



ATM Technology Fundamentals

This chapter provides a brief overview of ATM technology. It covers basic principles of ATM, along with the common terminology, and introduces key concepts you need to be familiar with when configuring ATM network equipment. If you already possess this basic knowledge, you can skip this chapter and go on to Chapter 2, “ATM Signaling and Addressing.”



Note

This chapter provides only generic ATM information. Subsequent chapters in this guide include implementation-specific information for the Catalyst 8540 MSR, Catalyst 8510 MSR, and LightStream 1010 ATM switch routers.

This chapter includes the following sections:

- What is ATM?, page 1-1
- ATM Basics, page 1-2
- Traffic Contracts and Service Categories, page 1-12
- Common Physical Interface Types, page 1-15

What is ATM?

Asynchronous Transfer Mode (ATM) is a technology designed for the high-speed transfer of voice, video, and data through public and private networks using cell relay technology. ATM is an International Telecommunication Union Telecommunication Standardization Sector (ITU-T) standard. Ongoing work on ATM standards is being done primarily by the ATM Forum, which was jointly founded by Cisco Systems, NET/ADAPTIVE, Northern Telecom, and Sprint in 1991.

A cell switching and multiplexing technology, ATM combines the benefits of circuit switching (constant transmission delay, guaranteed capacity) with those of packet switching (flexibility, efficiency for intermittent traffic). To achieve these benefits, ATM uses the following features:

- Fixed-size cells, permitting more efficient switching in hardware than is possible with variable-length packets
- Connection-oriented service, permitting routing of cells through the ATM network over virtual connections, sometimes called virtual circuits, using simple connection identifiers
- Asynchronous multiplexing, permitting efficient use of bandwidth and interleaving of data of varying priority and size

The combination of these features allows ATM to provide different categories of service for different data requirements and to establish a service contract at the time a connection is set up. This means that a virtual connection of a given service category can be guaranteed a certain bandwidth, as well as other traffic parameters, for the life of the connection.


ATM Basics

To understand how ATM can be used, it is important to have a knowledge of how ATM packages and transfers information. The following sections provide brief descriptions of the format of ATM information transfer and the mechanisms on which ATM networking is based.

ATM Cell Basic Format

The basic unit of information used by ATM is a fixed-size cell consisting of 53 octets, or bytes. The first 5 bytes contain header information, such as the connection identifier, while the remaining 48 bytes contain the data, or payload (see Figure 1-1). Because the ATM switch does not have to detect the size of a unit of data, switching can be performed efficiently. The small size of the cell also makes it well suited for the transfer of real-time data, such as voice and video. Such traffic is intolerant of delays resulting from having to wait for large data packets to be loaded and forwarded.

Figure 1-1 ATM Cell Basic Format



ATM Device Types


An ATM network is made up of one or more ATM switches and ATM endpoints. An ATM endpoint (or end system) contains an ATM network interface adapter. Workstations, routers, data service units (DSUs), LAN switches, and video coder-decoders (CODECs) are examples of ATM end systems that can have an ATM interface. Figure 1-2 illustrates several types of ATM end systems—router, LAN switch, workstation, and DSU/CSU, all with ATM network interfaces—connected to an ATM switch through an ATM network to another ATM switch on the other side.



Note

In this document the term ATM switch is used to refer generically to the network device that switches ATM cells; the term ATM switch router is used to refer to the Catalyst 8540 MSR, Catalyst 8510 MSR, and LightStream 1010 ATM switch.

Figure 1-2 ATM Network Devices



ATM Network Interface Types

There are two types of interfaces that interconnect ATM devices over point-to-point links: the User-Network Interface (UNI) and the Network-Network Interface (NNI), sometimes called Network-Node Interface. A UNI link connects an ATM end-system (the user side) with an ATM switch (the network side). An NNI link connects two ATM switches; in this case, both sides are network.

UNI and NNI are further subdivided into public and private UNIs and NNIs, depending upon the location and ownership of the ATM switch. As shown in Figure 1-3, a private UNI connects an ATM endpoint and private ATM switch; a public UNI connects an ATM endpoint or private switch to a public switch. A private NNI connects two ATM switches within the same private network; a public NNI connects two ATM switches within the same public network. A third type of interface, the Broadband Inter-Carrier Interface (BICI) connects two public switches from different public networks.

Your ATM switch router supports interface types UNI and NNI, including the PNNI routing protocol. For examples of UNI and NNI, see Chapter 3, "ATM Network Interfaces."

Figure 1-3 ATM Network Interfaces

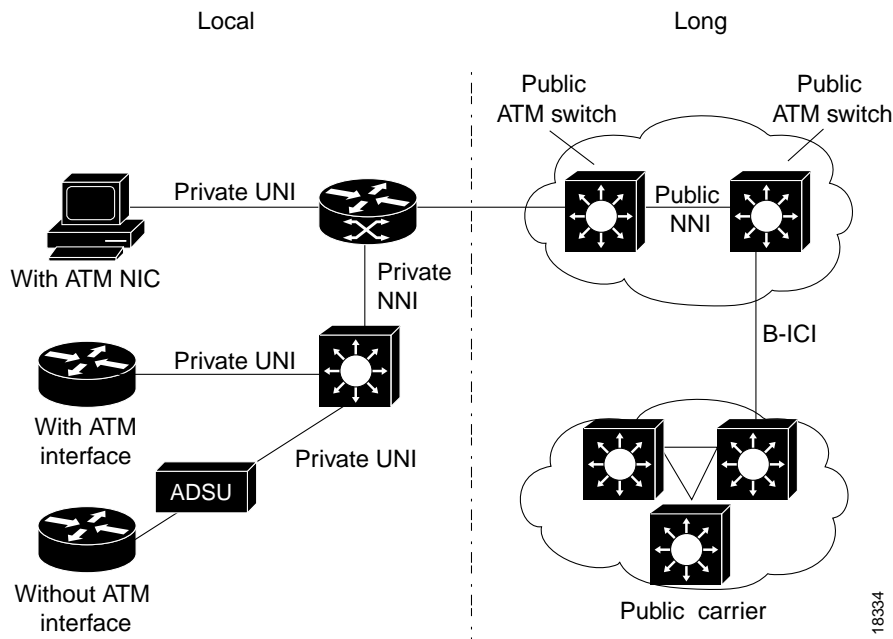


Figure 1-3 also illustrates some further examples of ATM end systems that can be connected to ATM switches. A router with an ATM interface processor (AIP) can be connected directly to the ATM switch, while the router without the ATM interface must connect to an ATM data service unit (ADSU) and from there to the ATM switch.

ATM Cell Header Formats

The ATM cell includes a 5-byte header. Depending upon the interface, this header can be in either UNI or NNI format. The UNI cell header, as depicted in Figure 1-4, has the following fields:

- Generic flow control (GFC)—provides local functions, such as flow control from endpoint equipment to the ATM switch. This field is presently not used.
- Virtual path identifier (VPI) and virtual channel identifier (VCI)—VPI identifies a virtual path leg on an ATM interface. VPI and VCI together identify a virtual channel leg on an ATM interface. Concatenating such legs through switches forms a virtual connection across a network.
- Payload type (PT)—indicates in the first bit whether the cell contains user data or control data. If the cell contains user data, the second bit indicates whether congestion is experienced or not, and the third bit indicates whether the cell is the last in a series of cells that represent a single AAL5 frame. (AAL5 is described in the “Service-dependent ATM Adaptation Layers” section on page 1-14.) If the cell contains control data, the second and third bits indicate maintenance or management flow information.
- Cell loss priority (CLP)—indicates whether the cell should be discarded if it encounters extreme congestion as it moves through the network.
- Header error control (HEC)—contains a cyclic redundancy check on the cell header.

Figure 1-4 ATM Cell Header—UNI Format

GFC 4	VPI 8	VCI 16	PT 3	CLP 1	HEC 8
32 bits					8 bits CRC

18335

The NNI cell header format, depicted in Figure 1-5, includes the same fields except that the GFC space is displaced by a larger VPI space, occupying 12 bits and making more VPIs available for NNIs.

Figure 1-5 ATM Cell Header—NNI Format

VPI 12	VCI 16	PT 3	CLP 1	HEC 8
32 bits				8 bits CRC

18336

ATM Services

There are three general types of ATM services:

- Permanent virtual connection (PVC) service—connection between points is direct and permanent. In this way, a PVC is similar to a leased line.
- Switched virtual connection (SVC) service—connection is created and released dynamically. Because the connection stays up only as long as it is in use (data is being transferred), an SVC is similar to a telephone call.
- Connectionless service—similar to Switched Multimegabit Data Service (SMDS)



Note

Your ATM switch router supports permanent and switched virtual connection services. It does not support connectionless service.

Advantages of PVCs are the guaranteed availability of a connection and that no call setup procedures are required between switches. Disadvantages include static connectivity and that they require manual administration to set up.

Advantages of SVCs include connection flexibility and call setup that can be automatically handled by a networking device. Disadvantages include the extra time and overhead required to set up the connection.


Virtual Paths and Virtual Channels

ATM networks are fundamentally connection oriented. This means that a virtual connection needs to be established across the ATM network prior to any data transfer. ATM virtual connections are of two general types:

- Virtual path connections (VPCs), identified by a VPI.
- Virtual channel connections (VCCs), identified by the combination of a VPI and a VCI.

A virtual path is a bundle of virtual channels, all of which are switched transparently across the ATM network on the basis of the common VPI. A VPC can be thought of as a bundle of VCCs with the same VPI value (see Figure 1-6).

Figure 1-6 ATM Virtual Path And Virtual Channel Connections



Every cell header contains a VPI field and a VCI field, which explicitly associate a cell with a given virtual channel on a physical link. It is important to remember the following attributes of VPIs and VCIs:

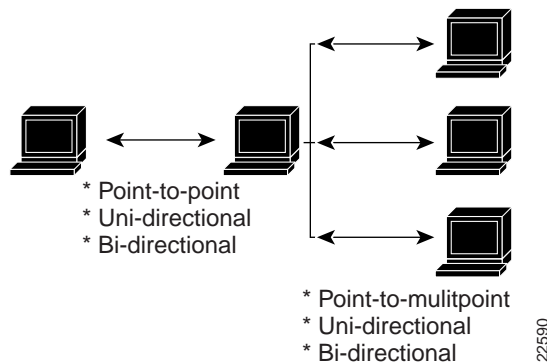
- VPIs and VCIs are not addresses, such as MAC addresses used in LAN switching.
- VPIs and VCIs are explicitly assigned at each segment of a connection and, as such, have only local significance across a particular link. They are remapped, as appropriate, at each switching point.

Using the VCI/VPI identifier, the ATM layer can multiplex (interleave), demultiplex, and switch cells from multiple connections.

Point-to-Point and Point-to-Multipoint Connections

Point-to-point connections connect two ATM systems and can be unidirectional or bidirectional. By contrast, point-to-multipoint connections (see Figure 1-7) join a single source end system (known as the root node) to multiple destination end-systems (known as leaves). Such connections can be unidirectional only, in which only the root transmits to the leaves, or bidirectional, in which both root and leaves can transmit.

Figure 1-7 Point-to-Point and Point-to-Multipoint Connections



Note that there is no mechanism here analogous to the multicasting or broadcasting capability common in many shared medium LAN technologies, such as Ethernet or Token Ring. In such technologies, multicasting allows multiple end systems to both receive data from other multiple systems, and to transmit data to these multiple systems. Such capabilities are easy to implement in shared media technologies such as LANs, where all nodes on a single LAN segment must necessarily process all packets sent on that segment. The obvious analog in ATM to a multicast LAN group would be a bidirectional multipoint-to-multipoint connection. Unfortunately, this obvious solution cannot be implemented when using AAL5, the most common ATM Adaptation Layer (AAL) used to transmit data across ATM networks.

AAL 5 does not have any provision within its cell format for the interleaving of cells from different AAL5 packets on a single connection. This means that all AAL5 packets sent to a particular destination across a particular connection must be received in sequence, with no interleaving between the cells of different packets on the same connection, or the destination reassembly process would not be able to reconstruct the packets.

This is why ATM AAL 5 point-to-multipoint connections can only be unidirectional; if a leaf node were to transmit an AAL 5 packet onto the connection, it would be received by both the root node and all other leaf nodes. However, at these nodes, the packet sent by the leaf could well be interleaved with packets sent by the root, and possibly other leaf nodes; this would preclude the reassembly of any of the interleaved packets.


Solutions

For ATM to interoperate with LAN technology, it needs some form of multicast capability. Among the methods that have been proposed or tried, two approaches are considered feasible (see Figure 1-8).

- **Multicast server.** In this mechanism, all nodes wishing to transmit onto a multicast group set up a point-to-point connection with an external device known as a multicast server. The multicast server, in turn, is connected to all nodes wishing to receive the multicast packets through a point-to-multipoint connection. The multicast server receives packets across the point-to-point connections, serializes them (that is, ensures that one packet is fully transmitted prior to the next being sent), and retransmits them across the point-to-multipoint connection. In this way, cell interleaving is precluded.
- **Overlaid point-to-multipoint connections.** In this mechanism, all nodes in the multicast group establish a point-to-multipoint connection with each other node in the group and, in turn, become a leaf in the equivalent connections of all other nodes. Hence, all nodes can both transmit to and receive from all other nodes. This solution requires each node to maintain a connection for each transmitting member of the group, while the multicast server mechanism requires only two connections. The overlaid connection model also requires a registration process for telling nodes that join a group what the other nodes in the group are, so that it can form its own point-to-multipoint connection. The other nodes also need to know about the new node so they can add the new node to their own point-to-multipoint connections.

Of these two solutions, the multicast server mechanism is more scalable in terms of connection resources, but has the problem of requiring a centralized resequencer, which is both a potential bottleneck and a single point of failure.

Figure 1-8 Approaches to ATM Multicasting



Applications

Two applications that require some mechanism for point-to-multipoint connections are:

- LAN emulation—in this application, the broadcast and unknown server (BUS) provides the functionality to emulate LAN broadcasts. See Chapter 6, “LAN Emulation and MPOA,” for a discussion of this protocol.
- Video broadcast—in this application, typically over a CBR connection, a video server needs to simultaneously broadcast to any number of end stations. See Chapter 9, “Circuit Emulation Services and Voice over ATM.”

Operation of an ATM Switch

An ATM switch has a straightforward job:

1. Determine whether an incoming cell is eligible to be admitted to the switch (a function of Usage Parameter Control [UPC]), and whether it can be queued.
2. Possibly perform a replication step for point-to-multipoint connections.
3. Schedule the cell for transmission on a destination interface. By the time it is transmitted, a number of modifications might be made to the cell, including the following:
 - VPI/VCI translation
 - setting the Early Forward Congestion Indicator (EFCI) bit
 - setting the CLP bit


The functions of UPC, EFCI, and CLP are discussed in Chapter 10, “Traffic and Resource Management.”

Because the two types of ATM virtual connections differ in how they are identified, as described in the “Virtual Paths and Virtual Channels” section on page 1-5, they also differ in how they are switched. ATM switches therefore fall into two categories—those that do virtual path switching only and those that do switching based on virtual path and virtual channel values.

The basic operation of an ATM switch is the same for both types of switches: Based on the incoming cell’s VPI or VPI/VCI pair, the switch must identify which output port to forward a cell received on a given input port. It must also determine the new VPI/VCI values on the outgoing link, substituting these new values in the cell before forwarding it. The ATM switch derives these values from its internal tables, which are set up either manually for PVCs, or through signaling for SVCs.

Figure 1-9 shows an example of virtual path (VP) switching, in which cells are switched based only on the value of the VPI; the VCI values do not change between the ingress and the egress of the connection. This is analogous to central office trunk switching.

Figure 1-9 Virtual Path Switching

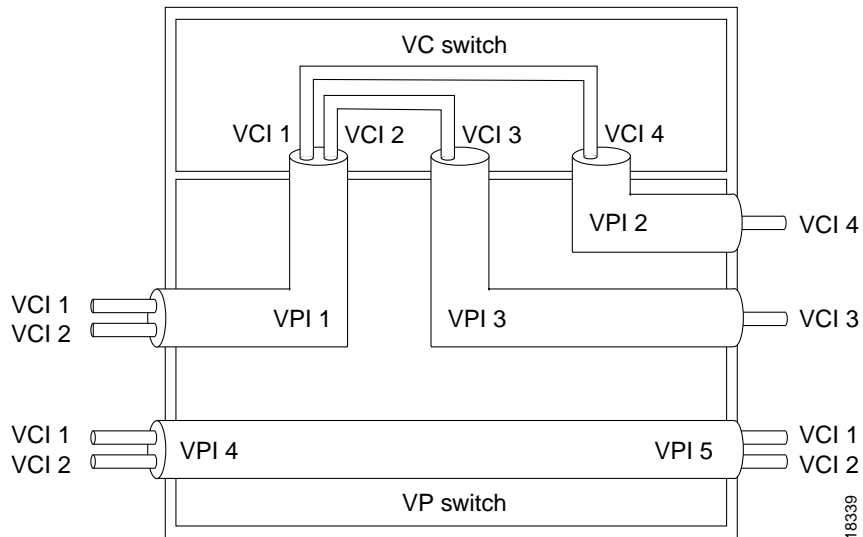


VP switching is often used when transporting traffic across the WAN. VPCs, consisting of aggregated VCCs with the same VPI number, pass through ATM switches that do VP switching. This type of switching can be used to extend a private ATM network across the WAN by making it possible to

support signaling, PNNI, LANE, and other protocols inside the virtual path, even though the WAN ATM network might not support these features. VPCs terminate on VP tunnels, as described in the “VP Tunnels” section on page 4-13 in the chapter “Virtual Connections.”

Figure 1-10 shows an example of switching based on both VPI and VCI values. Because all VCIs and VPIs have only local significance across a particular link, these values get remapped, as necessary, at each switch. Within a private ATM network switching is typically based on both VPI and VCI values.

Figure 1-10 Virtual Path/Virtual Channel Switching



Note

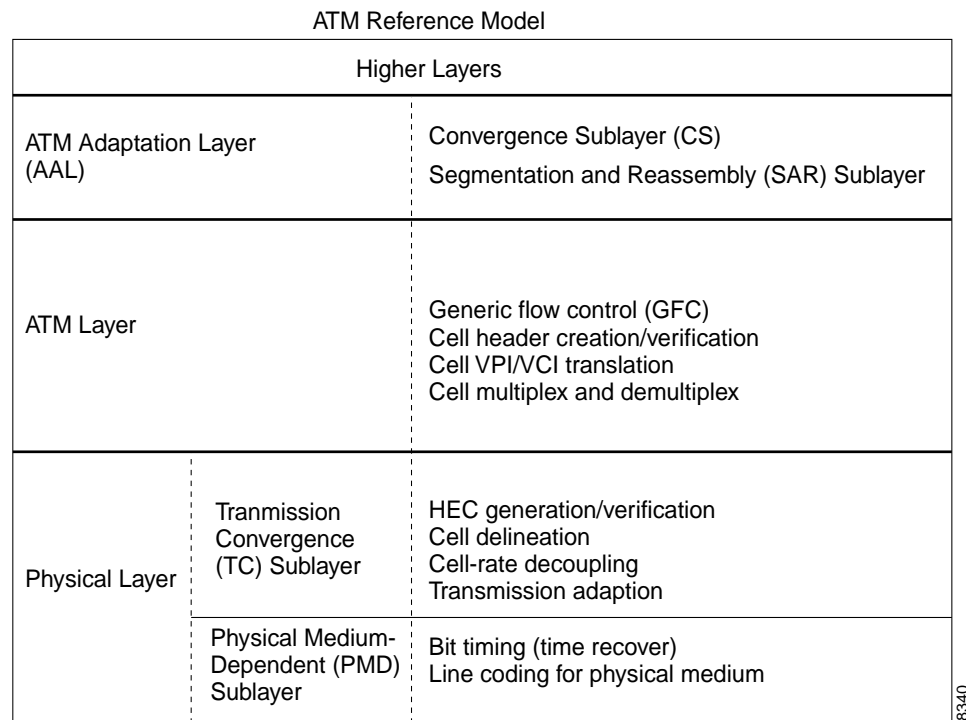
Your Cisco ATM switch router performs both virtual path and virtual channel switching.

18339

The ATM Reference Model

The ATM architecture is based on a logical model, called the ATM reference model, that describes the functionality it supports. In the ATM reference model (see Figure 1-11), the ATM physical layer corresponds approximately to the physical layer of the OSI reference model, and the ATM layer and ATM adaptation layer (AAL) are roughly analogous to the data link layer of the OSI reference model.

Figure 1-11 ATM Reference Model



18340

The layers of the ATM reference model have the following functions:

- Physical layer—manages the medium-dependent transmission. The physical layer is divided into two sublayers:
 - Physical medium-dependent sublayer—synchronizes transmission and reception by sending and receiving a continuous flow of bits with associated timing information, and specifies format used by the physical medium.
 - Transmission convergence (TC) sublayer—maintains ATM cell boundaries (cell delineation), generates and checks the header error-control code (HEC), maintains synchronization and inserts or suppresses idle ATM cells to provide a continuous flow of cells (cell-rate decoupling), and packages ATM cells into frames acceptable to the particular physical layer-implementation (transmission-frame adaptation).
- ATM layer—establishes connections and passes cells through the ATM network. The specific tasks of the ATM layer include the following:
 - Multiplexes and demultiplexes cells of different connections
 - Translates VPI/VCI values at the switches and cross connections

- Extracts and inserts the header before or after the cell is delivered to the AAL
- Maintains flow control using the GFC bits of the header
- ATM adaptation layer (AAL)—isolates higher-layer protocols from the details of the ATM processes by converting higher-layer information into ATM cells and vice versa. The AAL is divided into two sublayers:
 - Convergence sublayer (CS)—takes the common part convergence sublayer (CPCS) frame, divides it into 53-byte cells, and sends these cells to the destination for reassembly.
 - Segmentation and reassembly sublayer—segments data frames into ATM cells at the transmitter and reassembles them into their original format at the receiver.
- Higher layers—accept user data, arrange it into packets, and hand it to the AAL.

ATM Addressing

If cells are switched through an ATM network based on the VPI/VCI in the cell header, and not based directly on an address, why are addresses needed at all? For permanent, statically configured virtual connections there is in fact no need for addresses. But SVCs, which are set up through signaling, do require address information.

SVCs work much like a telephone call. When you place a telephone call you must have the address (telephone number) of the called party. The calling party signals the called party's address and requests a connection. This is what happens with ATM SVCs; they are set up using signaling and therefore require address information.

The types and formats of ATM addresses, along with their uses, are described in Chapter 2, "ATM Signaling and Addressing."

Traffic Contracts and Service Categories

ATM connections are further characterized by a traffic contract, which specifies a service category along with traffic and quality of service (QoS) parameters. Five service categories are currently defined, each with a purpose and its own interpretation of applicable parameters.

The following sections describe the components of the traffic contract, the characteristics of the service categories, and the service-dependent AAL that supports each of the service categories.

The Traffic Contract

At the time a connection is set up, a traffic contract is entered, guaranteeing that the requested service requirements will be met. These requirements are traffic parameters and QoS parameters:

- Traffic parameters—generally pertain to bandwidth requirements and include the following:
 - Peak cell rate (PCR)
 - Sustainable cell rate (SCR)
 - Burst tolerance, conveyed through the maximum burst size (MBS)
 - Cell delay variation tolerance (CDVT)
 - Minimum cell rate (MCR)
- QoS parameters—generally pertain to cell delay and loss requirements and include the following:
 - Maximum cell transfer delay (MCTD)
 - Cell loss ratio (CLR)
 - Peak-to-peak cell delay variation (ppCDV)

The Service Categories

One of the main benefits of ATM is to provide distinct classes of service for the varying bandwidth, loss, and latency requirements of different applications. Some applications require constant bandwidth, while others can adapt to the available bandwidth, perhaps with some loss of quality. Still others can make use of whatever bandwidth is available and use dramatically different amounts from one instant to the next.

ATM provides five standard service categories that meet these requirements by defining individual performance characteristics, ranging from best effort (Unspecified Bit Rate [UBR]) to highly controlled, full-time bandwidth (Constant Bit Rate [CBR]). Table 1-1 lists each service category defined by the ATM Forum along with its applicable traffic parameters and QoS characteristics.

Table 1-1 Service Categories and Characteristics

Service Category	Traffic Parameters	QoS Characteristics	
		Cell Loss	Cell Delay
CBR—constant bit rate	PCR	low	low
VBR-RT—variable bit rate real-time	PCR, SCR, MBS	low	low
VBR-NRT—variable bit rate non-real time	PCR, SCR, MBS	low	unspecified
ABR—available bit rate	PCR, MCR	unspecified	unspecified
UBR—unspecified bit rate	(no guarantees)	unspecified	unspecified

The characteristics and uses of each service category are summarized as follows:

- CBR service provides constant bandwidth with a fixed timing relationship, which requires clocking synchronization. Because CBR traffic reserves a fixed amount of bandwidth, some trunk bandwidth might go unused. CBR is typically used for circuit emulation services to carry real-time voice and video.
- VBR-RT service provides only a partial bandwidth guarantee. Like CBR, however, some bandwidth might still go unused. Typical applications include packetized voice and video, and interactive multimedia.
- VBR-NRT service provides a partial bandwidth guarantee, but with a higher cell delay than VBR-RT. This service category is suitable for bursty applications, such as file transfers.
- ABR provides a best effort service, in which feedback flow control within the network is used to increase bandwidth when no congestion is present, maximizing the use of the network.
- UBR service provides no bandwidth guarantee, but attempts to fill bandwidth gaps with bursty data. UBR is well suited for LAN protocols, such as LAN emulation. An additional category, UBR+, is a Cisco extension to UBR that provides for a nonzero MCR in the traffic contract.

Service-dependent ATM Adaptation Layers

For ATM to support multiple classes of service with different traffic characteristics and requirements, it is necessary to adapt the different classes to the ATM layer. This adaptation is performed by the service-dependent AAL.

The service-dependent AAL provides a set of rules for segmentation and reassembly of packets. The sender segments the packet and builds a set of cells for transmission, while the receiver verifies the integrity of the packet and reassembles the cells back into packets—all according to a set of rules designed to satisfy a particular type of service. Table 1-2 lists the four AAL types recommended by the ITU-T, along with the service categories commonly supported by each and the corresponding connection mode.



Note

The correspondence between AAL and service category is not a fixed one. For example, AAL5 can be used for CBR.

Table 1-2 Service-Dependent ATM Adaptation Layers and Service Categories

AAL	Service Category	Connection Mode and Characteristics
AAL1	CBR	Connection-oriented; supports delay-sensitive services that require constant bit rates and have specified timing and delay requirements, such as uncompressed video.
AAL2	VBR	Connection-oriented; supports services that do not require constant bit rates, such as video schemes that use variable bit rate applications. AAL2 is presently an incomplete standard.
AAL3/4	UBR	Connectionless; mainly used for SMDS applications.
AAL5	ABR, UBR, VBR	Connection-oriented and connectionless; supports services with varying bit rate demands; offers low bandwidth overhead and simpler processing requirements in exchange for reduced bandwidth capacity and error-recovery capability.

Common Physical Interface Types

ATM networks can use many different kinds of physical interfaces. The ATM Forum has defined a number of these interface types and is working on defining still others. In general, an interface type is defined by three characteristics:

- Data rate—the overall bandwidth, in Mbps, for a physical interface. Data rates for standard ATM physical interfaces range from 1.544 to 2488.32 Mbps.
- Physical medium—the physical characteristic of the link, which determines the type of signal it can carry. Physical media fall into two categories:
 - Optical, including multimode fiber and single-mode fiber
 - Electrical, including coaxial cable, unshielded twisted-pair (UTP), and foil twisted-pair (FTP, formerly shielded twisted-pair [STP])
- Framing type—how the ATM cells are framed to be carried over the physical medium. Framing types include the following:
 - ATM25, also called Desktop25—used for 25.6-Mbps connections over UTP-3, primarily for desktop connections
 - Transparent Asynchronous Transmitter/Receiver Interface 4B/5B (TAXI)—used for speeds of up to 100 Mbps over multimode fiber
 - Digital signal level 1 (DS-1)—used for 1.544-Mbps T1 and 2.108-Mbps E1 facilities
 - Digital signal level 2 (DS-3)—used for 44.736-Mbps T3 and 34.368-Mbps E3 facilities
 - Synchronous Optical Network (SONET)—used for high-speed transmission over optical or electrical media

Optical media SONET rates are designated OC- x ; electrical media rates are designated Synchronous Transport Signal (STS- x), where x designates a data rate. A near-equivalent standard, Synchronous Digital Hierarchy (SDH), specifies framing only for electrical signals. SDH rates are designated Synchronous Transport Module (STM- x).

Table 1-3 shows the most commonly used physical interface types for ATM.

Table 1-3 Common ATM Physical Interface Types

Framing/Interface Type	Data Rate (Mbps)	Physical Media
DS-1		
T1	1.544	twisted pair
E1	2.048	twisted pair and coaxial cable
DS-3		
T3	44.736	coaxial cable
E3	34.368	coaxial cable
ATM25	25.6	UTP-3
4B/5B (TAXI)	100	multimode fiber
SONET/SDH		
OC-3	155.52	multimode and single-mode fiber
STS-3c/STM-1	155.52	UTP-5
OC-12	622.08	single-mode fiber
OC-48	2488.32	single-mode fiber

A physical interface on an ATM switch must support all three characteristics—framing type, data rate, and physical medium. As Table 1-3 shows, an OC-3 interface—the most commonly used one for ATM—can run over multimode or single-mode fiber. If you planned to use an OC-3 SM fiber link, you would need a physical interface (port adapter or interface module) that supports the SONET framing at 155.52 Mbps over single-mode fiber.

The choice of physical interface depends upon a number of variables, including bandwidth requirements and link distance. In general, UTP is used for applications to the desktop, multimode fiber between wiring closets or buildings, and SM fiber across long distances.



Note

This guide does not discuss hardware. Refer to the *ATM Switch Router Software Configuration Guide* and to your hardware documentation for the characteristics and features of the port adapters and interface modules supported on your particular ATM switch router model.