Technical Information Bulletin 98-2

Selected ATM/Internet Protocol (IP) Technical Interface

Considerations

January 1998

ABSTRACT

Connection-oriented and connectionless transport modes are the two fundamental types of service used to support the flow of information in today's communications networks. Each of the two transport modes is designed to support specific applications and meet specific quality of service (QoS) requirements. However, because of its capability to support high data rates, provide gains in bandwidth use efficiency, and accommodate a broad range of traffic from diverse sources, Asynchronous Transfer Mode (ATM)-a connection-oriented transport mode-is gaining broad acceptance as the preferred transport mode for both connection-oriented and connectionless communications services across multiple public networks. This report provides: (1) a general overview of salient features of connection-oriented ATM architectures and of the Internet Protocol (IP) used to support the connectionless transfer of data; and (2) selected technical interface considerations with potential impact on the transfer of IP data over ATM virtual connections.

TABLE OF CONTENTS

LIST OF FIGURES

EXECUTIVE SUMMARY

PURPOSE

This report presents the results of an examination of selected technical interface considerations with potential impact on the transfer of connectionless Internet Protocol (IP) data over connection-oriented Asynchronous Transfer Mode (ATM) networks.

BACKGROUND

The direct operation of IP, a connectionless protocol over ATM-a connection-oriented transport mode-is one of several application areas currently being examined by both the ATM Forum's Multiprotocol Over ATM (MPOA) Working Group and the Internet Engineering Task Force (IETF)'s IP Over ATM Working Group. The current ATM Forum solution for the operation of IP over ATM requires that an end-to-end virtual channel/circuit (VC)-a communications channel that provides for the sequential transport of ATM cells-be established for the transport of IP packets. There is a certain amount of routing overhead and delay associated with the establishment of each VC. Currently, both the ATM Forum and the IETF are considering approaches to reducing the routing overhead and delay associated with the transmission of IP over ATM.

ATM/INTERNET PROTOCOL DATA TRANSFER TECHNICAL CONSIDERATIONS

The transfer of IP data over ATM requires the translation of IP from a connectionless data transfer mode to a connection-oriented ATM transfer mode. Major differences exist in the makeup, interfacing, and operation of connection-oriented and connectionless

telecommunications protocols. These differences include different addressing schemes, different routing procedures, different control measures, and other factors involving the general complexities of the two modes. The IETF's IP over ATM Working Group charter identifies the following as initial areas in which it plans to develop experimental protocols for the operation of internetwork protocols over ATM networks: (1) encapsulation, (2) multicasting, (3) addressing, (4) address resolution, (5) call setup, and (6) network management.

ENCAPSULATION

Encapsulation is enclosing data formatted for one protocol within another protocol in order to transmit the data across a network for which the original protocol was not designed. It is one method of supporting multiple protocols across a single connection. For non-similar protocols, encapsulation provides for the creation of an envelope into which other protocols are placed. The data within the envelope is transported transparent to the network. At the destination endpoint, the data is removed from the envelope and returned to its original form. Encapsulation of IP datagrams for transmission over ATM is accomplished by: (1) the segmentation by ATM Adaptation Layer-5 (AAL-5) of protocol data units (PDUs) received from upper layer protocols (ULPs) into 48-byte packets; (2) mapping the 48-byte packets into the information fields of the ATM cell; and (3) forwarding the information fields sequentially to the ATM layer to generate 53-byte ATM cells. The size of an IP datagram in most instances will be many times the size of an ATM cell. Consequently, following segmentation, it is expected that a single IP datagram will correspond to a number of ATM cells. The ATM Forum is currently considering approaches to reducing routing overhead and delay associated with the transmission of IP over ATM.

MULTICASTING

Multicasting is the process whereby a source interface to a group sends a single PDU to more than one selective destination simultaneously, using a single local transmit operation. The IP multicast model is receiver initiated. Receivers wishing to subscribe to an IP multicast group inform their local routers. In ATM networks, multicast connections are supported by a collection of sender (root)-initiated point-to-multipoint VCs and associated endpoints. This arrangement requires that the sender know each intended recipient, and explicitly establish a VC between itself as the root and each recipient as a leaf node. IP multicasting over ATM networks requires a mechanism for mapping IP group addresses to corresponding ATM addresses. In the ATM Forum's MPOA model, the mechanism used is the Multicast Address Resolution Server (MARS). The MARS acts as a registry, associating layer 3 multicast group identifiers with the ATM interfaces representing the group's members. A potential concern with the MPOA model multicast arrangement is the requirement to establish dedicated VCs between the root and each

recipient leaf node.

Considerable discussion is currently ongoing within the IETF concerning ways by which the overhead associated with the establishment and maintenance of separate multicast VCs might be avoided. The current focus of the discussion is on the relative merits of using "VC meshes" versus multicast servers (MCSs) to support network layer (e.g., IP) multicasting over ATM. Both seem to offer advantages and disadvantages. With the multicast VC mesh, each source establishes independent point-to-multipoint VCs to the set of leaf nodes it wishes to send messages. Interfaces for the leaf nodes originate and terminate VCs, as appropriate, for each active leaf nodethe term "VC mesh" is used to describe the resulting crisscross VC pattern.

ADDRESSING

Addressing is the assignment of a sequence of characters in a message, datagram, or cell header to designate its origin and destination. Each device within a network must have a unique and identifiable address in order to receive and transmit messages. Addressing in an integrated IP/ATM environment is currently being considered by the IETF Networking Working Group and the ATM Forum Technical Committee as a major issue. In a classical IP environment, the source and destination address fields of the IP datagram header contain the addressing information needed to route datagrams in a connectionless network. With IP over ATM, the IP addressing information is encapsulated within the ATM cell along with the datagram and becomes transparent to the ATM network. Consequently, IP over ATM requires a mechanism for mapping IP addresses to corresponding ATM addresses. For classical IP over ATM, an ATM address resolution protocol (ATMARP) server provides the mapping mechanism.

As ATM networks become more richly interconnected and as the use of switched ATM services become more widespread in Federal ATM networks, the establishment of connections involving multiple ATM service provider networks will become common-place. The corresponding need and complexity of address interworking of ATM service provider networks will place very heavy demands on address interworking between providers of ATM global services and on individual Federal ATM networks. Support of the flexible connectivity demands cited above will require seamless interworking of the Open Systems Interconnection (OSI) Network Service Access Point (NSAP) addressing schemes and formats used by endpoints in private networks, and the E.164 addressing schemes and interfaces used by ATM for public network addressing. Efforts are currently in progress within the ATM Forum Broadband-Integrated Services Digital Network (B-ISDN)-Inter-Carrier Interface (B-ICI) Working Group to define the protocol(s) for mapping and translation of non-E.164 NSAP addresses to native E.164 addresses across the B-ICI.

ADDRESS RESOLUTION

For IP over ATM, address resolution is the procedure by which LAN Emulation (LANE) clients associate a LAN destination media access control (MAC) address with a matching ATM address. LANE is a set of services, functional groups and protocols which provide for the emulation of LANs utilizing ATM as a backbone to allow connectivity among LAN and ATM end stations. IP in the classical mode is typically a shared media, therefore MAC is required. However, since ATM uses point-to-point links, MAC is not required. The transmission of IP over ATM requires a mechanism for associating destination MAC addresses with matching ATM addresses. ATM Address Resolution Protocol (ATMARP) servers provide the required mapping mechanism. Currently, ATM does not support multicast address services. Therefore, there are no mappings available from IP multicast addresses to ATM multicast services. However, mapping is available for classical IP employing address resolution protocol (ARP) over ATM. Each IP station connected to the ATM network must have an ATM hardware address and an ATMARP request address. The ATMARP server is responsible for resolving ATMARP requests of all IP members within logical IP subnetworks.

CALL SETUP

Call setup is the protocol that supports the establishment of a connection or call between different parties. Call setup includes both call request and call answer. In ATM networks, encapsulated traffic is transported over the connection transparent to the network. The ATM traffic descriptor, broadband bearer capability, and quality of service (QoS) parameter information elements (IEs) are mandatory in the SETUP message. If the network is able to provide the traffic parameter values specified in the ATM traffic descriptor IE, the alternative ATM traffic descriptor IE, or the minimum acceptable ATM traffic descriptor IE, as appropriate, the network will progress the connection establishment request onward. If the network is unable to provide the traffic parameter values specified in the ATM traffic descriptor IE, the alternative ATM traffic descriptor IE, or the minimum acceptable ATM traffic descriptor IE, as appropriate, the network will reject the connection establishment request.

NETWORK MANAGEMENT

Network management is the monitoring, analysis, evaluation, and control of network performance to ensure that the end user is provided an agreed level of performance, and ensure optimal use of network resources. The seamless architecture and mix of traffic service requirements in ATM networks pose a unique challenge for network management systems in just about every area of network management. Following encapsulation, ATM network management of encapsulated traffic is treated the same as all other traffic in the

same category. The ATM Forum's Network Management Working Group has generated a framework to address network management functionality. The Network Management model, often referred to as "M specs," addresses management interfaces for both private and public network services. The "M specs" consist of five management interfaces labeled M1 through M5. The interfaces collectively provide end-to-end ATM service monitoring and control information.

A standard Management Information Base (MIB) structure defines management items for network components that can be accessed by a network manager, and various fields and values that the agent software must keep track of in the managed device. The network management information maintained in the MIB is modeled in terms of managed objects. Managed objects are abstractions of data communications resources used to model information about the object and its role in the network. Included in this information are characteristics such as its name, the events it may generate, and the operations it may be requested to perform. As with other telecommunications networks, network management operations for networks supporting IP over ATM can generally be organized into the functional categories of: configuration management, fault management, performance management, accounting management, and security management.

SECTION 1.0

INTRODUCTION

1.1 PURPOSE	

This report presents the results of an examination of selected technical interface considerations with potential impact on the transfer of connectionless Internet Protocol (IP) data over connection-oriented Asynchronous Transfer Mode (ATM) networks.

1.2 SCOPE	

This report provides: (1) a general overview of salient features of a connection-oriented ATM architecture and of a connectionless IP-based architecture; and (2) selected technical interface considerations with potential impact on the transfer of IP data over ATM virtual connections.

1.3 BACKGROUND	

The direct operation of IP-a connectionless protocolover ATM-a connection-oriented transport modeis one of several application areas currently being examined by both the ATM Forum's Multiprotocol Over ATM (MPOA) Working Group and the Internet Engineering Task Force (IETF)'s IP Over ATM Working Group. Connection-oriented and connectionless service are the two fundamental types of transport modes used to convey information in today's communications networks. Services that are connection-oriented require the establishment of an end-to-end connection prior to the transmission of any information by the source. With the connectionless mode, each element of information is sent independently without regard to the transmission of previous or subsequent elements. The prior establishment of an end-to-end connection is not required.

Each of the two transport modes is designed to support specific applications and meet specific quality of service (QoS) requirements. Using the connectionless mode, initial overhead is generally less than for connection-oriented service. However, because of its capability to support high data rates, provide gains in bandwidth use efficiency, and accommodate a broad range of traffic from diverse sources, ATM is gaining broad acceptance as the preferred transport mode for both connection-oriented and connectionless services across multiple public networks.

ATM is connection-oriented in the sense that ATM networks require that an end-to-end virtual channel/circuit (VC) be established prior to the transmission of user data. A VC is a communications channel that provides for the sequential transport of ATM cells. The current ATM Forum solution for the operation of IP over ATM requires that a VC be

established for the transport of IP packets. There is a certain amount of routing overhead and delay associated with the establishment of each VC. For support of applications of long duration, the associated overhead and inherent delay might not be considered significant because of the relatively long time the connection is in use before being disestablished. However, for short-duration applications (e.g., IP) where VCs are established and broken with great frequency, the associated routing overhead and inherent delay might well be considered significant under conditions of severe traffic overload and/or network degradation. This effort examines selected technical interface considerations with potential impact on the transfer of IP data over ATM VCs.

1.4 ORGANIZATION	

This document is divided into the following subsequent sections:

Section 2.0, (*ATM/Internet Protocol Service Architectures*): Provides a general overview of the salient features of connection-oriented ATM and connectionless IP-based architectures.

Section 3.0, (*ATM/Internet Protocol Data Transfer Technical Considerations*): Examines selected technical interface considerations with potential impact on the transfer of connectionless IP data over ATM VCs.

1.5 REVISIONS	

This document will be updated as directed by the Technology and Standards Division (N6), Office of the Manager, National Communications System (OMNCS). Comments and recommendations should be forwarded to:

Office of the Manager, National Communications System

Attn: Technology and Standards Division (Mr. Dale Barr)

701 South Court House Road

Arlington, Virginia 22204-2198

SECTION 2.0

ATM/INTERNET PROTOCOL SERVICE ARCHITECTURES

The transfer of IP datagrams over ATM requires the translation from a connectionless data transfer mode to a connection-oriented ATM transfer mode. Major differences exist in the makeup, interfacing, and operation of connection-oriented and connectionless communications protocols. These differences include different addressing schemes, different routing procedures, different control measures, and other factors involving the general complexities of the two modes. This section provides a general overview of the salient features of connection-oriented ATM and connectionless IP communications architectures.

2.1 ATM ARCHITECTURE	

The Broadband Integrated Services Digital Network (B-ISDN) reference architecture is vertically layered and covers transport, switching, signaling and control, user protocols,

and applications of services. ATM is the underlying switching and multiplexing technology of the B-ISDN reference architecture. For the transfer of information over ATM networks, variable length user information is organized into fixed-length slots called cells. An ATM cell consists of a header containing information used to route the cells in the ATM network, and an information field containing payload information that is transported transparently by the network. The length of the header is five octets-a term for 8 bits that is sometimes used interchangeably with "byte." The length of the payload information field is 48 octets. ATM is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic. Figure 2-1 is a simplified depiction of the two fields of an ATM cell.

For ATM, the three primary layers of the B-ISDN reference architecture are: the Physical Layer, the ATM Layer, and the ATM Adaptation Layer (AAL). A discussion of the principal functions performed by these layers and an overview of how they interface is provided below.

Physical Layer: The Physical Layer is the lowest layer of the reference architecture. It provides for the transmission of cells over a physical medium connecting two ATM devices. This layer interfaces with the actual physical medium, extracts and inserts ATM cells within Time Division Multiplexed (TDM) frames, and passes these to and from the ATM Layer. The principal functions performed by this layer are bit transfer, line coding, electrical-optical transformation, cell rate decoupling, header error control (HEC), header sequence generation/verification, and the adaption, generation/recovery of transmission frames.

ATM Layer: The ATM Layer is the second layer of the reference architecture. It is independent of the Physical Layer. This layer performs multiplexing, switching, routing, and control actions based on information contained in the ATM cell header. It passes cells to and accepts cells from the AAL. Specific functions performed in this layer include encapsulation - taking data formatted for one protocol (e.g., IP) and enclosing it within another protocol. (Encapsulation includes cell header generation and extraction). Other functions performed in this layer are: generic flow control (GFC), translation of virtual path identifiers (VPIs)-an eight-bit field in the ATM cell header which indicates the virtual path over which the cell should be routedand virtual channel identifiers (VCIs)-a unique numerical tag in the ATM cell header that identifies the VC over which the cell is to traveland cell multiplexing/demultiplexing.
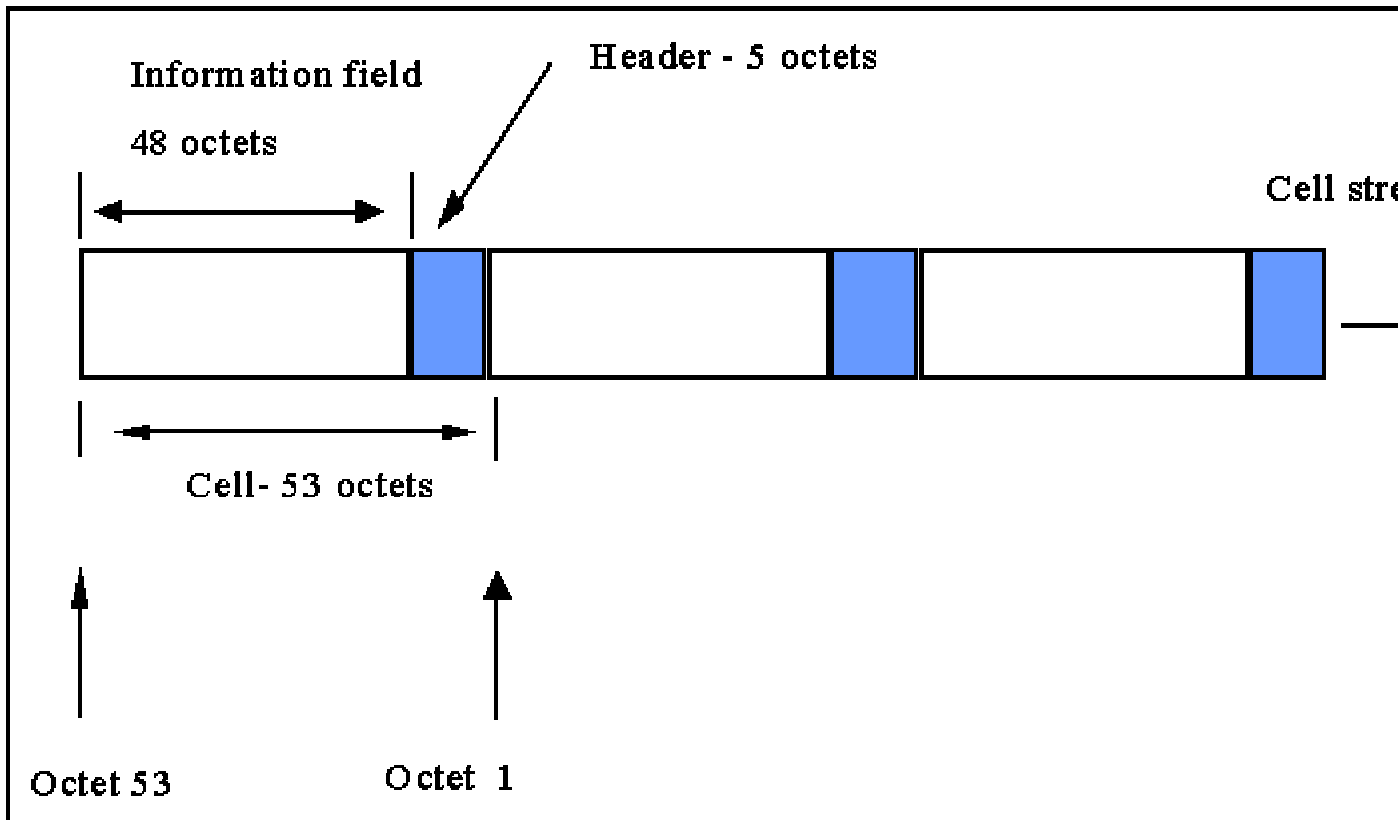
Figure 2-1. Basic ATM Cell

AAL: The AAL passes protocol data units (PDUs)-messages composed of payload and protocol-specific control informationbetween the ATM Layer and the higher layers. PDUs may be of variable length, or may be of fixed length different from the ATM cell length. When a VC is created, a specific AAL type is associated with the VC. The basic function of the AAL is to translate higher layer services into the size and format of an ATM cell, and to map higher layer PDUs into the information field of the ATM cell and vice versa. The AAL at the originating endpoint segments PDUs into ATM cells and marks the last cell of each PDU. At the destination endpoint, the AAL at that location uses the end of packet marker to reassemble the data from the

cells received. Independent functions required to support the above services have been organized in two logical sublayers: the convergence sublayer (CS) and the segmentation and reassembly sublayer (SAR). The CS performs required conversions between ATM and non-ATM formats. The CS is further divided into a common part (CPCS) and a service specific part (SSCS). The CPCS provides services such as padding and CRC checking. The SSCS is service dependent. Figure 2-2 shows the positional relationship between the AAL, CS, SAR and ATM layers/sublayers.

```
  A       C
  A       S                         CPCS
  L
                                    SSCS


                                SAR Sublayer


                                 ATM Layer
```
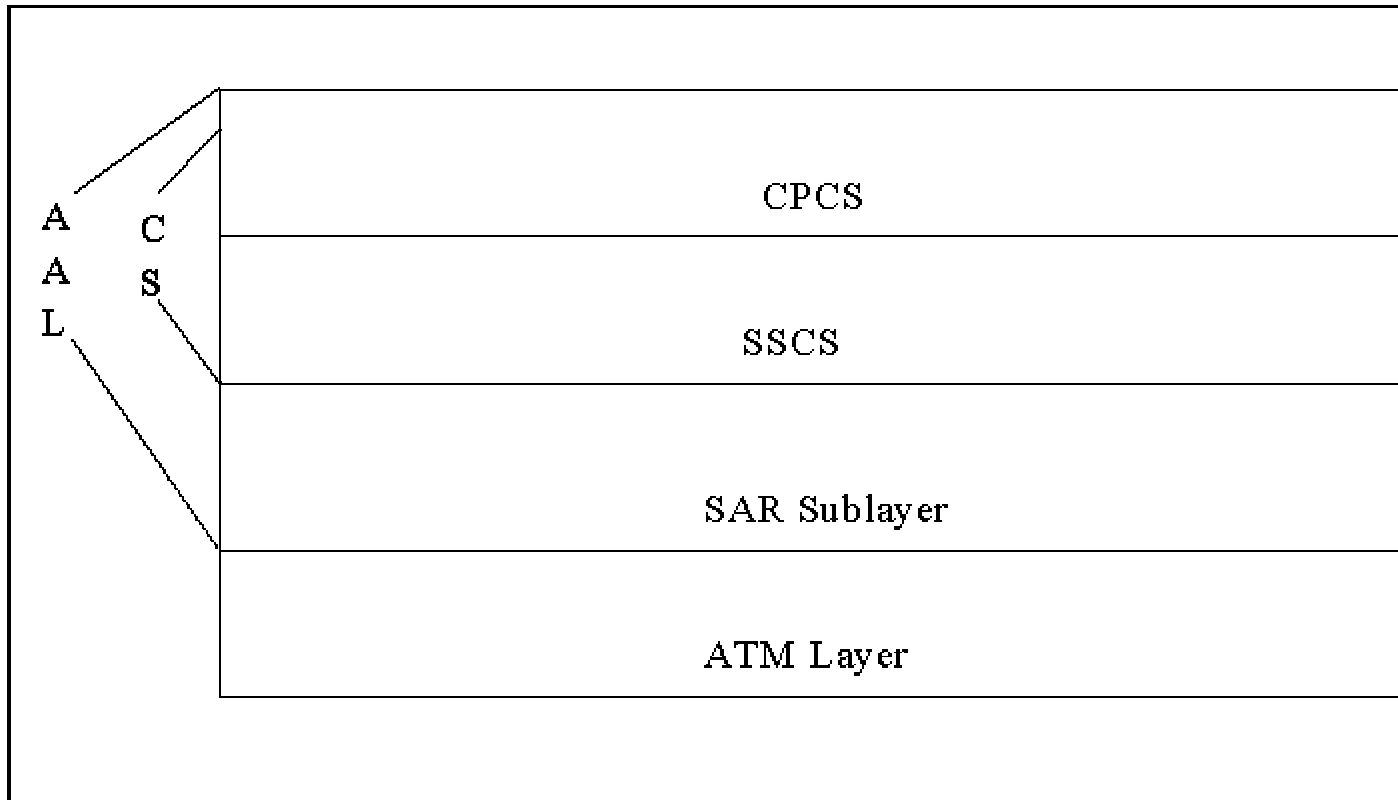
Figure 2-2. AAL, CS, SAR Relationship, and ATM Layer/Sublayer Position Relationship

The primary function of the SAR is to: (1) segment higher-layer PDUs into 48 byte packets and forward them to the ATM layer to generate 53 byte ATM cells, and (2) reassemble the contents of the ATM cell information fields received from the ATM layer into higher-layer PDUs.

There are four different AAL types or categories. The simplest of these is AAL-1. AAL-1 functions in support of the transport of constant bit rate (CBR), time-dependent traffic such as voice and video. AAL-2 is still in the stage of being defined by international standards bodies. It is intended that AAL-2 will be used to support transport of variable bit rate (VBR) video

transmission. AAL-3/4 was originally two separate AALs intended to support transport of VBR, delay-tolerant, connection-oriented and connectionless data traffic requiring some sequencing

and/or error detection support. However, as the specifications evolved, it became apparent that both services required similar procedures and as a result the specifications were merged to become the AAL-3/4 standard. AAL-5 was designed specifically to support transport of VBR, delay-tolerant connection-oriented data traffic requiring minimal sequencing or error detection support. This is typically the type of data traffic

found in current local area networks (LANs). AAL-5 evolved because AAL-3/4 was considered too complex and inefficient for LAN traffic. Because of smaller bandwidth overhead, simpler processing requirements, and reduced implementation complexity, the IETF's Request for Comment (RFC) 1483, *Multiprotocol Encapsulation over ATM Adaption Layer 5,* specifies the use of AAL-5 for the transmission of packets across an ATM VC. Consequently, where necessary, for the remainder of this analysis any discussion of AALs will be limited to AAL-5. The other three types of AALs are mentioned for introductory purposes only.

2.1.1 The ATM Cell Header	

The ATM cell header is the protocol control information located at the beginning of an ATM cell. There are two standardized ATM cell structure coding schemes used in the cell header. The user-to-network interface (UNI) coding scheme is used for the interface between the user or customer premises equipment (CPE) and the network switch.

The network-to-network interface (NNI) coding scheme is used for the interface between switches or between networks. The format for both cell structures is identical except for an additional field in the UNI coding scheme header for
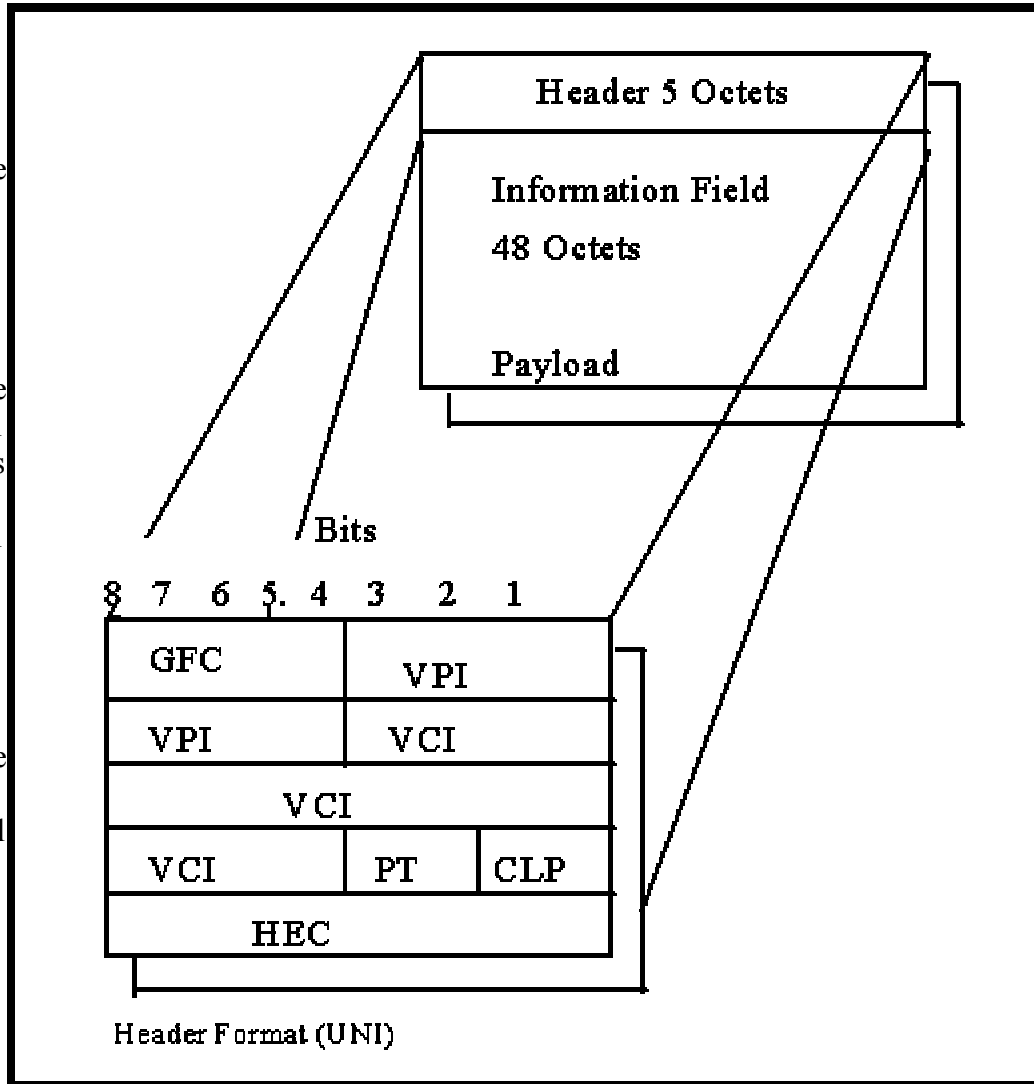
Figure 2-3. ATM Cell Header

providing: (1) local contention resolution and simple flow control for shared medium-access arrangements at the CPE; and (2) an increase in the number of bits in the VPI field for the NNI coding scheme. Figure 2-3 shows the six fields of the ATM cell header for the UNI coding scheme. A description of each field is provided below.

Generic Flow Control (GFC): This field provides local contention resolution and simple flow control for shared medium-access arrangements at the CPE. The value encoded in this field is not carried end-to-end.

Virtual Channel Identifier (VCI): This field is used to provide a unique numerical tag that identifies a virtual channel over which the cell is to travel. It is used to establish

connections using translation tables at switching nodes that map an incoming VCI to an outgoing VCI.

Virtual Path Identifier (VPI): This field is used to indicate the VP the cell is to be routed over.

Payload Type (PT): This field is used to differentiate cells that traverse the same virtual circuits.

Cell Loss Priority (CLP): This field is used to indicate two levels of priority for ATM cells. Depending on network conditions, cells with CLP set to CLP=1 may be discarded to preserve the cell loss ratio of cells with CLP set to CLP=0.

Header Error Control (HEC): This field is used to perform a cyclic redundancy check (CRC) calculation on the header to detect and correct errors.

2.1.2 ATM Cell Routing	

The translation of VPI and VCI values contained in the header of ATM cells provide the basis for switching in ATM networks. The VCI value is used to establish connections using translation tables at switching nodes that map an incoming VCI to an outgoing VCI. The VPI value is used to establish a virtual path connection (VPC) for one or more logically equivalent VCIs in terms of route and service characteristics. Based on assigned VCI/VPI values, virtual channel connections (VCCs)-a concatenation of virtual channel links (VCLs) that extends between the points where the ATM service users access the ATM layer-are established. By reserving capacity on a VPC in anticipation of later call arrivals, new VCCs can be established by executing simple control functions at the endpoints of the VPC. In such cases, no call processing is required at intermediate switching nodes. Additionally, because the functions needed to set up a path through the network are only executed once for all VCCs subsequently using that path, the load on the VCC control mechanism in the core network is also reduced.

At the VC level, ATM cells are routed and switched individually at ATM VC switches or the points of cross-connection for VCLs. This switching involves the translation of VCI

values of the incoming VCL into the appropriate VCI values of the outgoing VCL at the connection point. Thus the value of VCI can potentially change for each link as the VC spans multiple links. A virtual link is originated or terminated by the assignment or removal of the VCI value. When a virtual channel is created, the ATM switch creates and maintains a table entry that maps inbound VCIs on an inbound port to an outbound port

Routing functions for VPs are performed at ATM VP switches or points of cross-connection. VP routing involves translation of the VPI values of the incoming VP links into the VPI values of the outgoing VP links. A virtual path link (VPL) is originated or terminated by the assignment or removal of the VPI value. A concatenation of VPLs forms a VPC. Cell sequence integrity is

preserved by the ATM layer for cells belonging to the same VPC. Hence, cell sequence integrity

is preserved for each VCL within a VPC.

The ATM VP concept simplifies path network design and increases flexibility of path management for routing and capacity allocation. Because ATM multiplexing is non-hierarchical, path capacity does not need to be explicitly assigned at the VPC point at VP establishment time, but is handled by separate management procedures such as call admission control and usage monitoring-functions carried out at the ingress VPC end-point. Consequently, VPC points along the VP route are not affected by changes in VP capacity allocations which may be initiated by capacity management procedures. The independence of route and capacity management leads to two important features of ATM networks: adaptive network reconfiguration and dynamic bandwidth allocation, both of which are important in executing restoration strategies. Figure 2-4 shows a conceptual VC, VP, and transmission path relationship.
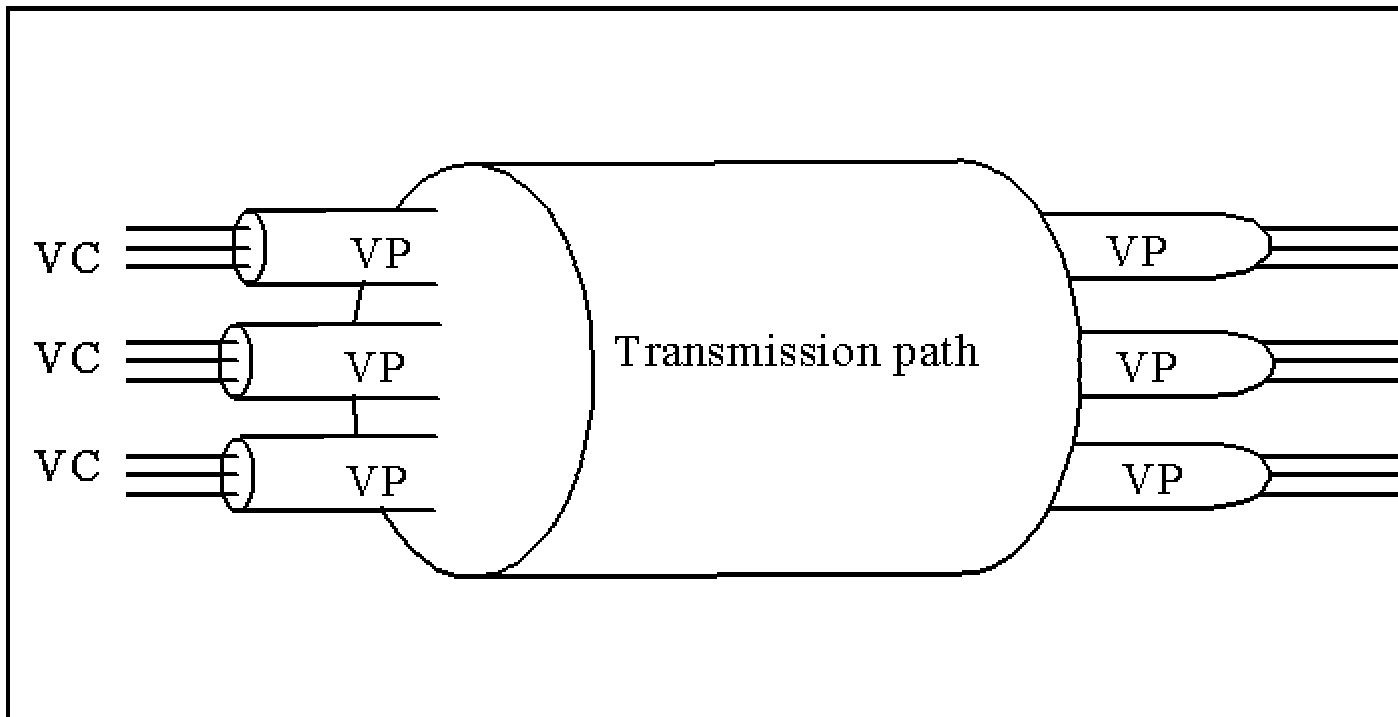
Figure 2-4. VC, VP, and Transmission Path Relationship

2.2 INTERNET PROTOCOL ARCHITECTURE	

Connectionless communications architectures describe a "packet-based" system whereby the

message or data stream is broken into a sequence of smaller self-sufficient units called packets or datagrams. The resulting datagrams containing full source and destination address information and a data segment are then individually routed through the supporting telecommunications

network without prior establishment of an end-to-end telecommunications connection. The datagrams are forwarded as variable-length PDUs consisting of a header followed by the data

being transported. The header contains the required source, destination and control information. The content of the data portion depends on the transport layer protocol in use and the application. At each node, the datagrams are received, stored, processed, and retransmitted as required until

they reach the matching destination address. In internetworks, network routers dynamically determine the best path for datagrams to follow to reach their destination based on the status and congestion of network links.

2.2.1 Internet Protocol	

IP operates at the network layer of the Open Systems Interconnection (OSI) model to route datagrams across an interconnected set of networks. It was originally developed by the Department of Defense to support interworking of dissimilar computers across networks. The network layer provides higher-level OSI layers independence from the routing and switching associated with establishing network connections. Network layer functions include addressing, end-point identification, and service selection. Although IP provides connectionless service, it does not contain the necessary mechanisms to ensure end-to-end reliability, flow control, sequencing, or other services commonly found in host-to-host-protocols. It must rely on the services of supporting networks or other protocols to guarantee specific qualities of service.

In a generic IP network, data from the sending host is joined to an IP header to create an IP datagram which is subsequently passed to the subnetwork. Prior to joining with the IP header, data from the sending host is fragmented, as required, to meet datagram size limitations of networks with which IP carries out an interface function. Each datagram is independent and has no relationship with other datagrams. At the subnetwork it is encapsulated in the appropriate data link layer frame and routed to successive subnetwork nodes in accordance with the destination address information contained in the header. At gateways, subnetwork information is removed from the datagram and the datagram is encapsulated in a datalink frame appropriate to the network it is about to enter. Eventually, the datagram arrives at its destination where the

fragments are reassembled and delivered to the intended host. Figure 2-5, is a simplified depiction of an IP datagram. A description of each IP datagram field is provided below:

Version: This field identifies the version of IP in use at a particular gateway or router.

Header Length: This field specifies the length of the IP Header in multiples of 4 octets.

Service: This field identifies QoS parameters which characterize service options provided by IP. Examples are reliability, throughput and delay.

Total Length: This field identifies the length of the datagram. The minimum length of a datagram is 576 octets. The maximum length of a datagram is 65,535 octets.

Identification: This field contains the sequence number of the datagram.

Flag: This field contains a code which permits or prohibits fragmentation of the datagram.

Offset: This field identifies a fragment as part of a complete datagram and indicates where in a PDU it belongs.

IP Datagram	IP Header		
		Version	Header Length
		Service	
		Total Length	
		Identification	
		Flag	Offset
		Life	
		Protocol	
		Header Checksum	
		Source Address	
		Destination Address	
		Options	
		Padding	
		Data	
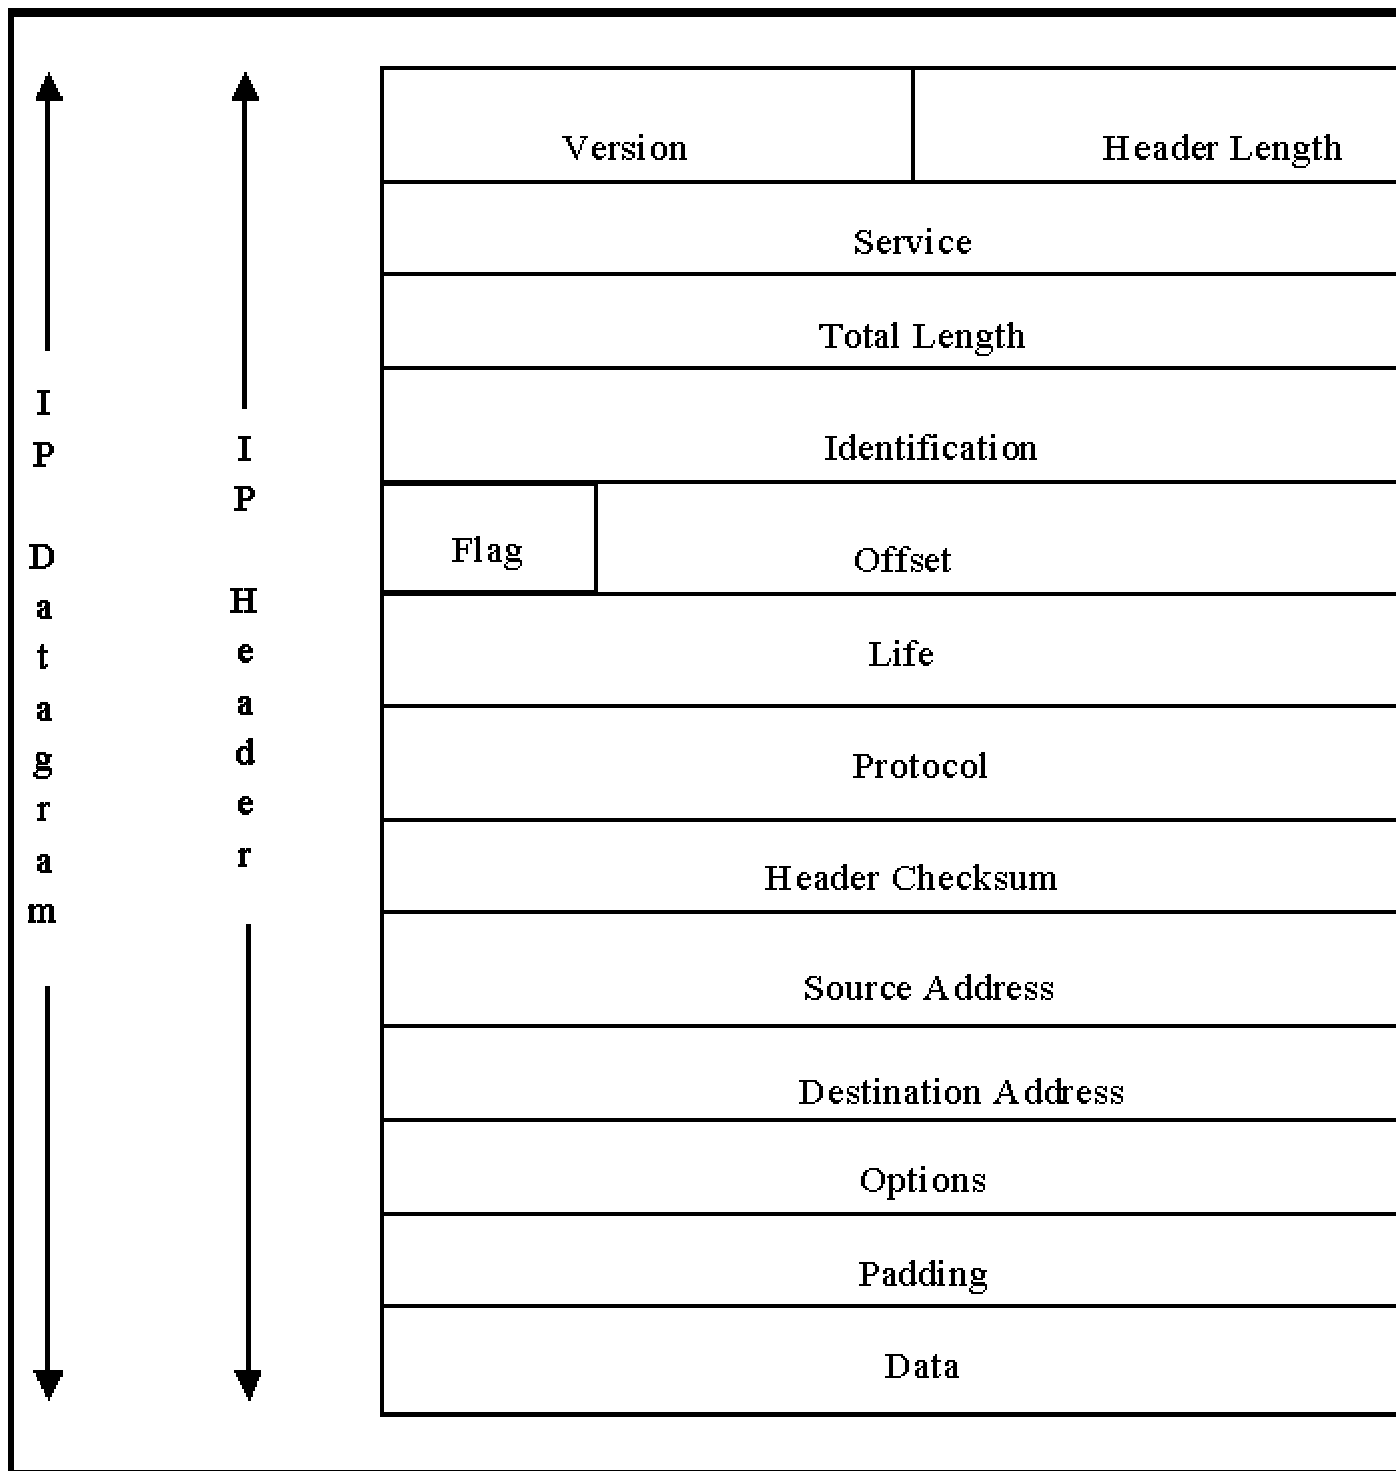
Figure 2-5. IP Datagram

Life: This field indicates how long the datagram has to live. It is set by the sender of the datagram and reduced at points along the route where it is processed. This field prevents

datagrams or datagram fragments from looping indefinitely.

Protocol: This field identifies the next-higher-layer level protocol the datagram expects at the destination host.

Header Checksum: This field is used to perform error checking on the header. If the checksum fails, the datagram is discarded by the entity detecting the error.

Source Address: This field contains the address of the originating host.

Destination Address: This field contains the address of the destination host.

Options: This is a variable length field used to request specific routing, handling, and other services.

Padding: This is a variable length field used to ensure that the header is a multiple of 32 bits.

Data: This is a variable length field containing the payload information.

2.2.2 Transmission Control Protocol (TCP)	

The User Data Protocol (UDP) and Transmission Control Protocol (TCP) are the two transport layer protocols used with IP. Transport layer functions include flow control, error detection, recovery, and other functions required to maintain the state of the connection; compensate for lost, missing, or destroyed datagrams; and optimize the use of network resources. The transport layer is composed of multiple protocols. Each protocol has a different set of characteristics, features and intended use. UDP is a connectionless protocol. It operates as a simple extension of IP. It lacks end-to-end capabilities such as sequencing and flow control. With UDP, measures to ensure reliable

transport of data are left to the application.

TCP-a connection-oriented protocolworks with IP to provide service guarantees between hosts. TCP builds its transport layer services on top of the network layer routing services provided by IP. It can support a wide range of upper-layer protocols (ULPs). In concert with underlying layers, TCP performs end-to-end functions to guarantee delivery of the exact data stream passed to it by the ULPs. TCP/IP traffic sources include at least three standardized ULPs: (1) Telenet-a terminal emulation protocol commonly used on the command line level of the Internet, (2) File Transfer Protocol (FTP)-a program that transmits files from one computer to another, and (3) Simple Mail Transfer Protocol (SMTP)-a standard application which allows electronic mail messages to be packaged and delivered to remote systems. ULPs channel data through TCP in streams for delivery to peer ULPs at destination hosts. TCP divides the data streams into encapsulated TCP segments called TCP data units (TCPDUs), containing appropriate addressing and control information. These TCPDUs are then passed to IP at the network layer for transmission within IP datagrams through the communications network to the TCP at the destination host. In IP datagrams, the TCP header follows the IP header and

supplies information specific to the TCP protocol. TCP guarantees delivery of the exact data stream by the use of sequence numbers and timers that require acknowledgments within designated time periods. Flow control is performed by the destination host sending a window value to the source host specifying the number of octets the source is authorized to send. Figure 2-6 provides a simplified view of a TCPDU. A description of each field is provided below.

```
                            ┌──────────────────────────────────────────┐
  ↑          ↑              │              Source Port                 │
  │          │              ├──────────────────────────────────────────┤
  T          │              │            Destination Port              │
  C          │              ├──────────────────────────────────────────┤
  P          │              │            Sequence Number               │
             T              ├──────────────────────────────────────────┤
             C              │         Acknowledgment Number            │
  D          P              ├─────────────────────┬────────────────────┤
  a                         │      Control        │      Reserved      │
  t          H              ├─────────────────────┴────────────────────┤
  a          e              │                Window                    │
             a              ├──────────────────────────────────────────┤
             d              │               Checksum                   │
  U          e              ├──────────────────────────────────────────┤
  n          r              │            Urgent Pointer                │
  i                         ├──────────────────────────────────────────┤
  t                         │                Options                   │
             │              ├──────────────────────────────────────────┤
  │          │              │                Padding                   │
  │          │              ├──────────────────────────────────────────┤
  ↓          ↓              │                 Data                     │
                            └──────────────────────────────────────────┘
```
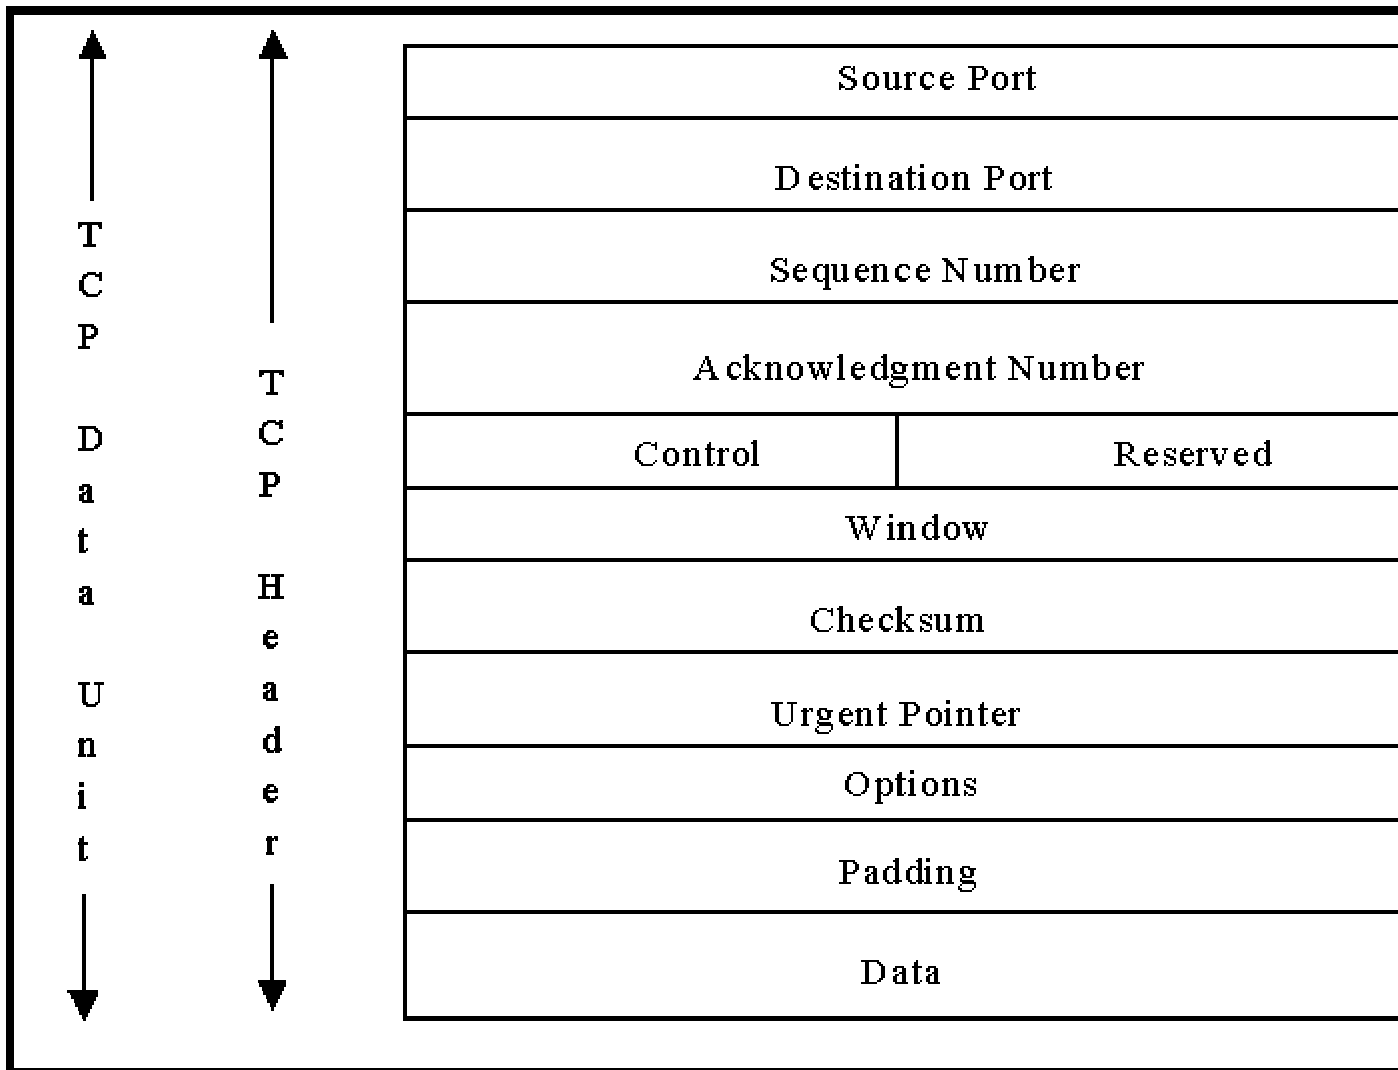
Figure 2-6. TCP Data Unit

Source Port: This field identifies the source ULP or service access point (SAP)-the boundary between the AAL and ATM layer-within the source node or host.

Destination Port: This field identifies the ULP or SAP within the destination node or host.

Sequence Number: This field contains the value of the sequence number of the first data octet in the information PDU.

Acknowledgment Number: This field contains the value of the next sequence number that the sender of the segment is expecting to receive.

Control: This field contains the control bits. Control bits are set to indicate urgency, request immediate delivery, request reset, synchronize sequence numbers, and end of data stream.

Reserved: This field is reserved for future use.

Window: This field is used for flow control. It specifies the number of data octets which the sender of this segment is willing to accept.

Checksum: This field contains a value used to ensure data is transmitted without error. It is created by adding the binary value of each alphanumeric character in a block of data and sending it with the data.

Urgent Pointer: This field indicates the current value of the urgent pointer as a positive offset from the sequence number of this segment. It is the sequence number of the octet following urgent data.

Options: This is a variable length field used to request specific receive buffer size, and other services. Options include "end of option list," "no-operation," and "maximum segment size."

Padding: This is a variable length field used to ensure that the TCP header ends and data begins on a 32-bit boundary.

Data: This is a variable length field used for payload data.

Figure 2-7 depicts the protocol model for connectionless IP over ATM. The protocol stack provides a mechanism for the transfer of application data to corresponding destination applications.

Figure 2-7. IP Over ATM Protocol Stack

SECTION 3.0

ATM/INTERNET PROTOCOL DATA TRANSFER TECHNICAL CONSIDERATIONS

The IETF's IP over ATM Working Group charter identifies the following as initial areas in which it plans to develop experimental protocols for the operation of internetwork protocols over ATM networks: (1) encapsulation, (2) multicasting, (3) addressing, (4) address resolution, (5) call setup, and (6) network management. Using these areas as focal points, the remainder of this section discusses significant technical considerations involved in interfacing and operating IP over ATM.

3.1 ENCAPSULATION

Encapsulation is enclosing data formatted for one protocol within another protocol in order to transmit the data across a network for which the original protocol was not designed. It is one method of supporting multiple protocols across a single connection. For non-similar protocols, encapsulation provides for the creation of an envelope into which other protocols are placed. The data within the envelope is transported transparent to the network. At the destination endpoint, the data is removed from the envelope and returned to its original form.

Encapsulation of IP datagrams for transmission over ATM is accomplished by: (1) the segmentation by ATM Adaptation Layer-5 (AAL-5) of PDUs received from ULPs into 48-byte packets; (2) mapping the 48-byte packets into the information fields of the ATM cell; and (3) forwarding the information fields sequentially to the ATM layer to generate 53-byte ATM cells. In the transmit direction, the encapsulation function includes the generation of an appropriate cell header for the information field by the ATM layer, less the HEC which is the responsibility of the physical layer. (The encapsulation function can also include the translation from a SAP identifier to a VPI and VCI). In the receive direction, the de-encapsulation function performed by the ATM layer includes the extraction of the ATM cell header and passing the cell information field to the AAL for mapping of the information field contents into appropriate PDUs for forwarding to ULPs (e.g., Telnet, FTP, and SMTP).

Fragmentation is an IP convention designed to compensate for the maximum frame size limitation of a network through which a datagram passes. The technique involves splitting a single datagram into two or more smaller fragment blocks, each of which is routed independently to the original destination. Upon arrival at the target host, the fragments are reassembled by IP before delivery to the local application. Encapsulation of IP datagrams for transmission over ATM requires fragmentation of the datagrams into packets with appropriate padding to ensure that integral multiples of 48 octets are presented to the ATM layer to prevent the need to send partially filled cells after segmentation.

The size of an IP datagram in most instances will be many times the size of an ATM cell. Consequently, following segmentation, it is expected that a single IP datagram will correspond to a number of ATM cells. The values contained in the VCI and VPI fields of

the ATM cell header are used by switches in ATM networks to perform the routing and switching functions. Incoming VCI and VPI values are translated by the switch to outgoing values. This translation function introduces a certain amount of delay and overhead in the switch for each cell. In ATM networks, the recurrence of cells containing information from an individual user is not necessarily periodic. Consequently, it appears that considerable savings in routing overhead and delay might be realized if all cells belonging to a particular datagram could be identified as such for routing over the same VC connection. Both the ATM Forum and the IETF are currently considering approaches to reducing both the routing overhead and the delay associated with the transmission of IP over ATM.

3.2 MULTICASTING	

Multicasting is the process whereby a source interface to a group sends a single PDU to more than one selective destination simultaneously, using a single local transmit operation . Broadcastingthe process whereby a source interface sends a single PDU to all destinations simultaneously, may be considered a special case of multicasting. The IP multicast model is receiver initiated. Receivers wishing to subscribe to an IP multicast group inform their local routers. Routers, using multicast routing protocols, maintain and disseminate membership information.

In ATM networks, multicast connections are supported by a collection of sender (root)-initiated point-to-multipoint, unidirectional VCs and associated endpoints. This arrangement requires that the sender know each intended recipient, and explicitly establish a VC between itself as the root and each recipient as a leaf node. A typical connection setup is achieved by first establishing a point-to-point connection between the root node and one leaf node. After the first setup is complete, additional leaf nodes are added to the connection by "Add-Party" message requests via the root nodes. Leaf nodes may be added or dropped at any time after establishing the connection.

IP multicasting over ATM networks requires a mechanism for mapping IP group addresses to corresponding ATM addresses. In the ATM Forum's MPOA model, the mechanism used is the Multicast Address Resolution Server (MARS). The MARS acts as a registry, associating layer 3 multicast group identifiers with the ATM interfaces

representing the group's members. Each ATM based host and router client communicates with the MARS by using a globally known VC. The MARS may reside within any ATM endpoint that is directly addressable by the endpoints it is serving. When a new host is added to the network, it must register with the MARS to provide a table entry that maps its corresponding IP address to its ATM address.

The MARS manages a cluster of ATM-attached endpoints. The IETF defines a cluster as "the set of ATM interfaces choosing to participate in direct ATM connections to achieve multicasting of AAL Service Data Units (SDUs)-units of interface information whose identity is preserved from one end of a layer connection to the otherbetween themselves." Endpoints wishing to join a multicast cluster must be configured with the ATM address of the node on which the cluster's MARS resides. Traffic between interfaces to different clusters passes through an inter-cluster devicee.g., an IP multicast router with logical interfaces into each cluster. The distribution of multicast group membership information between the MARS and the endpoints is accomplished via MARS messages. Endpoint address resolution entities query the MARS when a network level address needs to be resolved, and informs the MARS when they need to join or leave a particular group. It should be noted that an endpoint decision to join or leave a group is a local issue. It has no effect on other members of the multicast group.

A potential concern with the MPOA model multicast arrangement is the requirement to establish dedicated VCs between the root and each recipient leaf node. Considerable discussion is currently ongoing within the IETF concerning ways by which the overhead associated with the establishment and maintenance of separate multicast VCs might be avoided. It is generally agreed that for certain long-duration applications requiring QoS guarantees, the establishment of multicast VCs can be justified. However, for relatively short-duration applications lacking the requirement for QoS guarantees, the use of some type of shared service is also being discussed as an option.

The current focus of the discussion is on the relative merits of using "VC meshes" versus multicast servers (MCSs) to support network layer (e.g., IP) multicasting over ATM. Both seem to offer advantages and disadvantages. With the multicast VC mesh, each source establishes independent point-to-multipoint VCs to the set of leaf nodes it wishes to send messages. Interfaces for the leaf nodes originate and terminate VCs, as appropriate, for each active leaf node-the term "VC mesh" is used to describe the resulting crisscross VC pattern.

With the MCS, each source establishes a VC to the MCS. The MCS establishes a point-to-multipoint VC to the leaf nodes. AAL-SDUs arriving on incoming VCs are reassembled by the MCS and queued for transmission on its single outgoing point-to-multipoint VC. Reassembly of incoming AAL-SDUs by the MCS is required since AAL-5 does not support cell level multiplexing of different AAL-SDUs on a single outgoing

VC. IETF Network Working Group RFC 2022, entitled Support for Multicast over UNI 3.0/3.1 based ATM Networks, lists the following tradeoffs for each approach:

The VC Mesh: The VC mesh lacks the obvious single congestion point of an MCS. Throughput is likely to be higher, and end-to-end latency lower, because the mesh lacks the intermediate AAL-SDU reassembly that must occur in MCSs. The underlying ATM signalling system also has greater opportunity to ensure optimal branching points at ATM switches along the multicast trees originating on each source. However, resource consumption will be higher. Every group member's ATM interface must terminate a VC per sender (consuming on-board memory for state information and requiring and buffering in accordance with the vendor's particular architecture). With an MCS, only two VCs (one out, one in) are required. With a multicast server, the allocation of VC related resources is also lower within the ATM internetwork.

The Multicast Server: With regard to the signalling load, the MCS has the advantage over the VC mesh when faced with dynamic sets of receivers. Every time the membership of a multicast group changes (a leaf node needs to be added or dropped) only a single point-to-multipoint VC needs to be modified when using an MCS. This generates a single signalling event across the MCS's UNI. When a membership change occurs in a VC mesh, signalling events occur at the UNI of every traffic source. The transient signalling load is determined by the number of sources. However, MCS introduces a "reflected packet" problem which requires additional AAL-SDU information to be carried in order for network layer sources to detect returns of their own AAL-SDUs.

3.3 ADDRESSING	

Addressing is the assignment of a sequence of characters in a message, datagram, or cell header to designate its origin and destination. Each device within a network must have a unique and identifiable address in order to receive and transmit messages. Addressing in an integrated IP/ATM environment is currently being considered by the IETF Networking Working Group and the ATM Forum Technical Committee as a major issue.

In a classical IP environment, the source and destination address fields of the IP datagram header contain the addressing information needed to route datagrams in a connectionless network. IP routers using static or dynamic lookup or routing tables attempt to match the network address contained in the header of a datagram with a network address entity contained in the routing table. If the destination node is on its local network, the datagram is forwarded directly to the destination host. If the destination node is on some other network, the datagram is forwarded to the IP local router for examination of the IP address and forwarding, as appropriate. With IP over ATM, the IP addressing information is encapsulated within the ATM cell along with the rest of the datagram and becomes transparent to the ATM network. Consequently, IP over ATM requires a mechanism for mapping IP addresses to corresponding ATM addresses. For classical IP over ATM, an ATM address resolution protocol (ATMARP) server - to be discussed in Section 3.4 - similar to the previously discussed MARS, provides the mapping mechanism.

An ATM address may be either a native E.164 numbera public network addressing standard used by ATM for public network addressing - or a 20-octet ATM endpoint address based on the generic Open Systems Interconnection (OSI) Network Service Access Point (NSAP) network address encoding format. The ATM Forum's *ATM User-Network Interface (UNI) Signalling Specification, Version 4.0* specifies the use of the OSI NSAP format for endpoints in private networks. As specified in the *UNI Signalling Specification,* the ATM address structure for private networks consists of multiple fields. Two of these fieldsthe End System Identifier (ESI) and the Selector (SEL) fieldsare supplied by the user side of the UNI. All other fields form a network prefix for ATM addresses. Their values are supplied by the network side of the UNI. An ATM address for an endpoint on the user side of a Private UNI is obtained by appending values for the ESI and SEL fields to the network prefix for that UNI. The network side is allowed to supply multiple network prefixes for use at a single UNI. ATM addresses in public network may use either the private ATM address structure or the native E.164 addresses. For native E.164 addresses, the network side supplies the whole ATM address.

In ATM networks, both individual and group addresses are used to identify endpoints. An individual address is used to identify a single ATM end system, whereas an ATM group address is used to identify one or more ATM end systems. An ATM end system may join or leave a group at any time by using the client registration and deregistration procedures outlined in the ATM Forum's *Interim Link Management Interface Specification, Version 4.0*. End nodes maintain the lookup tables that translate addresses into circuit paths. These circuit path lookup tables differ at every node and are maintained in a quasi-real-time fashion by a routing protocol. Each node on the network must register its address in the address register in the router. The router maintains a regularly updated register of addresses and tables for all nodes in the network. Address registration is the dynamic exchange of network routing prefixes on the network side and end system identifiers on the host side. VP routing involves the translation of VPI values of the incoming VP links into the VPI values of the outgoing VP links. A VP is assigned a specific value of VPI

each time a VP is switched in the network.

As ATM networks become more richly interconnected and as the use of switched ATM services become more widespread in Federal ATM networks, the establishment of connections involving multiple ATM service provider networks will become common-place. The corresponding need and complexity of address interworking of ATM service provider networks will place very heavy demands on address interworking between providers of ATM global services and on individual Federal ATM networks. Support of the flexible connectivity demands cited above will require seamless interworking of the Open Systems Interconnection (OSI) Network Service Access Point (NSAP) addressing schemes and formats used by endpoints in private networks, and the E.164 addressing schemes and interfaces used by ATM for public network addressing. Efforts are currently in progress within the ATM Forum Broadband-Integrated Services Digital Network (B-ISDN)-Inter-Carrier Interface (B-ICI) Working Group to define the protocol(s) for mapping and translation of non-E.164 NSAP addresses to native E.164 addresses across the B-ICI. The B-ICI is an ATM Forum defined interface between public ATM networks to support user services across multiple public carriers.

3.4 ADDRESS RESOLUTION	

For IP over ATM, address resolution is the procedure by which LAN Emulation (LANE) clients associate a LAN destination media access control (MAC) address with a matching ATM address. LANE is a set of services, functional groups and protocols which provide for the emulation of LANs utilizing ATM as a backbone to allow connectivity among LAN and ATM end stations. A MAC is a data-link layer (Layer 2) protocol that governs access to transmission media. The destination client could be a workstation or a bridge which is providing proxy service, functional group and protocol client functions for the legacy MAC devices behind it. IP in the classical mode is typically a shared media, therefore MAC is required. However, since ATM uses point-to-point links, MAC is not required. The transmission of IP over ATM requires a mechanism for associating destination MAC addresses with matching ATM addresses. ATM Address Resolution Protocol (ATMARP) servers - much like the previously discussed MARS - provide the required mapping mechanism.

Currently, ATM does not support multicast address services. Therefore, there are no mappings available from IP multicast addresses to ATM multicast services. However, as described by M. Laubach, Hewlett-Packard Laboratories, in a January 1994 memorandum entitled *Classical IP and ARP over ATM*, mapping is available for classical IP employing ARP over ATM. For such configurations, each administrative entity configures its hosts and routers as closed logical IP subnetworks (LISs).

"Each LIS operates and communicates independently of other LISs on the same ATM network. Hosts connected to ATM communicate directly to other hosts within the same LIS. Communication to hosts outside of the local LIS is provided via an IP router. This router is an ATM endpoint attached to the ATM network that is configured as a member of one or more LISs."

Each IP station connected to the ATM network must have an ATM hardware address and an ATMARP request address. An ATMARP request address is the address of an individual ATMARP server located within the LIS. The server has responsibility for resolving ATMARP requests of all IP members within the LIS.

3.4.1 Address Resolution Mechanism	

In the non-broadcast, non-multicast environment that ATM networks currently provide, ATMARP servers are the focal points for resolving address resolution requests of IP members within the LIS. The ATMARP server for a logical IP subnet may also be an IP station/ATM endpoint. Each ATMARP client is administratively configured with the ATMARP server's ATM address. Within the LIS, ATMARP clients are responsible for opening VCs to the ATMARP server to register their ATMARP information and to gain and refresh their own ATMARP entry/information about other IP members. The information provided by the clients is used by the ATMARP server to generate replies to the ATMARP requests it receives.

At call setup, the ATMARP server transmits an initialization request message to the originating ATM station. After receiving the reply message from the originating station, the ATMARP server examines the IP MAC address and the ATM address contained in the reply message, and update the relevant mapping data contained in its ATMARP table. Upon receiving an ATMARP request message from a client for a mapping of an IP address to a corresponding physical address, the ATMARP server searches its table. If

the destination MAC address is registered with the ATMARP server, the ATMARP server generates a return message to the requesting ATMARP client providing the requested destination ATM address. If the destination MAC address is unknown to the ATMARP server or cannot be found, then the ATMARP server broadcasts the request containing the IP target destination address in question to other ATMARP clients. This is typically the case if the destination MAC address belongs to a workstation attached to a legacy LAN on the other side of a bridge. If an ATMARP client

recognizes the address, it replies to the request and update its own ATMARP entry/information cache with the information.

Address resolution for unregistered MAC addresses is slightly more complex. Only ATMARP clients that are directly attached to the ATM network are allowed to register their own MAC addresses. Devices such as ATM LAN bridges are allowed to register only their own MAC addresses, or if they are token-ring source-route bridges, then their route descriptors. For example, an Ethernet-to-ATM transparent bridge will not register any of the MAC addresses of real Ethernet-attached workstations that it has knowledge of with the ATMARP server. Instead, it responds to request messages that have been forwarded to it from the ATMARP server with its

own ATM address, the MAC address of the actual Ethernet-attached workstation, and a flag indicating that the MAC address is "remote" from the ATMARP client that responded.

In summary, since IP addresses are assigned independently of ATM addresses, each host is required to know both its IP and ATM address, and to respond appropriately to address resolution requests. Additionally, the ATMARP client is responsible for contacting the ATMARP server to register its own ATMARP information and to gain and refresh its ATMARP entry/information about other IP members. Lastly, it is the ATMARP client's responsibility to: (1) initiate the VC connection to the ATMARP server, (2) respond appropriately to ARP request packets received on any VC, (3) generate and transmit request packets to the ATMARP server, and process replies from the ATMARP server, as appropriate, and (4) provide an ATMARP table aging function to remove old entry data from ATMARP tables and refresh its ATMARP server information on a timely basis to ensure that it is current.

3.5 CALL SETUP	

Call setup is the protocol that supports the establishment of a connection or call between different parties. Call setup includes both call request and call answer. In ATM networks, encapsulated traffic is transported over the connection transparent to the network. There are two methods of performing call setup functions in ATM networks - one based on using permanent virtual circuits (PVCs) and the other based on using switched virtual circuits (SVCs). PVCs are administratively established communications channels with static routes defined in advance for each connection. Because of their pre-defined nature, PVCs have negligible provisioning time. Once established, the circuit between the designated points exists until disestablished. No call establishment procedures are required. However, bandwidth and QoS requirements are application-dependent, and some traffic does not require static connections (e.g., burst data, which typically only requires the error-free transmission of large bursts of data for short durations, and is generally tolerant of network delay). For short-duration applications (e.g., IP) where VCs are established and broken with great frequency, the use of SVCs provides a more efficient use of network resources.

SVCs are connections established, maintained, and released by control signalling. The user defines the endpoints when the call is initiated. Signaling procedures for such connections are defined in terms of messages and the information elements (IEs) used by network elements (e.g., end-point equipment, ATM switches) to characterize and establish the connection. Establishing a SVC connection to a host that is separated by multiple switch hops requires a hop-by-hop series of messages to be exchanged. The calling party commences the action by forwarding a SETUP message to the first switch containing all the information required to process the call. In particular, the called party address information is contained in the called party number IE, possibly supplemented by other IEs. The ATM traffic descriptor, broadband bearer capability, and QoS parameter IEs are mandatory in the SETUP message.

If the network is able to provide the traffic parameter values specified in the ATM traffic descriptor IE, the alternative ATM traffic descriptor IE, or the minimum acceptable ATM traffic descriptor IE, as appropriate, the network will progress the connection establishment request onward. If the network is unable to provide the traffic parameter values specified in the ATM

traffic descriptor IE, the alternative ATM traffic descriptor IE, or the minimum acceptable ATM traffic descriptor IE, as appropriate, the network will reject the connection establishment request.

If the switch accepts the call, it returns a CALL PROCEEDING message providing the proper VCI/VPI information to the calling party or previous hop switch. This continues

through each hop until the destination is reached. At the destination, the destination switch returns a CONNECT message that propagates hop-by-hop back to the calling party. During each propagation the appropriate VCI/VPI tables are updated. The disestablishment of connections is the inverse of connection establishment. It follows the same general procedures but uses RELEASE and RELEASE COMPLETE messages originated by the called party/destination switch instead of the SETUP, CALL PROCEEDING, and CONNECT messages used by the calling party/originating switch to establish the connection.

Figure 3-1 provides a simplified depiction of the signalling message flow call control sequence for establishment of SVC connections in ATM networks.
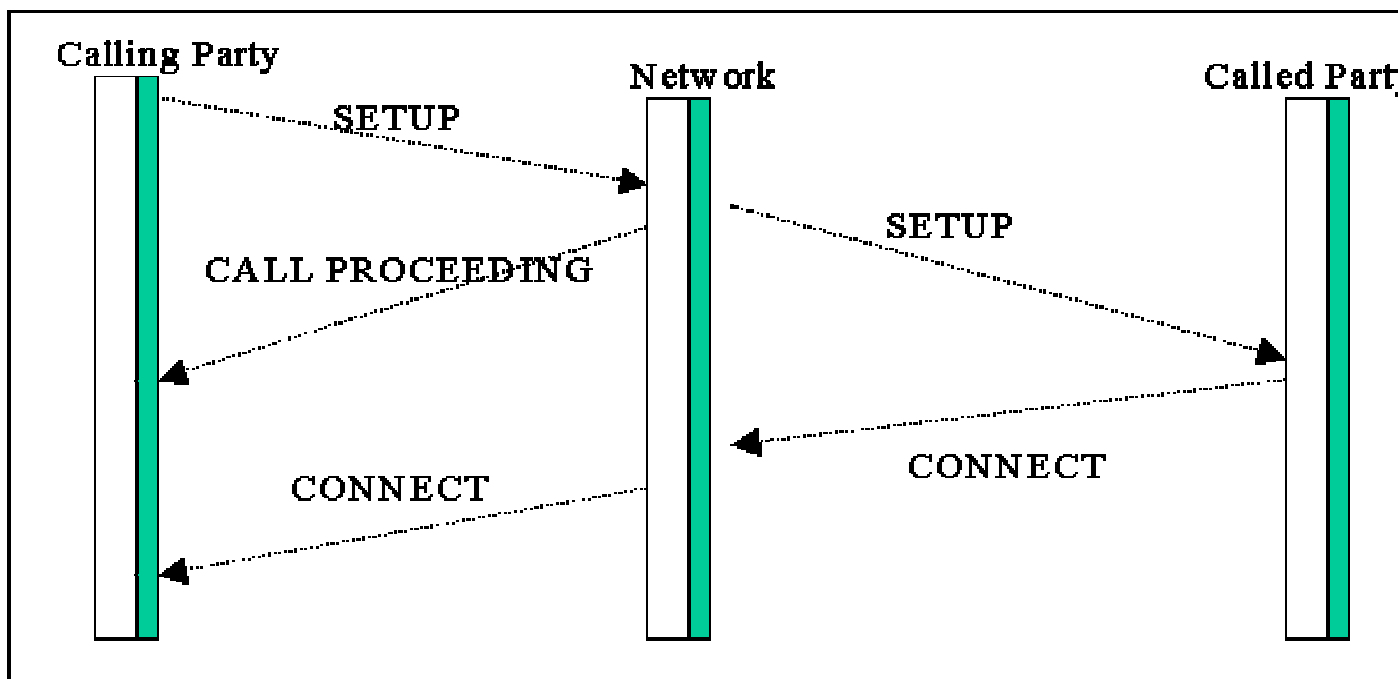

Figure 3-1. Call Control Sequence for SVC Connections

Return to Table the of Contents

3.6 NETWORK MANAGEMENT	

Network management is the monitoring, analysis, evaluation, and control of network performance to ensure that the end user is provided an agreed upon level of performance, and ensure optimal use of network resources. The seamless architecture and mix of traffic service requirements in ATM networks pose a unique challenge for network management systems in just about every area of network management. Following encapsulation, in ATM networks management of encapsulated traffic is treated the same as all other traffic in the same category. The ATM Forum's Network Management Working Group has generated a framework to address network management functionality in three areas: (1) a five layer ATM Management model to be used for managing ATM networks and services; (2) an Interim Local Management Interface (ILMI) used as an interim specification for network management functions between an end user and a public or private network, and between a public network and a private network based on a limited set of Simple Network Management Protocol (SNMP) capabilities; and (3) an operations, administration, and maintenance (OAM) facility used for end-to-end circuit management. The Network Management Working Group is currently considering a fourth area for incorporation in the frameworkone that involves signaling capabilities to embed management with the signaling protocol.

The Network Management model, often referred to as "M specs," addresses management interfaces for both private and public network services. The "M specs" consist of five management interfaces labeled M1 through M5. The following interfaces collectively provide end-to-end ATM service monitoring and control information: (1) M1 provides information for the management of ATM end devices; (2) M2 provides information for the management of private ATM networks or switches; (3) M3 provides information for the management of links between public and private networks; (4) M4 provides information for the management of public ATM networks; and (5) M5 provides information for the management of links between two public networks. Figure 3-2 is a simplified depiction of the above ATM network management interfaces.
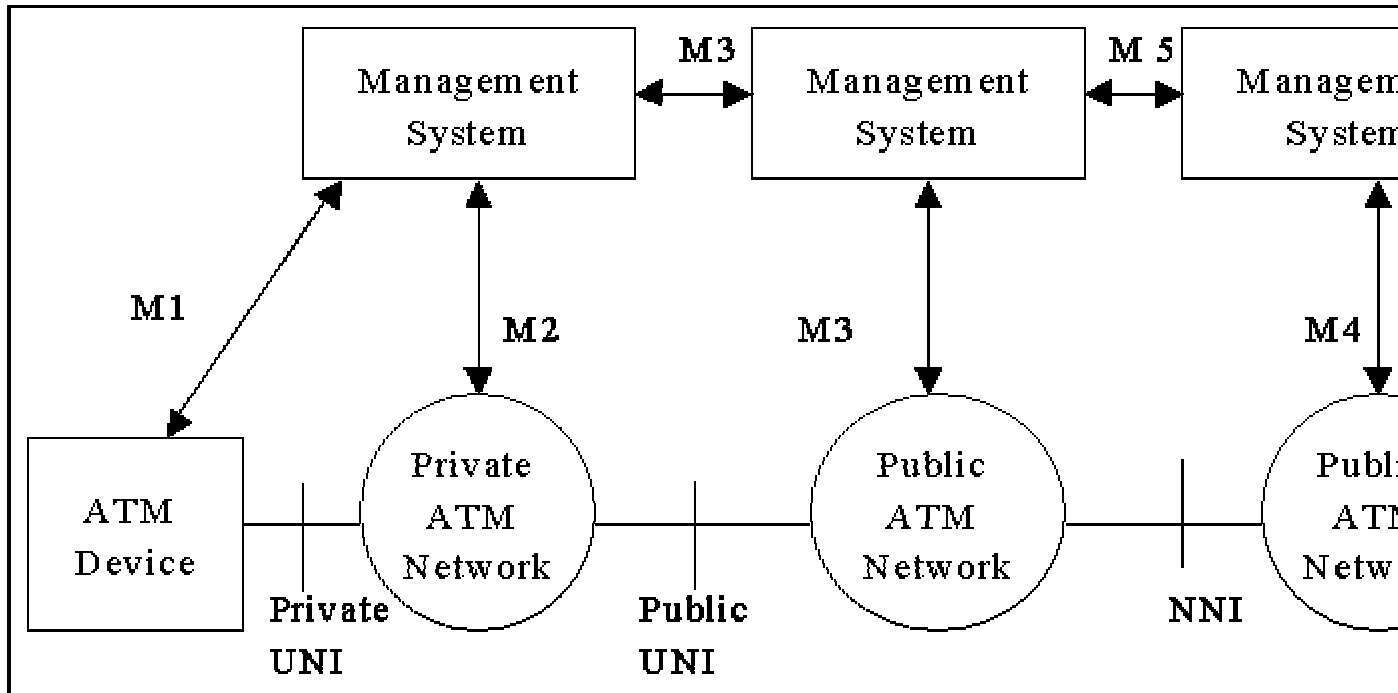
Figure 3-2. ATM Network Management Interfaces

The ILMI specification provides address registration, status, configuration, and control information for an ATM interface using SNMP for monitoring and control operations of ATM management information across the UNI. It is considered interim in the sense that it will

eventually be incorporated (or even phased out) in favor of standard "M specs" when completed. Along with performance information, an important ILMI function is to support user configuration of VP and VC connections. ILMI attributes are organized in a standard Management Information Base (MIB) structurea MIB defines management items for network components that can be accessed by a network manager, and various fields and values that the agent software must keep track of in the managed device.

SNMP is a connectionless protocol originally designed for the Department of Defense network to support TCP/IP. It is the IETF standard management protocol. A limited set of SNMP capabilities is used by network managers to manage objects modeled in the MIB.

The use of OAM cells by the OAM facility reduces the amount of management-related traffic

and minimizes the need to distribute MIBs by providing ATM network devices the capability to gather information about end-to-end connections. The ATM Forum has specified several OAM cells targeting fault management. Included are alarm indication

signal (AIS) cells, far end

reporting failure (FERF), and a loopback capability to ascertain whether idle connections are still up or have failed.

3.6.1 Network Management Operations	

Effective coordination of network management operations requires that all elements involved in the network view network management information identically. The network management information maintained in the MIB is modeled in terms of managed objects. Managed objects are abstractions of data communications resources used to model information about the object and its role in the network. Included in this information are characteristics such as its name, the events it may generate, and the operations it may be requested to perform. As with other telecommunications networks, network management operations for networks supporting IP over ATM can generally be organized into the following functional categories:

Configuration Management: The configuration management function is performed to ensure that the build state of the network is ascertainable and recordable by the execution of positive measures to exercise control over, identify, and collect data from and provide data to required network elements (NEs). In an ATM network configuration management deals with the allocation of bandwidth to VP and VC connections, the creation and deletion of VPs and VCs, and QoS provisioning.

Performance Management: In an ATM network, because of the increased number of connections that need to be managed, massive amounts of statistics must be collected to accurately gauge the performance of the network. The performance management function evaluates and reports upon the behavior of telecommunications equipment and the effectiveness of the network or network elements. From an ATM perspective this involves performance monitoring at the physical and ATM layers, traffic management, traffic monitoring and control, performance management control, and network data collection.

Fault management: The fault management function is concerned with the reporting of relevant events and conditions which occur within the network, and with defect detection and fault localization. Network defects include physical layer defects, loss of cell delineation, and loss of continuity (LOC). It is possible for defects to occur at the ATM connection level without the loss of the transmission path (e.g., a VPI/VCI routing table in an ATM node could become corrupted). Fault management institutes measures such as continuity checks, and the use of AIS

and remote defect indication (RDI) cells to detect ATM LOC defects and carry information to indicate the type and location of the detected defect to assist in fault localization and recovery.

Accounting Management: The large volumes, coupled with different types of traffic generated in an ATM network, place a great burden on accounting and billing. Private organizations, as well as carriers, need to determine and map network traffic to individual users. The accounting management function is concerned with the operation and maintenance of a billing system based on actual usage to generate accurate cost accounts. Automatic message accounting and customized billing are high on the list of measures that are rapidly being implemented by vendors in their accounting management solutions.

Security management: The security management function is performed to control access to the network and protects both the network and the network management subsystem against intentional or accidental abuse, unauthorized access and communication loss. Security management involves encryption, user authentication, password protection, and authorized usage of the network resource. Currently, ATM lacks any provisions for security and authentication. The IETF has identified security as a major interest area and is currently active in addressing this requirement.

LIST OF REFERENCES

1. Smith, T., and Armitage, G., October 1997, *IP Broadcast over ATM Networks,* Internet Engineering Task Force, Reston, VA.

2. Heinanen, J., July 1993, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, Internet Engineering Task Force RFC 1483, Telecom Finland.

3. Sonia Fahmy, et al., The ATM Forum, December 1997, *A Switch Algorithm for ABR Multipoint-to-Point Connections*, ATM Forum/97-1085R1, The ATM Forum, Mountain View, CA.

4. Huitema, C., 1996, *IPv6: The New Internet Protocol*, Prentice-Hall PTR, Upper Saddle River, NJ.

5. Larson, R., December 1995, *Methods of Implementing IP on ATM Networks*, Odyssea Systems Corporation.

6. The ATM Forum Technical Committee, July 1996, *ATM User-Network Interface (UNI) Signalling Specification - Version 4.0*, The ATM Forum, Mountain View, CA.

7. The ATM Forum Technical Committee, September 1996, *Integrated Local Management Interface (ILMI) Specification - Version 4.0*, The ATM Forum, Mountain View, CA.

8. The ATM Forum Technical Committee, August 1993, *Data Exchange Interface (DXI) Specification, Version 1.0 (af-dxi-0014.000)*, The ATM Forum, Mountain View, CA.

9. Information Science Institute, September 1981, *Transmission Control Protocol DARPA Internet Program Protocol Specification*, University of Southern California, Marina del Rey, CA.

10. Laubach, M., January 1994, Classical IP and ARP over ATM, Internet Engineering Task Force RFC 1577, Hewlett - Packard Laboratories.

11. Datapro Information Services Group, 1997, *LAN Emulation and Multiprotocol Over ATM*, McGraw-Hill, Inc., New York, NY.

12. Demopolous, D., October 1997, "The Converging Worlds of Layer 2 and Layer 3 Switching," *Telecommunications,* Horizon House, Norwood, MA.

13. Peterson, D., 1995, *TCP/IP Networking: A Guide to the IBM Environment,* McGraw-Hill, Inc., New York, NY.

14. Goldman, J., 1995, *Applied Data Communications,* Purdue University, John Wiley & Sons, New York, NY.

15. Carne, E., 1995, *Telecommunications Primer: Signals, Building Blocks and Networks,* Prentice Hall PTR, Upper Saddle River, NJ.

16. Kumar, B., 1994, *Broadband Communications: A Professional's Guide to ATM, Frame Relay, SMDS, SONET, and B-ISDN*, McGraw-Hill, Inc., New York, NY.

17. Freeman, R., 1996, *Telecommunications Systems Engineering: Third Edition*, John Wiley & Sons, New York, NY.

18. SETA Corporation, June 1997, *High Speed Networks Analysis Report: Section 2 - Call/Congestion Establishment and Congestion Control in ATM Networks,* SETA Corporation, McLean, VA.

19. Callon, R., et al., April 15-19, 1996, *Issues and Approaches for Integrated PNNI*, ATM Forum/96-0355, The ATM Forum, Mountain View, CA.

20. Onvural, R. O., 1994, *1994 Transport Protocols*, Datapro report, Artech House, Inc., New York, NY.

21. ATM Forum Technical Committee, October 1996, *MPOA Connection Management*, ATM Forum 96-1408, Montreux, Switzerland.

22. White, P., May 1997, *RSVP and Integrated Service in the Internet: A Tutorial*, IEEE Communications Magazine, New York, NY.

23. ATM Forum Technical Committee, February 5-9, 1996, *Issues in Extending Unicast and Multicast RSVP Flows Across ATM Networks*, ATM Forum/96-0094, Joint BOF (SAA, MPOA, LANE), The ATM Forum, Mountain View, CA.

24. Lea, Chin-Tau, April 1993, *A Multicast Broadband Packet Switch*, IEEE Transactions on Communications, New York, NY.

25. Anderson, J., et. al., December 1996, *Operations Standards for Global ATM Networks: Network Element View*, IEEE Communications Magazine, New York, NY.

26. Datapro Information Services Group, 1997, *Transmission Control Protocol/Internet Protocol (TCP/IP)*, McGraw-Hill, Inc., New York, NY.

27. Naugle, M., 1994, *Network Protocol Handbook*, McGraw-Hill, Inc., New York, NY.

28. ATM Forum Technical Committee, September 1994, *Beyond Classical IP-Integrated IP and ATM Architecture Overview*, ATM Forum,/94-0935, The ATM Forum, Mountain View, CA.

29. ATM Forum Technical Committee, June 8, 1994, *IP with ATM (IpwATM: Efficient IP- ATM Interworking Across LANs and WANs*, ATM Forum/94-0558, The ATM Forum, Mountain View, CA.

30. Matusow, D., 1996, *Methods of Integrating SNA and IP*, Datapro Information Services Group, McGraw-Hill, Inc., New York, NY.

31. Armitage, G., November 1996, Support for Multicast over UNI 3.0/3.1 based ATM Networks, Internet Engineering Task Force RFC 2022, Bellcore.

32. Shaver, T., June 10-14, 1996, *NSAP & Public E.164 Address Interworking Issues*, ATM Forum/96-0809, The ATM Forum, Mountain View, CA.

33. Datapro Information Services Group, 1997, *Technology Overview: B-ISDN Physical, ATM, and AAL Layers*, McGraw-Hill, Inc., New York, NY.

34. Desai, V., 1997, *Managing ATM,* Datapro Information Services Group, McGraw-Hill, Inc., New York, NY.

35. Freeman, R., 1995, *Practical Data Communications*, John Wiley & Sons, Inc., New York, NY.

36. ATM Forum Technical Committee, March 1994, *Proposal for Network Management Requirements of LAN Emulation/Interconnect Over ATM Services,* ATM Forum/94-0242, The ATM Forum, Mountain View, CA.

37. ATM Forum Technical Committee, May 1994, *PNNI & ILMI Requirements for Auto ATM Node Discovery by SNMP-based NMS*, ATM Forum/94-0499, The ATM Forum, Munich, Germany.

38. ATM Forum Technical Committee, June 1996, *Enhanced Lane Multicast Protocol*, The ATM Forum/96-0861, The ATM Forum, Orlando, FL.

39. ATM Forum Technical Committee, December 1997, *Requirements for Native Connectionless Service in ATM Networks*, The ATM Forum, Red Bank, NJ.

40. ATM Forum Technical Committee, April 1995, *Addressing and Peer Group Identifiers*, ATM Forum/95-0479, The ATM Forum, Denver, CO.

41. ATM Forum Technical Committee, April 1997, *Distributed Mechanisms for VCI Assignment in VP-based Multicast,* ATM Forum/97-0379, The ATM Forum, Chicago, IL.

42. ATM Forum Technical Committee, June 1996, *NSAP & Public E.164 Address Interworking Issues,* ATM Forum/96-0809, The ATM Forum, Orlando, FL.

ACRONYMS

AAL ATM Service Provider

AIS Alarm Indication Signal

ARP Address Resolution Protocol

ATM Asynchronous Transfer Mode

ATMARP ATM Address Resolution Protocol

B-ISDN Broadband-Integrated Services Digital Network

BUS Broadcast and Unknown Server

CBR Constant Bit Rate

CLNP Connectionless-mode Network Protocol

CLP Cell Loss Priority

CPCS Common Part Convergence Sublayer

CPE Customer Premises Equipment

CRC Cyclic Redundancy Check

CS Convergence Sublayer

CSU Channel Service Unit

DCE Data Communication Equipment

DFA DXI Frame Address

DTE Data Terminal Equipment

DXI Data Exchange Interface

E.164 Public network addressing standard

ESI End System Identifier

FERF Far End Reporting Failure

FTP File Transfer Protocol

GFC Generic Flow Control

HEC Header Error Control

IDU Interface Data Unit

IEEE Institute of Electrical and Electronic Engineers

IETF Internet Engineering Task Force

IE Information Element

ILMI Integrated Local Management Interface

IP Internet Protocol

Ipv6 Internet Protocol Next Generation

LAN Local Area Network

LANE LAN Emulation

LIS Logical IP Subnetwork

LMI Local Management Interface

LOC Loss Of Continuity

MAC Media Access Control

MARS Multicast Address Resolution Server

MCS Multicast Server

MIB Management Information Base

MPOA Multiprotocol Over ATM

N6 OMNCS Technology and Standards Division

NE Network Element

NEI Network Element Identifier

NEML Network Element Management Layer

NIST National Institute of Standards and Technology

NNI Network-Network Interface

NPC Network Parameter Control

NSAP Network Service Access Point

NS/EP National Security and Emergency Preparedness

OAM Operation, Administration, and Maintenance

OMNCS Office of the Manager, National Communications System

OSI Open Systems Interconnection

PDU Protocol Data Unit

PNNI Private Network-Network Interface

PT Payload Type

PTI Payload Type Identifier

PVC Permanent Virtual Circuit

QoS Quality of Service

RDI Remote Defect Indication

RFC Request For Comment

RSVP Resource Reservation Protocol

SAP Service Access Point

SAR Segmentation and Reassembly

SDU Service Data Unit

SEL Selector

SMTP Simple Mail Transfer Protocol

SNMP Simple Network Management Protocol

SSCS Service Specific Convergence Sublayer

STD Source Traffic Descriptor

SVC Switched Virtual Circuit

TCP Transmission Control Protocol

TCPDU TCP Data Unit

TCP/IP Transmission Control Protocol/Internet Protocol

TDM Time Division Multiplex

UDP User Datagram Protocol

ULP Upper-Layer Protocol

UNI User-Network Interface

UPC Usage Parameter Control

VBR Variable Bit Rate

VC Virtual Channel/Circuit

VCC Virtual Channel Connection

VCI Virtual Channel Identifier

VP Virtual Path

VPC Virtual Path Connection

VPI Virtual Path Identifier

VPL Virtual Path Link