

IPng Support for ATM Services

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

This document was submitted to the IETF IPng area in response to RFC 1550. Publication of this document does not imply acceptance by the IPng area of any ideas expressed within. Comments should be submitted to the big-internet@munari.oz.au mailing list.

Executive Summary

This white paper describes engineering considerations for IPng as solicited by RFC 1550 [1]. IPng should provide support for existing and emerging link technologies that it will be transported over. Link technologies like Ethernet simply multiplex traffic from upper layer protocols onto a single channel. "Sophisticated" link technologies like ATM are emerging in the marketplace allowing several virtual channels to be established over a single wire (or fiber) potentially based on an applications' network performance objectives.

Support for both "sophisticated" (ATM) and existing link technologies needs to be considered in an IPng candidate. End-to-end applications will communicate through a network where IPng packets travel across subnetworks such as Ethernet and Hippi and also more "sophisticated" link levels such as ATM. Though initial support for IPng over ATM subnetworks will not facilitate a virtual circuit per application, the hooks to provide such a mapping should be in place while also maintaining support for the transport of IPng packets across conventional subnetworks. Application support for QOS-based link level service requires that the following types of ATM information be mappable (or derivable) from the higher level protocol(s) such as IPng: source and destination(s) addresses, connection quality of service parameters, connection state, and ATM virtual circuit identifier. Some of these mappings may be derivable from information provided by proposed resource reservation protocols supporting an integrated services Internet [4]. However, the ATM virtual circuit identifier should be efficiently derivable from IPng packet

information.

An IPng candidate should provide evidence that the mapping from an applications' IPng packets to ATM virtual circuit(s) can be accomplished in a heterogeneous Internet architecture keeping in consideration the gigabit/sec rates that IPng/ATM subnetworks will eventually be operating at.

1. Introduction

This paper describes parameters that are needed to map IPng (or any protocol operating above the link level) to ATM services. ATM is a "sophisticated" link level technology which provides the potential capability for applications at the TCP/UDP level to map to a single ATM virtual circuit for transport across an ATM network(s) customized to the network performance and traffic requirements for that application. This is a step above many of today's existing link technologies which can only support a single level of network performance that must be shared by all applications operating on a single endpoint.

The future Internet will be comprised of both conventional and "sophisticated" link technologies. The "sophisticated" features of link layers like ATM need to be incorporated into an internet where data travels not only across an ATM network but also several other existing LAN and WAN technologies. Future networks are likely to be a combination of subnetworks providing best-effort link level service such as Ethernet and also sophisticated subnetworks that can support quality of service-based connections like ATM. One can envision data originating from an Ethernet, passing through an ATM network, FDDI network, another ATM network, and finally arriving at its destination residing on a HIPPI network. IPng packets will travel through such a list of interconnected network technologies as ATM is incorporated as one of the components of the future Internet.

To support per application customizable link level connections, four types of ATM information should be derivable from the higher level protocol(s) like IPng. This ATM information includes: source and destination ATM addresses, connection quality of service parameters, connection state, and an ATM virtual circuit identifier which maps to a single IPng application (i.e., single TCP/UDP application). Some of these mapping could potentially be derivable through information provided by proposed resource reservation protocols supporting an integrated services Internet [4]. However, the ATM virtual circuit identifier needs to be efficiently mappable from IPng packet information.

Organization of this white paper is as follows. First the characteristics of ATM are described focusing on functions that are not provided in today's LAN technologies. This section provides background information necessary for the following section describing the parameters needed to map IPng services to ATM services.

2. Terminology

In this white paper, the term "application" refers to a process or set of collective processes operating at the TCP/UDP level or above in the protocol stack. For example, each instance of "telnet" or "ftp" session running on an end station is a distinct application.

3. Characteristics of ATM Service

ATM has several characteristics which differentiates it from current link level technologies. First of all, ATM has the capability of providing many virtual channels to transmit information over a single wire (or fiber). This is very similar to X.25, where many logical channels can be established over a single physical media. But unlike X.25, ATM allows for each of these channels or circuits to have a customizable set of performance and quality of service characteristics. Link level technologies like Ethernet provide a single channel with a single performance and quality of service characteristic. In a sense, a single ATM link level media appears like an array of link level technologies each with customizable characteristics.

ATM virtual circuits can be established dynamically utilizing its signaling protocol. ATM signaling is a source initiated negotiation process for connection establishment. This protocol informs elements in the network of the characteristics for the desired connection. ATM signaling does not provide any guidelines for how network elements decide whether it can accept a call or where a signaling request should be forwarded if the end destination (from the link level perspective) has not been reached. In short, ATM signaling does not support any routing functionality of network admission control.

ATM signaling establishes a "hard state" in the network for a call. "Hard state" implies that the state of a connection in intermediate switching equipment can be set and once established it will be maintained until a message is received by one of the ends of the call requesting a change in state for the connection [2]. As a result, an ATM end system (this could be a workstation with an ATM adapter or a router with an ATM interface) receives guaranteed service from the ATM network. The ATM network is responsible for maintaining the connection state. The price the ATM termination points pay for this guarantee is the responsibility of changing the state of the

connection, specifically informing the ATM network to establish, alter, or tear-down the connection.

Each ATM end point in a network has an ATM address associated with it to support dynamic connection establishment via signaling. These addresses are hierarchical in structure and globally unique [3]. As a result, these addresses are routable. This allows ATM networks to eventually support a large number of ATM endpoints once a routing architecture and protocols to support it become available.

The ATM User-Network Interface (UNI) signaling protocol based on ITU-TS Q.93B allows many different service parameters to be specified for describing connection characteristics. [3] These parameters can be grouped into several categories: ATM adaptation layer (AAL) information, network QOS objectives, connection traffic descriptor, and transit network selector. The AAL information specifies negotiable parameters such as AAL type and maximum packet sizes. The network QOS objectives describe the service that the ATM user expects from the network. Q.93B allows for one of five service classes to be selected by the ATM user. The service classes are defined as general traffic types such as circuit emulation (class A), variable bit rate audio and video (class B), connection-oriented data transfer (class C), connectionless data transfer (class D), best effort service (class X), and unspecified [3]. Each of these categories are further specified through network provider objectives for various ATM performance parameters. These parameters may include cell transfer delay, cell delay variation, and cell loss ratio. The connection traffic descriptor specifies characteristics of the data generated by the user of the connection. This information allows the ATM network to commit the resources necessary to support the traffic flow with the quality of service the user expects. Characteristics defined in the ATM Forum UNI specification include peak cell rate, sustainable cell rate, and maximum and minimum burst sizes [3]. Lastly, the transit network selection parameter allows an ATM user to select a preferred network provider to service the connection [3].

4. Parameters Required to Map IPng to ATM

There are several parameters required to map ATM services from a higher level service like IPng. These ATM parameters can be categorized in the following manner: addressing parameters, connection QOS-related parameters, connection management information, and ATM virtual circuit identifier. The first three categories provide support for ATM signaling. The last parameter, a connection identifier that maps IPng packets to ATM virtual circuits, provides support for an ATM virtual circuit per application when the end-to-end connection travels across an ATM subnetwork(s) (this does not assume that ATM is the only type of subnetwork that this connection

travels across). Below, mapping issues for each of these parameters will be described.

4.1. Addressing

ATM supports routable addresses to each ATM endpoint to facilitate the dynamic establishment of connections. These addresses need to be derived from a higher level address such as an IPng address and IPng routing information. This type of mapping is not novel. It is a mapping that is currently done for support of current IP over link technologies such as Ethernet. An IP over ATM address resolution protocol (ARP) has been described in the Internet Standard, "Classical IP over ATM" [5]. In addition, support for IP routing over large ATM networks is being worked in the IETF's "Routing over Large Clouds" working group.

4.2. Quality of Service

As described in section 3, an ATM virtual circuit is established based upon a user's traffic characteristics and network performance objectives. These characteristics which include delay and throughput requirements can only be defined by the application level (at the transport level or above) as opposed to the internetworking (IPng) level. For instance, a file transfer application transferring a 100 MB file has very different link level performance requirements than a network time application. The former requires a high throughput and low error rate connection whereas the latter could perhaps be adequately serviced utilizing a best-effort service. Current IP does not provide much support for a quality of service specification and provides no support for the specification of link level performance needs by an application directly. This is due to the fact that only a single type of link level performance is available with link technologies like Ethernet. As a result, all applications over IP today receive the same level of link service.

IPng packets need not explicitly contain information parameters describing an application's traffic characteristics and network performance objectives (e.g., delay = low, throughput = 10 Mb/s). This information could potentially be mapped from resource reservation protocols that operate at the IP (and potentially IPng) level [4].

4.3. Connection Management

The establishment and release of ATM connections should ultimately be controlled by the applications utilizing the circuits. As described in section 3, ATM signaling establishes a "hard state" in the network which is controlled by the ATM termination points [2]. Currently, IP

provides no explicit mechanism for link level connection management. Future support for link level connection management could be accomplished through resource reservation protocols and need not necessarily be supported directly via information contained in the IPng protocol.

4.4. Connection Identifier

A mapping function needs to exist between IPng packets and ATM so that application flows map one-to-one to ATM virtual circuits. Currently, application traffic flows are identified at the transport level by UDP/TCP source and destination ports and IP protocol identifiers. This level of identification should also be available at the IPng level so that information in the IPng packets identify an application's flow and map to an ATM virtual circuit supporting that flow when the IPng packets travels across an ATM subnetwork(s).

Using the current IP protocol, identifying an application's traffic flow requires the combination of the following five parameters: source and destination IP addresses, source and destination UDP/TCP ports, and IP protocol identifier. This application connection identifier for IP is complex and could potentially be costly to implement in IP end stations and routers. The IPng connection identifier should be large enough so that all application level traffic from an IPng end point can be mapped into the IPng packet. Currently, ATM provides 24 bits for virtual circuit identification (VPI and VCI). This provides sufficient capacity for 2^{24} (16,777,216) connections [6]. The actual number of bits that are used for the ATM virtual circuit however is established through negotiation between the ATM endpoint and ATM network. This number is useful as an upper bound for the number of mappings that are needed to be supported by IPng.

An IPng candidate should be able to identify how IPng packets from an application can map to an ATM virtual circuit. In addition, this mapping should be large enough to support a mapping for every IPng application on an end system to an ATM virtual circuit. Careful consideration should be given to complexity of this mapping for IPng to ATM since it needs to eventually support gigabit/sec rates.

References

- [1] Bradner, S., and A. Mankin, "IP: Next Generation (IPng) White Paper Solicitation", RFC 1550, Harvard University, NRL, December 1993.
- [2] Clark, D., "The Design Philosophy of the DARPA Internet Protocols", Proc. ATM SIGCOMM '88, August 1988.
- [3] "ATM User-Network Interface Specification, Version 3.0", ATM Forum, September 10, 1993.
- [4] Zhang, L., Estrin, D., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification", Work in Progress, October 1993.
- [5] Laubach, M., "Classical IP and ARP over ATM", RFC 1577, Hewlett-Packard Laboratories, January 1994.
- [6] "Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements", Bellcore Technical Advisory TA-NWT-001113, Issue 1, June 1993.

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Christina Brazdziunas
Bellcore
445 South Street
Morristown, NJ 07960

Phone: (201) 829-4173
EMail: crb@faline.bellcore.com

