, NY 10532

Phone: +1-914-784-7722
E-mail: kandlur@watson.ibm.com

Doug Williams
T. J. Watson Research Center
IBM Corporation
30 Saw Mill River Rd.
Hawthorne, NY 10532

Phone: +1-914-784-5047

E-mail: dougw@watson.ibm.com