CHAPTER 1

ATM Technology

This chapter describes the Asynchronous Transfer Mode (ATM) technology on which the LightStream 2020 multiservice ATM switch (LS2020 switch) is based. If you are already familiar with ATM, you can skip this chapter.

ATM is a communications standard based on cell relay techniques. The next sections discuss cell relay and ATM technology. They also contrast ATM techniques with time-division multiplexing (TDM) and other packet-handling technologies.

Cell Relay Packet Handling

Cell relay is a flexible and responsive method for multiplexing all forms of digital traffic (data, voice, image, and video). Cell relay can handle rapid changes in the quantity and pattern of the traffic in the network. All traffic is placed in fixed-length packets of information (cells) and switched at high speeds. Cell relay is generally acknowledged as the best multiplexing technology for modern communication applications because it combines the strengths of TDM and conventional packet switching. Using cell relay packet-handling techniques, a mixture of bursty and delay-sensitive traffic can be processed simultaneously, while at the same time providing the services required by each traffic type.

Also, because cell relay processing is based on the use of small packets, the process technology is adaptable and cost effective for a wide range of interface speeds.

Technologies Compared

ATM technology first appeared in the Broadband Integrated Services Digital Network (BISDN). However, ATM is now recognized as a useful technology in and of itself and is based on the specifications and standards being developed by ITU-T (International Telecommunications Union Telecommunication Standardization Sector), ANSI (American National Standards Institute), and the ATM Forum.

Note The ITU-T carries out the functions of the former Consultative Committee for International Telegraph and Telephone (CCITT).

Each ATM cell contains a header and the data to be transferred. Cells are switched in the network based on routing information contained in the cell headers. ATM transports all types of traffic (data, voice, image, and video) using the same cell format.

ATM contrasts with TDM in the way it allocates communications channels. In TDM, communications channels are divided into fixed periods of time called frames. The frames are divided into a fixed number of time slots of equal duration (see Figure 1-1). Each user is assigned certain time slots within each frame. As Figure 1-1 indicates, a user can be given more than one time slot in a frame.

Figure 1-1 User Assignments on Communications Channel Using TDM

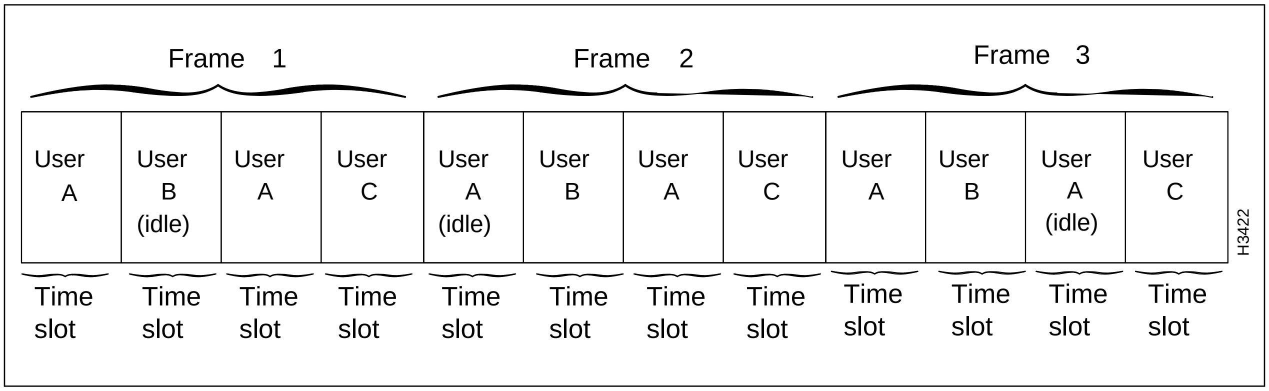

The time slots allocated for each user occur at precisely the same time in every frame. Since the time slots are synchronous, TDM is sometimes referred to as Synchronous Transfer Mode (STM).

Users can access the communications channel only when a time slot that has been allocated to them is available. For example, User A can send messages over the communications channel only during the time slot(s) designated for User A. If no traffic is ready to send when the designated time slot occurs, that time slot is unused. If a user has a burst of traffic that exceeds the capacity of the designated time slots, additional slots cannot be used, even if they are idle. As a result, a long delay could result before the burst of traffic is transferred over the TDM network.

In ATM, access to the communications channel is more flexible**.** Any user needing the communications channel can use it whenever it is available. In contrast to TDM, ATM imposes no regular pattern on the way users are given access to the communications channel. ATM is also described as providing *bandwidth on demand.*

In other packet-handling technologies, such as High-Level Data Link Control (HDLC), any user can gain access to the communications channel, but a user who has a long message to send can prevent other users from gaining access to the channel until the entire message has been passed. However, with ATM, every message is divided into small, fixed-length cells. Thus, no single user can monopolize access to the communications channel while other users have messages to send (see Figure 1-2).

Figure 1-2 User Assignments on ATM Communications Channel

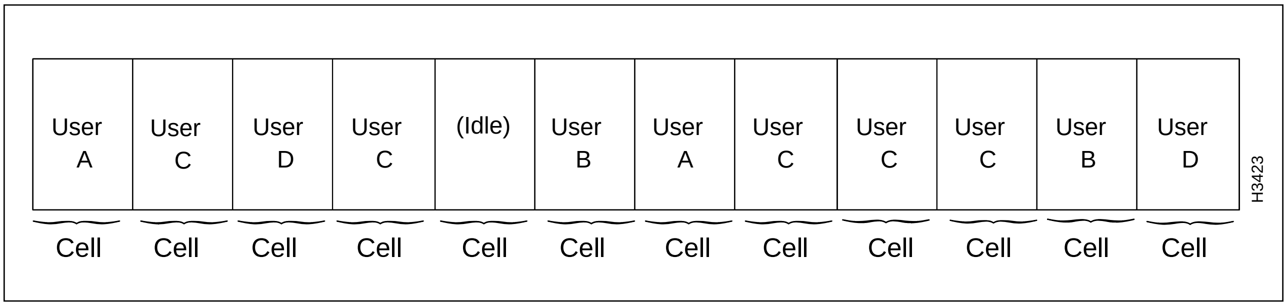

Benefits of ATM

ATM technology offers four principal benefits:

- **Bandwidth efficiency**—By allowing any user access to the network whenever resources are available, ATM provides better use of bandwidth for bursty traffic. By breaking traffic into small, fixed-length cells, ATM prevents a user with a long message from blocking access to the network.
- **Scalable technology—**ATM is flexible and accommodates a wide range of traffic rates and applications. ATM interface standards exist for data rates as low as 1.5 Mbps and as high as 1.2 Gbps.
- **Application transparency**—The ATM cell size is a compromise between the long frames of data applications and the short, repetitive frames required in voice applications. Due to its asynchronous nature, ATM supports traffic at rates and degrees of burstiness compatible with the applications being run, not at rates convenient to the network.
- **Networking advantages**—ATM is a simple, fast, switching and routing process based on the virtual channel identifier (VCI) in the cell address. Within the network, no processing occurs above the cell level, thus simplifying and increasing message-handling speed. ATM message-handling techniques foster the creation of efficient, self-routing switches that can grow in size and speed to meet the user's future communications requirements.

Fitting ATM into the OSI Model

ATM standards define protocols that operate at Layer 2 (the data link layer) of the International Standards Organization (ISO) seven-layer Open Systems Interconnection (OSI) reference model. Figure 1-3 shows the layered architecture of the OSI model.

Figure 1-3 OSI Model

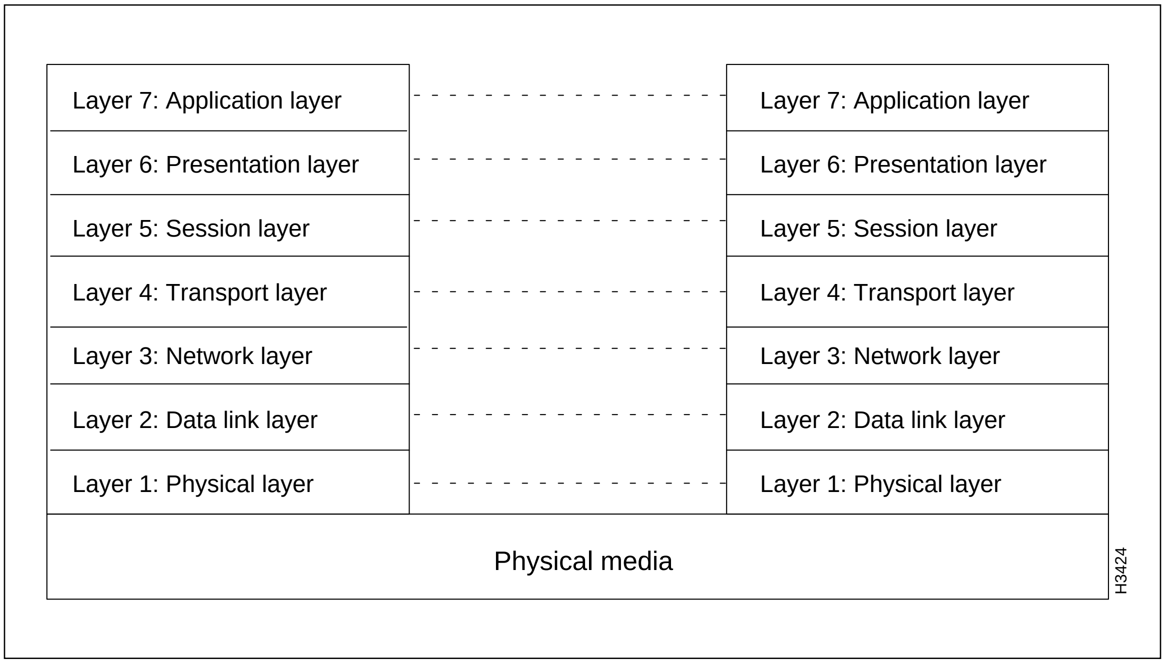

The data link layer is concerned with data transmission between two network switches. This layer is not concerned with the transmission of an entire message between a source and a destination switch—this responsibility belongs to Layer 3 (the network layer). Rather, the data link layer transports portions of messages (cells, in the case of ATM) between two points in the network. These points may be the source and the destination of the message, or they may be only intermediate hops between the source and the destination.

The data link layer may divide higher level data into smaller units (cells, in this case), whose sizes are compatible with overall network requirements. Layer 2 data units contain a cell header, an information field, and some method of checking for transmission errors.

Placing User Data into ATM Cells

Before frames can be transported across an ATM network, they must be divided into ATM cells. The processes that divide the frames into cells occur at the data link layer (Layer 2). Layer 2 is divided into two parts: the ATM adaptation layer (AAL) and the ATM layer. Once frames are divided into ATM cells, the cells can be transferred to Layer 1, the physical layer (see Figure 1-4).

Figure 1-4 Layer 2 Data Link Layer

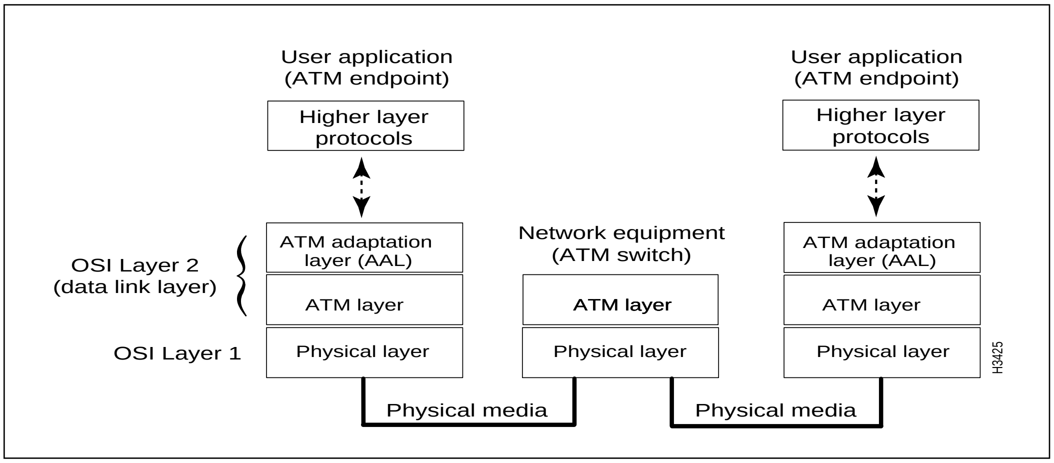

ATM Label Switching

ATM uses label switching, a technique in which a simple label is placed in the header of each cell. The label provides information used in transporting the cell across the next hop in the network. Networks that do not use label switching usually require each packet (or cell) to contain the explicit address of the final destination. ATM uses label switching because it is simpler, thereby making faster switching possible.

Here is how label switching works:

1. A switching unit reads an incoming cell from a particular port. The incoming cell has a routing label.
2. The switching unit uses the combination of the input port on which the cell was received and the information in the label to determine where the cell should go next. It does this by referring to a routing table that correlates the incoming port and label with an outgoing port and label.
3. The switch replaces the incoming label with a new outgoing label and sends the cell through the outgoing port, which is connected to another switching device. (The new outgoing label is taken from the routing table.)
4. This process is repeated until the cell reaches its final destination in the ATM network.

For example, suppose that your network includes a switching unit called Boston. A number of data paths go through the Boston switch. When those data paths are created, a routing table is set up within the Boston switch. The table in the Boston switch has one entry for every data path that goes through the switch. The entries in the table map the incoming port and label to an outgoing port and label for each data path, as shown in Table 1-1.

Table 1-1 Sample Routing Table in Boston Switch

Port In	Label In	Port Out	Label Out
1	L	6	Z
1	M	7	X
2	N	7	Y

When the Boston switch receives an incoming cell on port 1 with label M, it consults the routing table and finds that label M should be replaced with label X, and that the cell should be passed out of the Boston switch on port 7. The cell is then transported to the switch in the network that is connected to port 7 of the Boston switch, as shown in Figure 1-5.

Figure 1-5 Cell Passing through Boston Switch

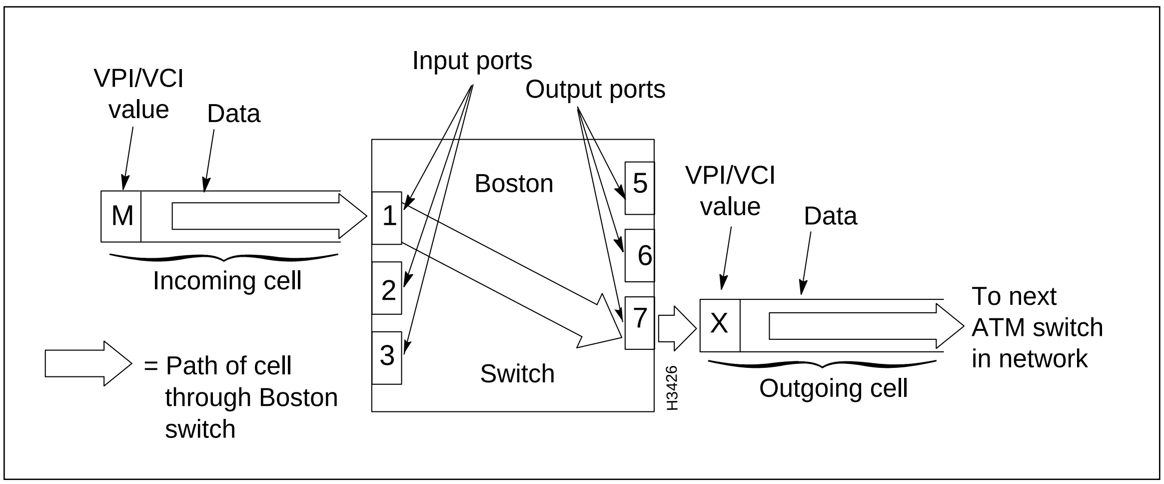

In all cases, transporting cells through the use of label switching requires a connection. Information about the connections is provided in the routing tables (sometimes called lookup tables) of switching and multiplexing units. ATM uses virtual channel connections and virtual paths to accomplish routing functions.

Virtual Channel Connections and Virtual Paths

A virtual channel connection (VCC) is a series of virtual channel links (VCLs) between two ATM points. A VCL is a means of bidirectional transport of ATM cells between a point where a virtual channel identifier (VCI) value is assigned and the point where the same value is either reassigned or terminated. The VCI identifies the VCL to which a cell belongs and determines where the cell should go next. Figure 1-6 shows the relationship between VCLs and VCCs in an ATM network.

Figure 1-6 Relationship between VCLs and VCCs in ATM Network

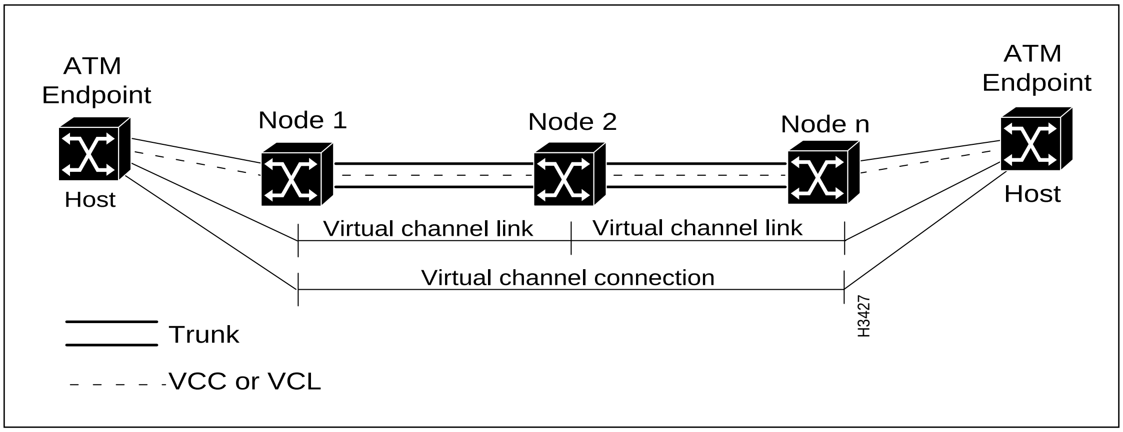

VCCs are sometimes transported within virtual paths (VPs). A VP is identified by its virtual path identifier (VPI). VPs provide a convenient way of bundling traffic directed to the same destination or traffic requiring the same Quality of Service (QoS) in the network (see Figure 1-7).

Figure 1-7 VCCs Transported within VPs

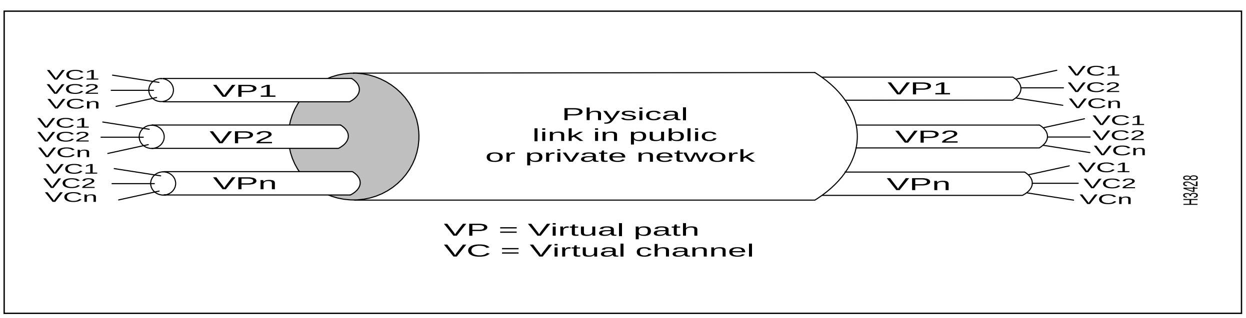

The ATM Cell

The ATM cell is the fixed-length transmission unit defined by the ATM standard. An ATM cell contains two major types of information: the payload and the header. The payload is the information to be transferred through an ATM network. It can include data, voice, image, or video. The header is the information used to route the cell through the network and to ensure that the cell is forwarded to its destination.

Every ATM cell is 53 bytes long. The first 5 bytes contain header information, and the remaining 48 bytes contain the payload (see Figure 1-8).

Figure 1-8 ATM Cell

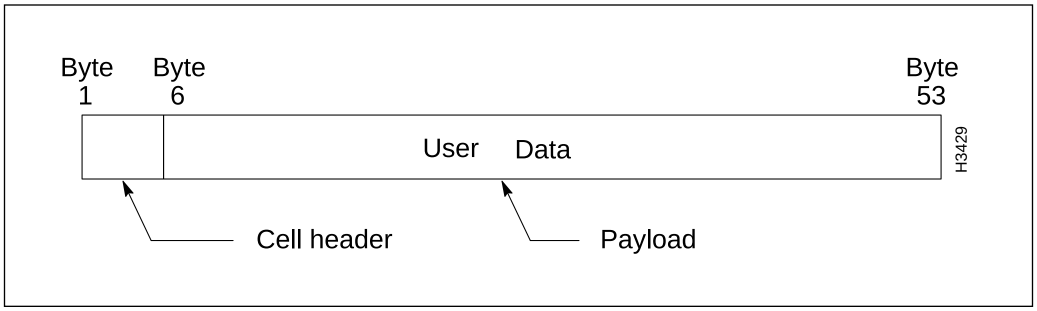

The 5-byte header (see Figure 1-9) contains several different fields (seeTable 1-2). The 48 bytes following the header (the payload) contain user data.

Figure 1-9 UNI ATM Cell Header Format

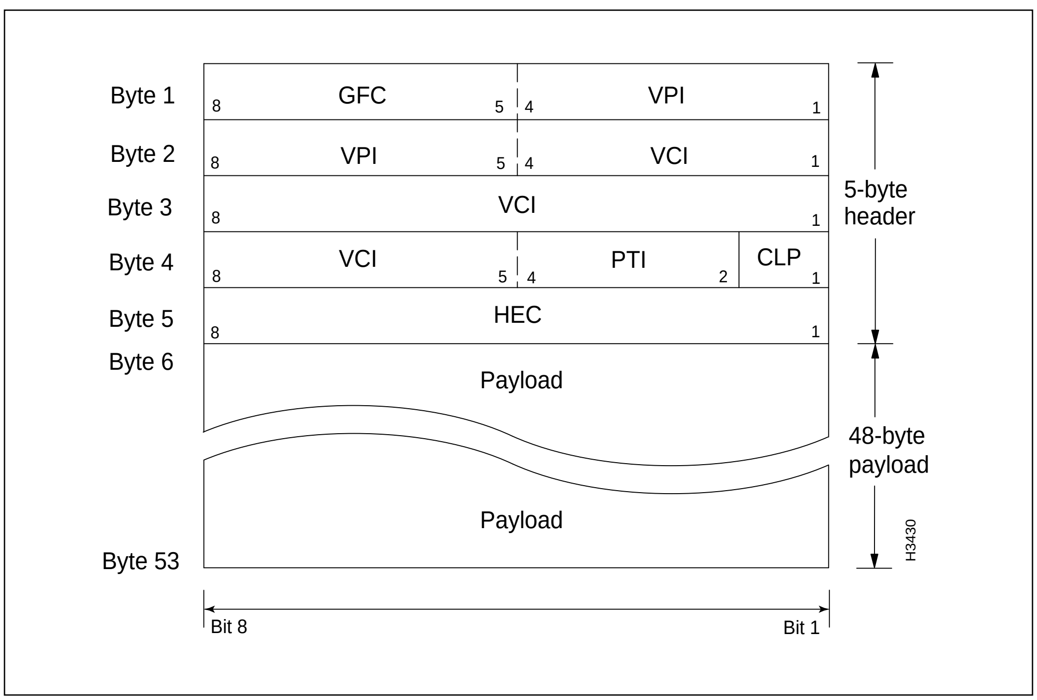

Table 1-2 Fields in ATM Cell Header

Header Field Name	Location in Header	Description
GFC (Generic Flow Control)[1]	First four bits of Byte 1	Controls the flow of traffic across the user network interface (UNI) and thus into the ATM network.
VPI (Virtual Path Identifier)	Second four bits of Byte 1 and the first four bits of Byte 2	Identifies a particular VPC. A VPC is a group of virtual connections carried between two points and may involve several ATM links. VPIs provide a way to bundle traffic heading to the same destination.
VCI (Virtual Channel Identifier)	Second four bits of Byte 2, Byte 3, and the first four bits of Byte 4	Identifies a particular VCC. A VCC is a connection between two active, communicating ATM entities. The VCI consists of a concatenation of several ATM links.
PT (Payload Type)	The fifth, sixth, and seventh bits of Byte 4	Indicates the type of information in the payload field. ATM cells carry different types of information that may require different handling by the network or terminating equipment.
CLP (Cell Loss Priority)	Eighth bit of Byte 4	Indicates the cell loss priority set by the user. This bit indicates the eligibility of the cell for discard by the network under congested conditions. If the bit is set to 1, the cell may be discarded by the network if congestion occurs.
HEC (Header Error Control)	Byte 5	Contains an error-correcting code calculated across the previous four bytes of the header. The HEC detects multiple-bit header errors and can be used to correct single-bit errors. The HEC provides protection against incorrect delivery of messages caused by address errors. The HEC does not provide any protection for the payload field itself.

1. For a network-to-node (NNN) interface, there is no GFC field. These four bits are part of the VPI field.

The ATM Adaptation Layer

The AAL accepts frames from higher OSI layers and adapts them to the 48-byte segments that are placed into the payload field of ATM cells. The ATM layer accepts the 48-byte segments, adds the 5-byte header, and produces ATM cells to be transferred to the physical layer, as illustrated in Figure 1-10.

Figure 1-10 ATM Adaptation Layer Functions

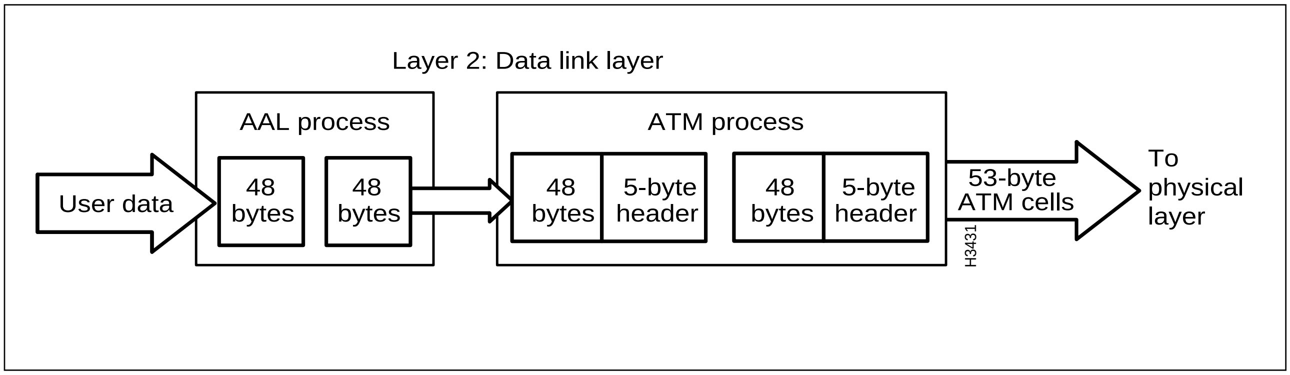

When ATM cells are transferred through a network, each cell is processed in isolation from all other cells. All processing decisions are made based on the cell header; no processing of the data in the payload field occurs.

Figure 1-11 presents some examples of AAL processing.

Figure 1-11 AAL Processing Examples

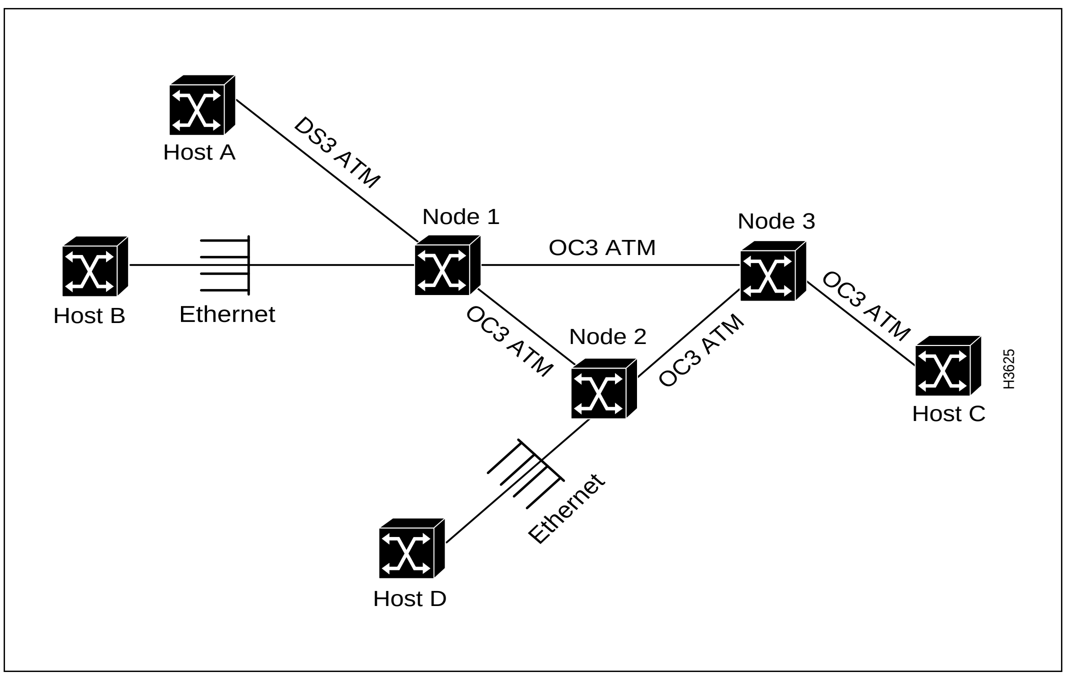

Hosts A and C are connected to the network through ATM interfaces, so they do all their AAL processing internally. The network does not do any processing for hosts A and C. Hosts B and D are connected to native Ethernet interfaces on Nodes 1 and 2. Therefore, Node 2 does all the AAL processing for Host D. Node 3 does no AAL processing.

Depending on the type of traffic entering the ATM network, the AAL uses one of four different AAL types to divide the traffic into small segments. These types are classified according to the timing relationship between the source and destination, the constant or variable bit rate, and the mode (connection-oriented or connectionless). The AAL types defined in the ATM standard are listed in Table 1-3.

Table 1-3 AAL Types

AAL Type	Examples of Traffic Type
1	Circuit emulation, constant bit rate video
2	Variable bit rate video and audio
3/4	Connection-oriented or connectionless data transfer (AAL 3/4 has cell-by-cell error checking and multiplexing)
5	Connectionless data transfer (AAL 5 has lower overhead than AAL 3/4)

The AAL is divided into two sublayers: the convergence sublayer (CS) and the segmentation and reassembly sublayer (SAR), (see Figure 1-12).

Figure 1-12 Information Flow Through AAL

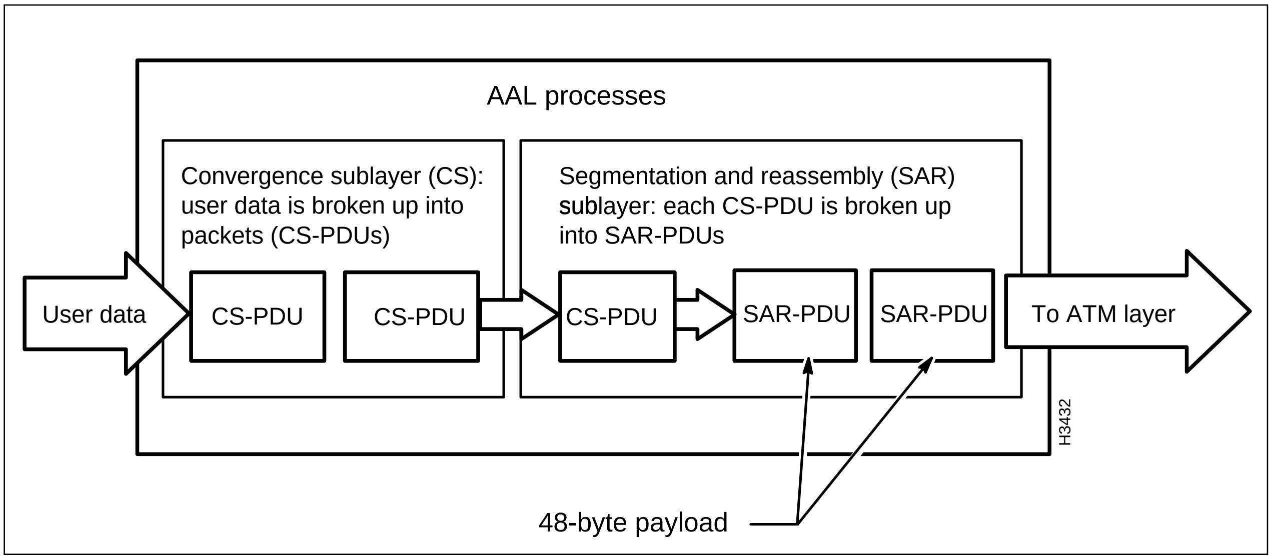

The convergence sublayer (CS) accepts higher layer traffic for transmission across the network. Depending on the AAL type, header and/or trailer fields are added to the packet. The packet is then segmented by the SAR sublayer to form 48-byte payloads (also known collectively as SAR-PDUs).

Upon receipt of cell payloads, the AAL removes any AAL-specific information from each payload and reassembles the entire packet before passing it to a higher layer (see Figure 1-13).

Figure 1-13 SAR Portion of AAL Process

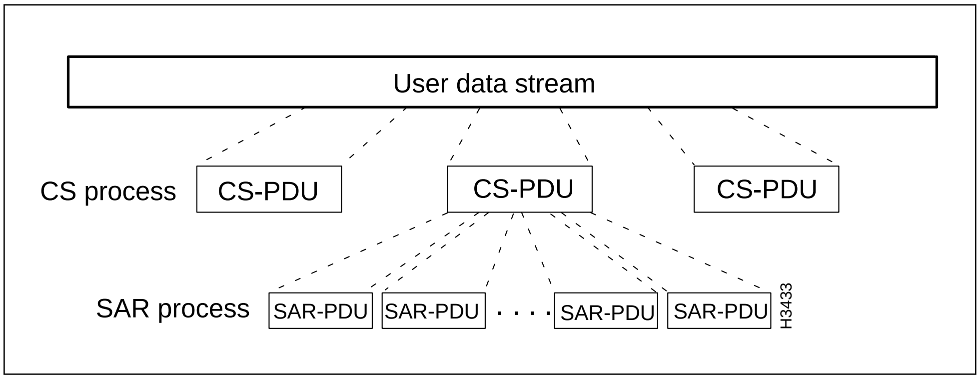

The ATM Layer

The ATM layer accepts the 48-byte SAR-PDUs from the SAR process, adds a 5-byte header to each, and produces ATM cells for transfer to the physical layer (see Figure 1-14).

Figure 1-14 ATM Layer Process

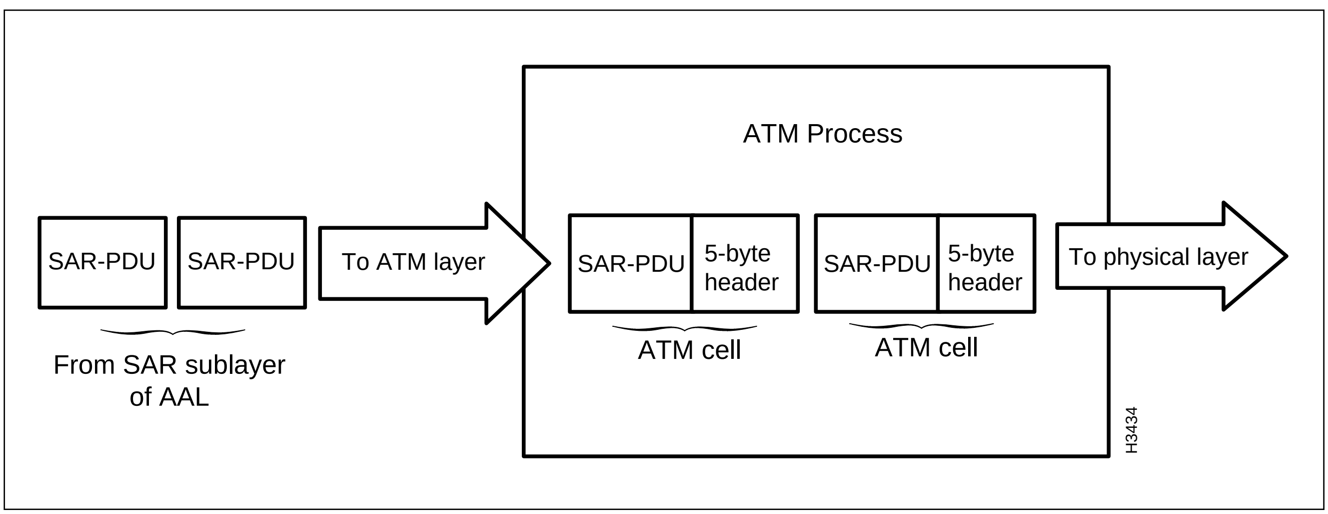

Placing Cells onto Physical Transport Medium

Once the data is packaged into 53-byte ATM cells, the cells are transferred to the physical layer, where they are placed onto a physical transport medium, such as fiber optic cable or coaxial cable. The process of placing cells onto the physical medium takes place in two sublayers: the physical medium dependent (PMD) sublayer and the transmission convergence (TC) sublayer.

Each PMD is specific to a particular physical medium and includes definitions of proper cabling as well as bit timing. The TC sublayer generates and receives transmission frames and performs all overhead functions associated with the transmission frame. The TC sublayer performs a convergence function by receiving a bit stream from the PMD and extracting cells.

Although PMD operation depends on the physical medium, the following TC functions remain common to all physical layers:

- Cell delineation—Extraction of cells from the bit stream received from the PMD
- Cell rate decoupling—Adaptation of the speed of the ATM layer cell stream to the rate of the physical interface
- Header error control (HEC) generation and checking—Performed when the TC sublayer checks where each received cell starts and ends by calculating the HEC for that cell
- Various operation and maintenance (OAM) functions