

Native ATM Support for ST2+

Status of This Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

As the demand for networked realtime services grows, so does the need for shared networks to provide deterministic delivery services. Such deterministic delivery services demand that both the source application and the network infrastructure have capabilities to request, setup, and enforce the delivery of the data. Collectively these services are referred to as bandwidth reservation and Quality of Service (QoS).

The IETF is currently working on an integrated services model to support realtime services on the Internet. The IETF has not yet focused on the integration of ATM and its inherent QoS and bandwidth allocation mechanisms for delivery of realtime traffic over shared wires. (ATM hardware and interfaces provide the network infrastructure for the deterministic data delivery, however the host resident protocol stacks and applications need more attention.)

Current IETF efforts underway in the IP over ATM (ipatm) working group rely on intserv, rsvp and ST2 to address QoS issues for ATM. As such, RFC 1577 and the ATM Forum's Lan Emulation do not provide direct QoS and bandwidth allocation capabilities to network applications. Without providing a mapping of reservations-style QoS to ATM signalling, ATM will remain a 'wire' rather than a shared media infrastructure component.

This memo describes a working implementation which enables applications to directly invoke ATM services in the following environments:

- ATM to internet,
- internet to ATM, and
- internet to internet across ATM.

Table of Contents

1.0	Introduction.....	2
2.0	ST-2 and ST-2+.....	5
3.0	Implementation Issues for Reservations over ATM.....	6
3.1	Addressing.....	6
3.2	Changes to Bandwidth and QoS.....	6
3.3	Multicasting.....	7
3.4	Receiver Initiated JOIN Requests to Multicast Groups.....	8
3.5	Computation of QoS Parameters.....	8
3.6	Use of HELLOs.....	9
4.0	Reservation Signalling with ATM.....	9
4.1	Embedded Reservation Signalling within Q.2931.....	10
4.2	In-Band Reservation Signalling.....	11
4.3	Dedicated Virtual Circuits for Reservation Signalling.....	12
4.4	Reservation Signalling via IP over ATM or LAN Emulation...	13
4.5	Summary of Reservation Signalling Options.....	14
5.0	Mapping Reservation QoS to ATM QoS.....	15
5.1	CPCS-SDU Size Computation.....	16
5.2	PCR Computation.....	17
5.3	Maximum End to End Transit Delay.....	17
5.4	Maximum Bit Error Rate.....	18
5.5	Accumulated Mean Delay.....	18
5.6	Accumulated Delay Variance (jitter).....	18
6.0	Data Stream Transmission.....	18
7.0	Implementation Considerations and Conclusions.....	19
8.0	Security Considerations.....	20
9.0	References.....	20
10.0	Author's Address.....	21

1.0 Introduction

The ATM Forum and the IETF seem to approach ATM networking differently.

The ATM forum appears to believe that host systems require no protocols beyond OSI layer 2 to deal with ATM. They define a layer 2 API and Q.2931 signaling for all new applications.

LAN Emulation, a mechanism to make the ATM interface appear to be a LAN/internet, is intended to support 'legacy' network applications. LAN emulation does not provide applications any visibility of the ATM features, nor does it provide a mechanism to allow applications to request specific ATM services. With LAN Emulation, application traffic shares virtual circuits with no policing or guarantees of service. LAN Emulation simply extends LAN characteristics to ATM.

Thus far, the IETF, through RFC 1577[1] treats an ATM network as a wire. The ipatm working group has explicitly left issues of specific QoS handling out of their specifications and working documents. Current approaches do not give the application access to individual virtualcircuits and their associated guaranteed bandwidth and QoS. Instead, all IP traffic between two hosts shares virtual circuits with no granularity assigned to application-specific traffic or QoS requirements.

Thus, neither LAN Emulation nor RFC 1577 (IP over ATM) uses the features of ATM that make it a unique and desirable technology. RFC 1821 (Integration of Realtime Services in an IP-ATM Network Architecture) [2] raises many of the issues associated with current IETF efforts towards integrating ATM into the Internet, but it does not propose any solutions.

This document offers a framework for provision of native ATM circuits for applications which require bandwidth guarantees and QoS. It identifies the requirements of a native ATM protocol which is complementary to standard IP and describes one working implementation.

This document recognizes the fact that it is critical that such a native ATM protocol is consistent in the four topologies described in [2]:

- * Communication across an ATM-only network between two hosts directly connected to the ATM network,
- * Communication between ATM connected hosts which involves some non-ATM subnets,
- * Communication between a host on a non-ATM subnet and a host directly connected to ATM,
- * Communication between two hosts, neither of which has a direct ATM connection, but which may make use of one or more ATM networks for some part of the path.

That is, to the host systems, the underlying type of network remains transparent even when QoS is involved in internet, ATM, and mixed networking environments. To make this consistency possible, the 'native ATM' protocol must also be:

- * Multicast capable, to optimize transmission overhead and support ATM multipoint facilities,
- * Routable, to enable transmissions across subnets and internets,
- * QoS knowledgeable, to take advantage of ATM QoS facilities,
- * Capable of Bandwidth/QoS Reservation to allocate proper facilities for application traffic as it travels across

different types of networks: to effectively extend virtual circuits across internets, and

- * Capable of policing to ensure proper packet scheduling behavior and to protect guaranteed services at merge points.

Clearly the protocol should support reservations. Reservation protocols enable creation of 'virtual circuits' with guaranteed bandwidth and QoS on the LAN or internet, and simultaneously can act as signaling mechanisms to routers or ATM interfaces to request provisioning of circuits. Use of a reservation protocol makes characteristics of mixed networks (LANs, internet, ATM, ISDN) transparent to the host systems. That is, a reservation will allow the host or router to provision ATM circuits which match the reservation, but in mixed networks, will allow routers and host to provide bandwidth reservation and QoS across the non-ATM interfaces as well. Effectively, the reservation maps ATM virtual circuits to reservations on subnets and internets.

This creates a consistent End-to-End, QoS-guaranteed service for mixed network topologies.

While it is beyond the scope of this document, the same requirements apply to mixed ISDN networks and are currently being explored by the ITU for their H.323, H.223, and T.123 standards.

Arguably, the reservation protocol that provides this end-to-end guaranteed service should be connection-oriented to facilitate mapping of real connections (ATM or ISDN) with virtual connections on the LAN/internet. [2] points out the shortcomings of IP and RSVP [3] in the ATM environment. Most notable among these are the difficulty of mapping connectionless traffic to ATM connections, the constant softstate refreshes of RSVP (and merging of RESV messages), the receiver orientation of RSVP, and the dependence on IP multicast.

[6] is an excellent document that proposes solutions to many of the issues raised in [2], but the solutions recommend modifications to the current RSVP and ATM implementations. Recently, issues of incompatibility with the current IP over ATM model, VC explosions due to use of multicast groups and VC explosions due to features associated with heterogeneous receivers suggest that the current version of RSVP may be inappropriate for ATM implementations.

Since ATM is connection-oriented, hard state, and origin-oriented for transmission, signaling, and multicast, and is bandwidth and QoS knowledgeable, perhaps the simplest and most elegant approach to a native protocol for ATM would include a protocol that shares these characteristics.

In surveying protocols described in IETF RFCs and Internet Drafts, only two seem to meet these requirements: Experimental Internet Stream Protocol: Version 2 (RFC 1190) [4] and Internet Stream Protocol Version 2+ (RFC 1819) [5]; ST2 and ST2+ respectively.

2.0 ST2 and ST2+

Both ST2 and ST2+ have been given the Internet Protocol Version 5 (IPv5) designation. In fact, ST2+ is an updated version of ST2. Both protocols are origin-oriented reservation and multicast protocols that provide bandwidth and QoS guarantees through internets. Unlike IPv4 or IPv6, ST2 and ST2+ are connection-oriented, subscribing to the philosophy that once a connection is established, protocol and routing overhead can be substantially reduced. This carries forward to QoS and Bandwidth Reservation as well, simplifying the implementation of QoS guarantees. THESE PROTOCOLS WERE INTENDED TO COMPLEMENT STANDARD CONNECTIONLESS IP, RECOGNIZING THAT WHILE MOST INTERNET TRAFFIC BENEFITS FROM CONNECTIONLESS NETWORKING, PERFORMANCE AND QoS GUARANTEES COULD BE ACHIEVED MOST EASILY WITH INTERNET CONNECTIONS.

Both ST2 and ST2+ really consist of two protocols: SCMP and ST. SCMP is analogous to ICMP in that it is the control and signaling protocol, while ST is the low-overhead streaming protocol. ST-2 uses standard IP addresses during connection setup, but then reduces header overhead by including a stream identifier in each data packet.

ST2+ includes simplification of many of the original ST2 features as well as clarification of the ST2 specification. Among these simplifications and clarifications are:

- 1) Much simpler connection setup.
- 2) Flow Specification independence and consolidation of experimental Flow Specifications.
- 3) Clarification on the implementation of Groups of Streams.
- 4) Clarification of leaf-initiated JOINS in multicast trees (several ST2 implementations had done this).

While there continues to be a dramatic increase in the use of ST2 for videoconferencing, video on demand, telemetry applications and networked virtual reality, ST2+ has no commercial implementations and is not yet supported by any router vendors. This is because ST2+ was released as an RFC late in the summer of 1995. It is expected that several implementations will appear over the coming months. As such, the approach described in this document applies to both protocols, and, in fact, would be valid for any other similar protocol used to establish 'native' ATM circuits. Since ST2 and ST2+ are so similar, this document will refer to 'the ST2 protocols'

generically in describing an implementation approach to both. Where particular features of ST2+ are required or affect implementation, 'ST2+' will be used specifically.

3.0 Implementation Issues for Reservations over ATM

As described above, ST is a connection-oriented, hard state, origin-oriented multicast protocol and thus maps fairly well to ATM. However, ST-2 has several features that may be difficult to support in the current version of ATM signaling with Q.2931 and UNI 3.1. Among these are:

- 1) Addressing.
- 2) Changes to Bandwidth and QoS.
- 3) Multicasting.
- 4) Receiver initiated JOINS to multicast groups.
- 5) Computation of certain QoS parameters.
- 6) Use of HELLOs.

The degree of difficulty in supporting these functions is dependent on the signaling mechanism chosen. See Section 4 for descriptions of possible signaling approaches and their respective impact on the features listed above.

3.1 Addressing

Of course mapping an Internet address to ATM address is always problematic. It would be possible to set up a well known ARP server to resolve the IP addresses of targets. However, the widespread deployment of IP over ATM and LAN emulation in host-based ATM drivers, and the assumption that most host systems will be running some IP applications that do not need specific QoS and bandwidth provisioning, suggests that use of ARP facilities provided by IP over ATM and LAN Emulation is the most obvious choice for address resolution.

It should be noted that ATMARP returns the ATM address. For some implementations (particularly kernel-based protocols), an NSAP address is also required. Since these addresses are often difficult to get from the ATM network itself in advance of the connection, it may be necessary to invoke out-of-band signaling mechanisms to pass this address, or it may be better to create an NSAP address server.

3.2 Changes to Bandwidth and QoS

Both ST-2 and ST-2+ allow the origin to dynamically change the QoS and Bandwidth of a particular stream. At this time Q.2931 and UNI 3.1 do not support this feature. Until this capability is available,

full support of the SCMP CHANGE message for dedicated ATM circuits (one reservation = one ATM circuit) can only be implemented by tearing down the existing VC for a stream and establishing a new one if efficient use of ATM resources are to be preserved.

Of course, the CHANGE message can simply be passed across the ATM virtual circuit to the hosts or routers. This would allow the hosts to relax resource requirements locally, and permit routers to relax access to downstream circuits, but the ATM VC itself, would still retain excessive bandwidth.

In addition, if the implementation allows sharing of virtual circuits by multiple streams, the bandwidth/QoS of individual streams within the VC can be CHANGED.

3.3 Multicasting

ST-2 and ST-2+ support origin-oriented multicasting. That is, the origin of a stream explicitly specifies the addresses of the targets it wants involved in the connection. In addition, the origin can Add or drop targets as desired. Aside from receiver-initiated JOINS (discussed in section 3.4), there is a one to one mapping between ST-2 multicast and ATM multipoint connections. Origin-initiated additions can be accomplished through an ADDPARTY, and drops can be done through DROPPARTY.

A key goal in implementation of a native ATM protocol is to ensure consistent implementation for unicast and multicast data transfers. One difficulty in doing this with ATM Virtual Circuits is the fact that point-to-point circuits are duplex, while multipoint circuits are simplex. This means that for multicast connections to be mapped to multipoint ATM Virtual Circuits, any two-way, end-to-end signaling must be done out of band. An alternative is to let the local reservation agent act as a split/merge point for the connection by establishing point-to-point Virtual Circuits for each member of the multicast group directly connected to the ATM network. For multicast group members not directly connected to the ATM network, traffic can be multicast to the router connected at the edge across a single virtual circuit associated with the reservation.

Section 4 describes alternative mechanisms for implementing signaling.

Included in each discussion is the optimal means for mapping multicast to ATM point-to-point or multipoint circuits.

Note that the fact that ST-2 does not rely on IP multicast is a strong advantage in implementation of a native protocol for ATM. The

one-to-one mapping of ST-2 multicast connections to ATM multipoint virtual circuits minimizes the number of circuits required to support large multicast groups.

3.4 Receiver Initiated JOINS to Multicast Groups

ST-2+ provides an in-band mechanism to permit receivers to join an existing stream. Based on an origin-established authorization level, the JOIN can be refused immediately, can be allowed with notification of the origin, or can be allowed without notifying the origin. This capability is made available through a new SCMP JOIN message. If the receiver knows the IP address of the origin and the Stream ID, he can join the stream if authorized to do so.

Note that since the JOIN flows from the receiver to the origin, there will be issues in trying to support this feature with Q.2931 and UNI 3.1. The JOIN may have to be sent out of band depending on the signaling mechanism chosen (section 4) because of the uni-directional flow for point to multipoint ATM connections. This is supposed to change with availability of UNI 4.0.

ST-2 did not support receiver initiated JOINS (unlike ST-2+). However, most implementations created an out-of-band, or SCMP extension to support this facility. Again, depending on the SCMP signaling mechanism chosen, this feature may be difficult to support.

3.5 Computation of QoS Parameters

The recommended flow specifications (flowspecs) for ST-2 and ST-2+ include parameters that are not currently available to ATM virtual circuits through Q.2931 and UNI 3.1. The mapping of packet rate to cell rate, packet delay to cell delay, and other translatable QoS parameters is described in section 5. However, the ST-2 flowspecs also include parameters like accumulated end-to-end delay and accumulated jitter. These parameters assume that the SCMP messages follow the same path as the data. Depending on the signaling mechanism chosen, this may not be true with ATM and thus certain QoS parameters may be rendered useless.

It should also be noted that since ST-2 connections are simplex, all QoS parameters are specified separately for each direction of data transfer. Thus two connections and two QoS negotiations are required for a duplex connection. To take advantage of the full duplex nature of point-to-point ATM connections, special multiplexing of ST connections would be required by ST-2 agents.

3.6 Use of HELLOs

Both ST-2 and ST-2+ support HELLO messages. HELLOs are intended to assure that the neighboring agent is alive. Failure to respond to a HELLO indicates that the connection is down and that the reservation for that particular link should be freed.

While the ATM network will notify an ST-2 agent if the network connection is down, there is still the possibility that the connection is intact but that the ST-2 agent itself is down. Knowledge of the neighboring agent's status is increasingly important when multiple ST-2 connections share virtual circuits, when the neighboring agents are routers, and when there are multiple dedicated virtual circuits between agents.

As such, HELLO is a desirable feature. Note that some signaling schemes (section 4), provide less than optimal support for HELLO.

4.0 Reservation Signaling with ATM

Use of Permanent Virtual Circuits (PVCs) for reservation signaling presents no problem for ST-2, ST-2+, or RSVP. Each circuit is considered to be a dedicated link to the next hop. If the PVCs are to be shared, reservation protocols can divide and regulate the bandwidth just as they would with any other link type.

Where ATM connections become more interesting is when the ATM network takes on the role of an extended LAN or internet. To do this, Switched Virtual Circuits are used to establish dynamic connections to various endpoints and routers. The ITU-TS Q.2931 SETUP message is used to request a connection from the network with specific bandwidth and QoS requirements, and a CONNECT message is received by the origin to indicate that connection establishment is complete.

For IP over ATM and LAN Emulation, SVCs are established between endpoints and data traffic for a given destination shares the SVCs. There is no mechanism to allow specific QoS guarantees for the traffic, nor is there a mechanism to set up virtual circuits with specific bandwidth and QoS for a particular type of traffic. This is what reservation protocols will attempt to do. The goal is to use reservations to request establishment of individual virtual circuits with matching bandwidth and QoS for each reservation. This will guarantee the requirements of the application while taking full advantage of the ATM network's capabilities.

There are four possible mechanisms to perform reservation signaling over ATM:

- 1) Embedding reservation signaling equivalents within the ATM Q.2931 controls.
- 2) Signaling in-band with the data.
- 3) Signaling over dedicated signaling VCs.
- 4) Implicitly sharing existing VCs for IP over ATM or LAN Emulation.

Note that ATM circuits are not necessarily reliable. As such, the reliability mechanisms provided by SCMP must be maintained to assure delivery of all reservation signaling messages.

4.1 Embedded Reservation Signaling Equivalents within ATM Q.2931 Controls

The basic idea in embedding reservation signaling within the ATM controls is to use the Q.2931 SETUP and CONNECT messages to establish both reservations and dedicated data paths (virtual circuits) across the ATM network. This eliminates the need for dedicated signaling channels, in-band signaling, or out of band mechanisms to communicate between endpoints. Since SETUP and CONNECT include bandwidth and QoS information, the basic concept is sound. In fact, this approach will speed network connection by preventing multiple passes at establishing a reservation and associated connection. This normally results from the fact that most higher layer protocols (network and transport) first require a link to signal their connection requirements. As such, with ATM, the ATM virtual circuit must be established before the network and/or transport protocols can do their own signaling.

Embedded reservation signaling allows the reservation information to be carried in the SETUP and CONNECT messages, allowing the reservation protocol to do its signaling simultaneously with the ATM signaling.

[7] describes a clever way of combining the reservation signaling with the ATM control plane signaling for ST-2. This 'simultaneous connection establishment' process will optimize the establishment of circuits and minimize connection setup time while simultaneously eliminating unnecessary network layer signaling in ST-2. To be effective, [7] requires enhancements to Q.2931 signaling and to the ST-2 protocol implementations. In addition, it currently only applies to point-to-point connections and will not work with multipoint largely due to the simplex nature of multipoint communication in current ATM implementations.

Implementation of multicast for Embedded Reservation Signaling is done as described above: the reservation agent at the edge of the ATM network must create point-to-point virtual circuits for each target that is directly connected to the ATM network, and for each router

that supports downstream targets. This ensures two-way signaling between targets and the origin.

Signaling itself is quite simple:

CONNECT maps directly to one or more (multicast) Q.2931
SETUPS and CONNECTs.
ACCEPT maps directly to Q.2931 CONNECTACK.
CHANGE/CHANGE REQUEST are not supported.
DISCONNECT maps directly to Q.2931 RELEASE.
HELLOs are not needed.

Unfortunately, the flowspec in the reservation protocol CONNECT message cannot be passed across the ATM network in the signaling messages and thus must be regenerated by the receiving agent.

In addition, User Data, which can be sent in most SCMP messages cannot be supported without substantial changes to current Q.2931 signaling.

One of the additional complexities with embedding the reservation signaling occurs in heterogeneous networks. Since ATM signaling only operates point to point across the ATM network itself, if the endpoints reside on other types of networks or subnets, the routers at the edge of the ATM networks must generate and regenerate endpoint-based signaling messages on behalf of the host reservation agents. In particular, CONNECT and ACCEPT messages and their associated flowspecs must be regenerated. Refer to Section 5 for details on the QoS mappings and on which QoS parameters can be recreated for the generated flowspecs.

This approach is worth revisiting as an optimal signaling method in pure ATM network environments once ATM signaling capabilities expand.

However, for heterogeneous networks, other signaling mechanisms may be more appropriate.

4.2 In-Band Reservation Signaling

In-Band Reservation Signaling is the easiest signaling mechanism to implement. When the applications requests a reservation, the reservation agent simply sets up ATM virtual circuits to the endpoints with the QoS specified in the CONNECT request. When ACCEPTed, all subsequent data transmissions proceed on the virtual circuits.

Once again, to support multicast, the reservation agent must create individual point-to-point virtual circuits to the targets which are

directly connected to the ATM network, as well as to routers which can access downstream targets.

Since signaling is done in-band, all reservation signaling messages can be passed between agents. However, some minimal additional bandwidth must be allocated in the Q.2931 SETUP to allow for the signaling messages themselves.

Note that the primary disadvantage to In-Band Reservation Signaling is the fact that it does not make use of the multipoint capabilities of ATM and will thus overreserve ATM network bandwidth and create a larger than necessary number of virtual circuits.

4.3 Dedicated Reservation Signaling Virtual Circuits

One mechanism that can be used to take advantage of the full data transmission capabilities of ATM networks is to use Dedicated Virtual Circuits for reservation signaling. This guarantees a two-way signaling pipe between the endpoints in a connection while enabling the data transmission to take advantage of the multipoint capabilities of ATM. Data and Signaling are done over separate virtual circuits.

When an application requests a reservation, the reservation agent reviews the list of targets in the CONNECT request. For any targets which have no current signaling virtual circuits established, the agent establishes UBR (unspecified bit rate) virtual circuits and forwards the CONNECT message to the targets over these virtual circuits. ATMARP is used to resolve any endpoint addresses. For any targets for which there already exist signaling virtual circuits, the agent simply forwards the CONNECT message over the existing virtual circuit.

Once an ACCEPT message is received, the agent issues a Q.2931 SETUP to the associated target. Upon receipt of a CONNECTACK, data can begin to flow. As additional ACCEPTs are received, the Q.2931 ADDPARTY message is used to add a target to the multicast and multipoint connection. Depending on the cause of any ADDPARTY failure, the agent may attempt to establish a dedicated point-to-point virtual circuit to complete the multicast group.

DISCONNECT requests result in Q.2931 DROPPARTY messages and will cause a member to be dropped from a multicast and multipoint connection. When all targets are dropped from a multipoint connection, a RELEASE can be issued to take down the virtual circuit.

Signaling virtual circuits are shared among reservations while data circuits are dedicated to a particular reservation. Once all

reservations to a given endpoint are terminated, the signaling virtual circuit to that endpoint can be RELEASED.

Note that this approach would allow the NSAP address to be passed as user data in the ACCEPT message to enable a kernel-based reservation protocol to establish the dedicated data circuit. In addition, because the connectivity to the endpoint is identical to that of the data circuit, this approach assures the fact that accumulated information in the flowspecs retains its validity.

4.4 Reservation Signaling via IP over ATM or LAN Emulation

As described in the previous section, it would be possible to set up unique SVCs for SCMP signaling, however, since the streaming, connection-oriented data transport offered by ST-2 is intended to be complementary to IP and other connectionless protocol implementations, it would be simpler and more elegant to simply use classical IP over ATM (RFC 1577) mechanisms, or to use LAN Emulation. The widespread deployment of IP over ATM and LAN emulation in host-based ATM drivers, and the assumption that most host systems will be running applications that do not need specific QoS and bandwidth provisioning, makes this the most straightforward (if not performance optimal) solution for signaling. Once an end-to-end acceptance of a reservation request is completed via normal LAN or IP transmission, then a unique direct virtual circuit can be established for each data flow.

If LAN Emulation is used, as long as the ST-2 implementation allows for different paths for SCMP and data, there would be no changes to the signaling mechanisms employed by the reservation agent.

For IP over ATM, all SCMP messages would be encapsulated in IP as described in both RFC 1190 and RFC 1819. This is required because current ATM drivers will not accept IPv5 packets, and most drivers do not provide direct access to the shared signaling virtual circuits used for IP.

In either case, LAN Emulation or IP over ATM, the reservation agent would handle SCMP messages as it normally does. However, once the first ACCEPT is received for a reservation request, a dedicated virtual circuit is established for the data flow. Subsequent ACCEPTs will result in the use of ADDPARTY to add multicast targets to the multipoint virtual circuit. In fact, processing of multipoint/multicast is identical to that described in section 4.3.

Once again, the use of an out-of-band signaling mechanism makes it possible to carry the NSAP address of the target in the ACCEPT message.

One potential drawback to using LAN Emulation or SCMP messages encapsulated in IP over ATM, is the fact that there is no guarantee that the connectivity achieved to reach the target via signaling has any relationship to the data path. This means that accumulated values in the flowspec may be rendered useless.

In addition, it is possible that the targets will actually reside outside the ATM network. That is, there may be no direct ATM access to the Targets and it may be difficult to identify ATM addresses of the associated ATM connected routers. This approach will involve some additional complexity in routing to the targets. However, since ST-2 is intended to run with IP, if ATM vendors would accept IPv5 packets or would allow direct access to the IP over ATM signaling virtual circuits, this approach would be optimal in minimizing the number of virtual circuits required.

4.5 Summary of Reservation Signaling Approaches

Embedded Reservation Signaling (section 4.1) is ideal for homogeneous ATM connections, but requires extensions to existing ATM signaling to support multipoint connections. In-Band Reservation Signaling (section 4.2) is the easiest to implement, but cannot employ multipoint connections either.

Perhaps the simplest way to do this is similar to what is suggested in [6]: separate the reservation signaling from the actual data flows, mapping the data flows directly to ATM circuits while doing the signaling separately.

While there is significant complexity in doing this for IP traffic and RSVP, the ST2 protocols lend themselves to this quite well. In fact, because SCMP reservation signaling results in streaming, multicast connections, the 'Shortcut' mechanism described in [6], which can bypass routers where direct ATM connections are possible, is automatically available to ST2 streams.

Using Reservation Signaling over LAN Emulation or IP over ATM (section 4.4) is one multipoint-capable approach to implement in hosts since most ATM drivers shipping today provide both IP over ATM and LAN Emulation, as well as associated address resolution mechanisms. However, it is not complete in its ability to accurately depict flowspec parameters or to resolve host ATM addresses. In addition, to be optimal, ATM vendors would either have to support IPv5 in their drivers or allow direct access to the IP signaling virtual circuits. Thus the current ideal approach to implementation of the ST2 protocols over ATM is to use shared Dedicated Reservation Signaling Virtual Circuits (section 4.3) for signaling of reservations, and then to establish appropriate multipoint ATM

virtual circuits for the data flows.

5.0 Mapping of Reservation QoS to ATM QoS

QoS negotiation in ST-2 (and ST-2+) is done via a two-way negotiation.

The origin proposes a QoS for the connection in a Flow Specification (Flowspec) associated with the CONNECT message. Most of the network-significant QoS parameters in the Flowspec include both a minimum and a desired value. Each ST agent along the path to the Target validates its ability to provide the specified QoS (at least the minimum value for each), updates certain values in the Flowspec, and propagates the CONNECT until it reaches the Target. The Target can either ACCEPT the Flowspec or REFUSE it if it cannot meet at least the minimum QoS requirements. Negotiation takes place as part of the process in that the Target can specify changes to the desired QoS values as long as the new value meets at least the minimum requirements specified by the Origin system. In addition, both the Target and the Origin can assess actual network performance by reviewing the values that are accumulated along the path.

The primary Reservation QoS parameters that impact an ATM network are:

ST-2 (RFC 1190)	ST-2+ (RFC 1819)
Desired PDU Bytes, Limit on PDU Bytes (minimum).	Desired Message Size, Limit on Message Size.
Desired PDU Rate, Limit on PDU Rate (minimum). Minimum Transmission Rate in Bytes.	Desired Rate, Limit on Rate.
Limit on Delay (maximum). Maximum Bit Error Rate.	Desired Delay, Limit on Delay.
Accumulated Delay. Accumulated Delay Variance (Jitter).	

Q.2931 ATM signaling offers the following QoS parameters:

- Cumulative Transit Delay,
- Maximum End to End Transit Delay.
- Forward Peak Cell Rate (PCR),
- Backward Peak Cell Rate (PCR).

- Forward Maximum CPCS-SDU size,
- Backward Maximum CPCS-SDU size.

- Forward QoS Class,
- Backward QoS Class.

- B-LLI (one byte user protocol information).

As previously noted, reservation protocols (ST and RSVP) make QoS reservations in one direction only. Thus, depending on the type of signaling used (see Section 4), the 'Backward' ATM parameters may not be useful. In particular, if Multipoint ATM connections are used to map multicast reservations, these parameters are not available.

However, it would be possible to implement a multiplexing scheme to enable reservations to share bi-directional point-to-point ATM connections if the reservation agent creates a split/merge point at the ATM boundary and sets up only point-to-point VC connections to targets.

The CPCS-SDU parameters are AAL Parameters which are used by the AAL entity to break packets into cells. As such, these parameters are not modified by the network and could conceivably be used for additional end-to-end signaling, along with the B-LLI.

Finally, QoS Class is somewhat limited in its use and implementation. While IP over ATM recommends use of Class 0 (Unspecified QoS), this is not sufficient for guaranteed connections. Instead, Class 1 with CLP=0 will provide at least minimum QoS services for the traffic.

5.1 CPCS-SDU Size Computation

The CPCS-SDU size computation is the easiest QoS mapping. Since ST-2 does not require a Service Specific Convergence Sublayer (SSCS), if AAL 5 is used, the ST packet size plus 8 bytes (for the AAL 5 Trailer) will be the CPCS-SDU size. Note that the ST-2 packet size also includes an 8-byte header for ST-2. Thus the CPCS-SDU size is:

$$\text{CPCS-SDUsize} = \text{PDUbytes} + 8 + 8.$$

For ST-2+, the header is larger than for ST-2, so the CPCS-SDU size is:

$$\text{CPCS-SDUsize} = \text{PDUbytes} + 12 + 8.$$

5.2 PCR Computation

The Peak Cell Rate (PCR) computation is only slightly more complex. The PCR will be the peak packet rate divided by the ATM payload size.

Since PDU rates in ST-2 are specified in tenths of packets per second, AAL 5 requires an 8 byte trailer, and the ATM payload size is 48 bytes, the computation for PCR proceeds as follows:

The requested maximum byte transmission rate for ST-2 is:

$$\text{PDUbytes} * \text{PDUrate} * 10.$$

Accounting for the AAL 5 and ST headers, the maximum byte rate is:

$$\text{Bytes per second} = (\text{PDUbytes} + 8 + 8) * \text{PDUrate} * 10.$$

Translating into cells and eliminating the possibility of a fractional PDU:

$$\text{PCR} = ((\text{PDUbytes} + 8 + 8 + 48) / 48) * \text{PDUrate} * 10.$$

For ST-2+, not only is the header size 12 bytes, but the Rate is in messages per second, not tenths of packets per second. Thus, the PCR for ST-2+ is:

$$\text{PCR} = ((\text{PDUbytes} + 12 + 8 + 48) / 48) * \text{PDUrate}.$$

5.3 Maximum End to End Transit Delay.

The End to End Transit Delay is a little more complex. The requested end to end delay must account for not only the PDU size as requested by the user, but the additional 8-byte AAL 5 header as well. The translation of the user-requested LimitOn Delay is preserved as long as the delay computation is based on the CPCS-SDU size instead of the PDU size.

In addition to the end to end delay introduced by the ATM network, there is additional delay created by the fragmentation of packets. Reassembly of these packets can only be accomplished at the rate at which they are received. The time (in milliseconds) required to receive a cell (inter-cell arrival time) is:

$$T = 1000 / \text{PCR}.$$

The number of cells in a CPCS-SDU is:

$$C = (\text{CPCS-SDUsize} + 48) / 48.$$

Thus the delay for a packet is:

$$\text{LimitonDelay} = (C - 1) * T + \text{MaxCellTransitDelay}.$$

Therefore, the requested Maximum End to End Transit delay is:

$$\text{MaxCellTransitDelay} = \text{Limiton Delay} - (C-1) * T.$$

5.4 Maximum Bit Error Rate

Q.2931 signaling does not offer the ability to directly specify the requested bit error rate or a corresponding cell error rate. Instead, this service is supposed to be offered through selection of QoS class.

Since these classes have few actual implementations, at this time, there is no effective mapping for bit error rate.

5.5 Accumulated Mean Delay

ST allows accumulation of the Mean Delay generated by each ST agent node and intervening circuits. With an ATM circuit each agent should factor in the overhead of the ATM connection. The delay associated with the ATM circuit is reflected in the Q.2931 CONNECT message as the Cumulative Transit Delay. Since this is a cell-based computation, the delay experienced for an ST packet, including the CPCS-SDU header and ST header is, as computed in Section 5.3:

$$\text{Delay} = (C - 1) * T + \text{CummulativeTransit Delay}.$$

5.6 Accumulated Delay Variance (Jitter)

Cell Delay Variance is not currently available as a Q.2931 parameter.

Thus, we can assume that the reassembly of cells into packets will be consistent, since the cell transmission rate should be constant for each packet. As such, except as noted by the specific ATM service, the ST agent should use its standard mechanisms for tracking packet arrival times and use this for Accumulated Delay Variance.

6.0 Data Stream Transmission

Once virtual circuits for data transmission are established through one of the mechanisms described in section 4, the ST data must be

transmitted over the connection. RFC 1483 describes mechanisms for encapsulating packet transmissions over AAL5. While the LLC encapsulation could be used, it is not necessary. If it is used, the computations in section 5 should be redone to include the LLC headers in addition to the AAL5 trailer currently used. These new values should be substituted for the QoS values in the SETUP message.

Instead, ST data packets can be encapsulated in standard AAL5 format with an 8 byte trailer and sent directly over the data virtual circuit. The mechanisms for computing the QoS values in the SETUP message are described in section 5.

7.0 Implementation Experience and Conclusions

All of the signaling mechanisms described in Section 4 were implemented and tested in a mixed ATM network/routed LAN environment.

Initially it appeared that the best approach was to do signaling via IP over ATM or LANE. However, because it required IP encapsulation of the SCMP packets (for IP over ATM), and because some applications use the accumulated values in the flowspecs (which are not guaranteed to be accurate in LANE and IP/ATM), using virtual circuits dedicated to SCMP signaling turned out to be the best implementation for taking full advantage of the ATM features.

Also, the issue of mapping ATM address to E.164 NSAP addresses was resolved through an external signaling mechanism (the User Data field of the ST-2 CONNECT and ACCEPT messages). It appears that ATM vendors need to implement a consistent addressing mechanism throughout their interfaces.

From a performance point of view, using ST over ATM provided more than triple the performance of raw IP. The differences became increasingly clear as more simultaneous applications were run. This resulted in dedicated virtual circuits for the ST traffic while the IP traffic suffered (saw inconsistent performance) over shared circuits. Even more dramatic were results in mixed network environments where all traffic shared the same LAN/router connections, and, when both IP and ST traffic was sent, the ST traffic maintained its quality while the IP traffic saw increasing variation in performance.

Clearly, using a connection-oriented, origin-oriented reservation protocol to provide consistent end-to-end guaranteed QoS and bandwidth in mixed ATM/internet environments is not only feasible, it results in dramatic performance and quality improvements for transmission of realtime traffic.

8.0 Security Considerations

This memo raises no security considerations. However, with their connection-oriented and origin controlled natures, ST-2 and ST-2+ lend themselves to better internet security. Discussion of this is beyond the scope of this document.

9.0 References

- [1] Laubach, M., "Classical IP and ARP over ATM", RFC 1577, Hewlett Packard Laboratories, December, 1993.
- [2] Borden, M., Crawley, E., Davie, B., and S. Batsell, "Integration of Real-time Services in an IP-ATM network Architecture", RFC 1821, August 1995.
- [3] Braden, R., Zhang, L., Estrin, D., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP Version 1 Functional Specification)", Work in Progress, November 1995.
- [4] Topolcic, C., "Experimental Internet Stream Protocol: Version 2 (ST-II)", RFC 1190, October 1990.
- [5] DelGrossi, L., and L. Berger, "Internet SStream Protocol Version 2+", RFC 1819, July 1995.
- [6] V. Firoiu, R. Guerin, D. Kandlur, A. Birman "Provisioning of RSVP-based Services over a Large ATM Network", IBM T.J. Watson Research Center, October 1995.
- [7] S. Damaskos, A. Anastassios Gavras, "Connection Oriented Protocols over ATM: A Case Study", German National Research Corporation for Mathematics and Data Processing (GMD) and Research Centre for Open Communications Systems (FOKUS), February 1994.
- [8] Heinanen, J., "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 1483, July 1993.
- [9] M. Graf, T. Kober, H. Stuttgen, "ST-II over ATM Implementation Issues", IBM European Networking Center, October 1995.

10.0 Author's Address

Steve Jackowski
NetManage Incorporated
269 Mt. Hermon Road, Suite 201
Scotts Valley, Ca 95066

Phone: (408) 439-6834
Fax: (408) 438-5115
EMail: Stevej@NetManage.com

