

Designing ATM Internetworks

Asynchronous Transfer Mode (ATM) is an evolving technology designed for the high-speed transfer of voice, video, and data through public and private networks in a cost-effective manner. ATM is based on the efforts of Study Group XVIII of the International Telecommunication Union Telecommunication Standardization Sector (ITU-T, formerly the Consultative Committee for International Telegraph and Telephone [CCITT]) and the American National Standards Institute (ANSI) to apply very large-scale integration (VLSI) technology to the transfer of data within public networks. Officially, the ATM layer of the Broadband Integrated Services Digital Network (BISDN) model is defined by CCITT I.361. Current efforts to bring ATM technology to private networks and to guarantee interoperability between private and public networks is being done by the ATM Forum, which was jointly founded by Cisco Systems, NET/ADAPTIVE, Northern Telecom, and Sprint in 1991.

This chapter describes current ATM technologies that network designers can use in their networks today. It also makes recommendations for designing non-ATM networks today that those networks can take advantage of ATM in the future without sacrificing current investments in cable.

This chapter focuses on the following topics:

- ATM Overview
- LAN Emulation
- ATM Data Exchange Interface
- ATM Interface Processor Card
- Cisco LightStream 100
- ATM Media

ATM Overview

ATM combines the strengths of time-division multiplexing (TDM)—whose fixed time slots are used by telephone companies to deliver voice without distortion—with the strengths of packet-switching data networks—whose variable size data units are used by computer networks, such as the Internet, to deliver data efficiently. While building on the strengths of TDM, ATM avoids the weaknesses of TDM (which wastes bandwidth by transmitting the fixed time slots even when no one is speaking) and PSDNs (which cannot accommodate time-sensitive traffic, such as voice and video, because PSDNs are designed for transmitting bursty data). By using fixed-size cells, ATM combines the isochronicity of TDM with the efficiency of PSDN.

ATM Cell Format

ATM cells consist of 5 bytes of header information and 48 bytes of payload data. Two fields in the ATM header are used to route cells through ATM networks:

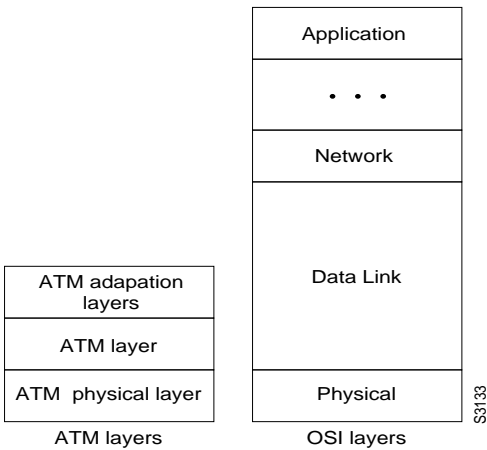
- VPI—The virtual path identifier.
- VCI—The virtual channel identifier.

The VPI and VCI fields of the cell header identify the next network segment that a cell needs to transit on its way to its final destination. A *virtual channel* is equivalent to a virtual circuit—that is, both terms describe a logical connection between the two ends of a communications connection. A *virtual path* is a logical grouping of virtual circuits that allows an ATM switch to perform operations on groups of virtual circuits.

ATM Functional Layers

ATM consists of three functional layers—the ATM physical layer, the ATM layer, and the ATM adaptation layer—that correspond roughly to layer 1 and parts of layer 2 (such as error control and data framing) of the *Open System Interconnection* (OSI) reference model, as shown in Figure 5-1.

Figure 5-1 Relationship of ATM Functional Layers to the OSI Reference Model



The ATM physical layer controls transmission and receipt of bits on the physical medium. It also keeps track of ATM cell boundaries and packages cells into the appropriate type of frame for the physical medium being used. Above the physical layer is the *ATM layer*, which is responsible for establishing virtual connections and passing ATM cells through the ATM network. Immediately above the ATM layer is the *ATM adaptation layer*, which translates between the larger *service data units* (SDUs) (for example, video streams, and data packets) of upper-layer processes and ATM cells. Several ATM adaptation layers are currently specified. They are analyzed later in this chapter. Cisco routers currently support ATM adaptation layer 5 (AAL5), which is used to transfer most non-SMDS data, such as classical IP over ATM and local-area network (LAN) emulation, and AAL3/4 for SMDS traffic.

ATM Addressing

The ATM Forum has adapted the subnetwork model of addressing in which the ATM layer is responsible for mapping network layer addresses to ATM addresses. Several ATM address formats have been developed. Public ATM networks will typically use E.164 numbers, as also used by Narrowband ISDN (N-ISDN) networks.

The ATM address formats are modeled on ISO Network Service Access Point (NSAP) addresses, but they identify SubNetwork Point of Attachment (SNPA) addresses. Incorporating the MAC address into the ATM address makes it easy to map ATM addresses into existing LANs.

Note For overview information about ATM, see the *Internetworking Technology Overview* publication.

LAN Emulation

LAN Emulation (LANE) is a service that provides interoperability between ATM-based workstations and devices connected to existing legacy LAN technology. The ATM Forum has defined a standard for LANE that provides to workstations attached via ATM the same capabilities that they are used to obtaining from legacy LANs.

LANE uses MAC encapsulation (OSI Layer 2) because this approach supports the largest number of existing OSI Layer 3 protocols. The end result is that all devices attached to an emulated LAN appear to be on one bridged segment. In this way, AppleTalk, IPX, and other protocols should have similar performance characteristics as in a traditional bridged environment. In ATM LANE environments, the ATM switch handles traffic that belongs to the same emulated LAN (ELAN), and routers handle inter-ELAN traffic.

LANE components include the following:

- *LAN emulation client (LEC)*—End systems that support LANE, such as network interface card (NIC)-connected workstations, Catalyst 5000 switches, and Cisco 7000 series routers, require the implementation of a LEC. The LEC emulates an interface to a legacy LAN to the higher-level protocols. It performs data forwarding, address resolution, and registration of MAC addresses with the LANE server and communicates with other LECs via ATM virtual channel connections (VCCs).
- *LAN emulation server (LES)*—The LES provides a central control point for all LECs. LECs maintain a Control Direct VCC to the LES to forward registration and control information. The LES maintains a point-to-multipoint VCC, known as the *Control Distribute VCC*, to all LECs. The Control Distribute VCC is used only to forward control information. As new LECs join the ATM emulated LAN (ELAN), each LEC is added as a leaf to the control distribute tree.
- *Broadcast and unknown Server (BUS)*—The BUS acts as a central point for distributing broadcasts and multicasts. ATM is essentially a point-to-point technology without “any-to-any” or “broadcast” support. LANE solves this problem by centralizing the broadcast support in the BUS. Each LEC must set up a Multicast Send VCC to the BUS. The BUS then adds the LEC as a leaf to its point-to-multipoint VCC (known as the *Multicast Forward VCC*).

The BUS also acts as a multicast server. LANE is defined on ATM adaptation layer 5 (AAL5), which specifies a simple trailer to be appended to a frame before it is broken into ATM cells. The problem is that there is no way to differentiate between ATM cells from different senders when multiplexed on a virtual channel. It is assumed that cells received will be in sequence, and when

the End of Message (EOM) cell arrives, you should just have to reassemble all of the cells that have already arrived. (For additional information about how AAL5 works, see the discussion of ATM in the *Internetworking Technology Overview* publication.)

The BUS takes the sequence of cells on each Multicast Send VCC and reassembles them into frames. When a full frame is received, it is queued for sending to all of the LECs on the Multicast Forward VCC. This way, all the cells from a particular data frame can be guaranteed to be sent in order and not interleaved with cells from any other data frames on the point-to-multipoint VCC.

- *LAN emulation configuration server (LECS)*—This server maintains a database of LECs and the ELANs that they belong to. It accepts queries from LECs and responds with the appropriate VLAN identifier, namely the ATM address of the LES that serves the appropriate VLAN or ELAN. This database is maintained by the network administrator.

Note that because LANE is defined at OSI Layer 2, the LECS is the only security checkpoint available. Once it has been told where to find the LES and it has successfully joined the ELAN, the LEC is free to send any traffic (whether malicious or not) into the bridged ELAN. The only place for any OSI Layer 3 security filters is in the router that routes this ELAN to other ELANs. Therefore, the larger the ELAN, the greater the exposure to security violations.

How LAN Emulation Works

In a typical LANE operation, the LEC must first find the LECS to discover which ELAN it should join. Specifically, the LEC is looking for the ATM address of the LECS that serves the desired ELAN. To find the ATM address of the LECS, the LEC does the following:

- 1 Queries the ATM switch via Interim Local Management Interface (ILMI). The switch has a MIB variable set up with the ATM address of the LECS. The LEC can then use UNI signaling to contact the LECS.
- 2 Looks for a fixed ATM address that is specified by the ATM Forum as the LECS ATM address.
- 3 Accesses PVC 0/17, a “well-known” PVC.

The LEC creates a signaling packet with the ATM address of the LECS. It signals a Configure Direct VCC and then issues an LE_CONFIGURE_REQUEST on that VCC. The information in this request is compared with the data in the LECS database. The source ATM address is most commonly used to place a LEC into a specific ELAN. There may also be a default LES for those LECs that are not found in the database.

If a matching entry is found, a successful LE_CONFIGURE_RESPONSE is returned with the ATM address of the LES that serves the desired ELAN.

In Cisco’s first implementation of LANE, the LECS resides on a subinterface of an ATM Interface Processor (AIP). It is likely that this function will move to the network management station at some time in the future because that is where all of the VLAN membership information is likely to reside.

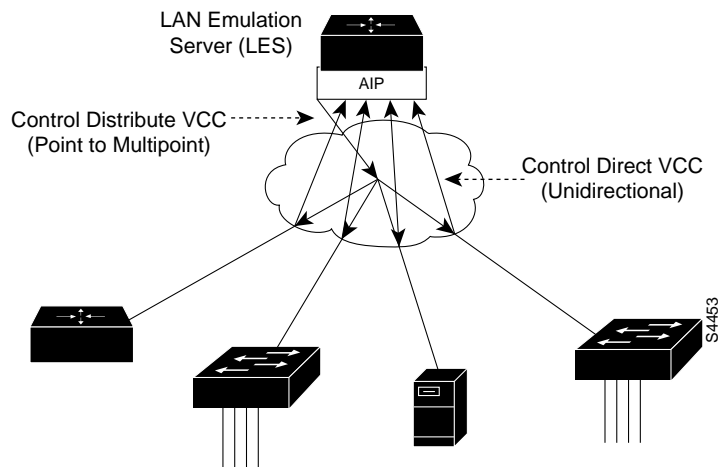
Joining the LES

Once the LEC has discovered the ATM address of the desired LES, it drops the connection to the LECS, creates a signaling packet with the ATM address of the LES, and signals a Control Direct VCC. Upon successful VCC setup, the LES sends an LE_JOIN_REQUEST.

This request contains the LEC ATM address as well as a MAC address that the LEC wants to register with the ELAN. This information is maintained so that no two LECs can register the same MAC or ATM addresses.

Upon receipt of the LE_JOIN_REQUEST, the LES checks with the LECS via its own open connection with the LECS and verifies the request, thus confirming the client's membership. Upon successful verification, the LES adds the LEC as a leaf of its point-to-multipoint Control Distribute VCC. Finally, the LES issues the LEC a successful LE_JOIN_RESPONSE that contains a LANE client ID (LECID), which is an identifier that is unique to the new client. This ID is used by the LEC to filter its own broadcasts from the BUS.

Figure 5-2 LES Connections

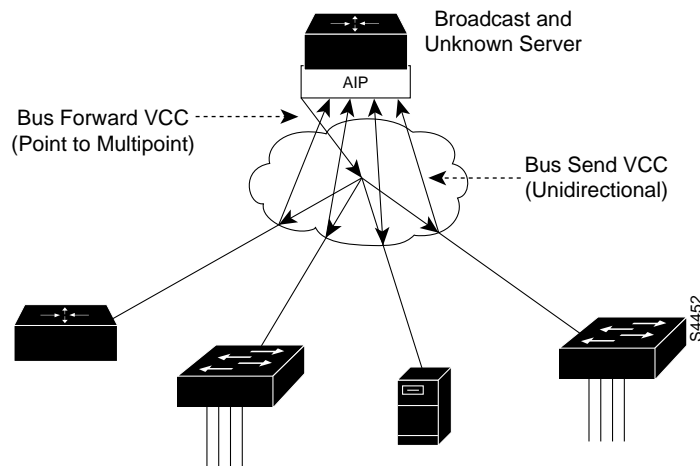


Finding the BUS

After the LEC has successfully joined the LES, its first task is to find the ATM address of the BUS and join the broadcast group. The LEC creates an LE_ARP_REQUEST packet with the MAC address 0xFFFFFFFF. This special LE_ARP packet is sent on the Control Direct VCC to the LES. The LES recognizes that the LEC is looking for the BUS, responds with the ATM address of the BUS, and forwards that response on the Control Distribute VCC.

Joining the BUS

When the LEC has the ATM address of the BUS, its next action is to create a signaling packet with that address and signal a Multicast Send VCC. Upon receipt of the signaling request, the BUS adds the LEC as a leaf on its point-to-multipoint Multicast Forward VCC. At this time, the LEC has become a member of the emulated LAN.

Figure 5-3 BUS Connections

Address Resolution

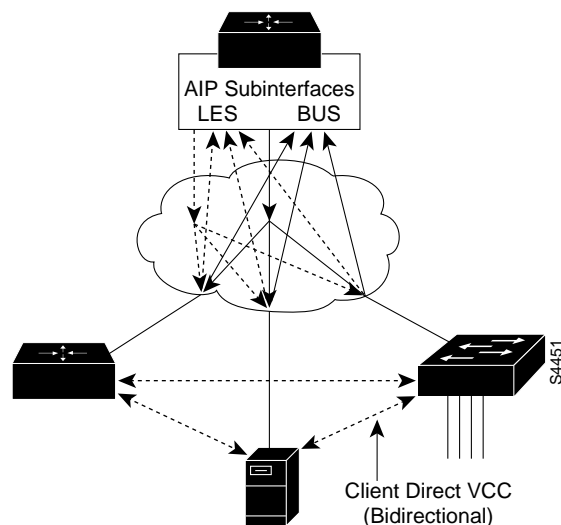
The real value of LANE is the ATM forwarding path that it provides for unicast traffic between LECs. When a LEC has a data packet to send to an unknown destination, it issues an `LE_ARP_REQUEST` to the LES on the Control Direct VCC. The LES forwards the request on the Control Distribute VCC, so all LEC stations hear it. In parallel, the unicast data packets are sent to the BUS, to be forwarded to all endpoints. This “flooding” is not the optimal path for unicast traffic, and this transmission path is rate-controlled to 10 packets per second (per the LANE standard). Unicast packets continue using the BUS until the `LE_ARP_REQUEST` has been resolved.

If bridging or switching devices with LEC software participate in the ELAN, they translate and forward the ARP on their LAN interfaces. One of the LECs should issue an `LE_ARP_RESPONSE` and send it to the LES, which forwards it to the Control Distribute VCC so that all LECs can learn the new MAC-to-ATM address binding.

When the requesting LEC receives the `LE_ARP_RESPONSE`, it has the ATM address of the LEC that represents the MAC address being sought. The LEC should now signal the other LEC directly and set up a Data Direct VCC that will be used for unicast data between the LECs.

While waiting for `LE_ARP` resolution, the LEC forwards unicasts to the BUS. With `LE_ARP` resolution, a new “optimal” path becomes available. If the LEC switches immediately to the new path, it runs the risk of packets arriving out of order. To guard prevent this situation, the LANE standard provides a flush packet.

When the Data Direct VCC becomes available, the LEC generates a flush packet and sends it to the BUS. When the LEC receives its own flush packet on the Multicast Forward VCC, it knows that all previously sent unicasts must have already been forwarded. It is now safe to begin using the Data Direct VCC.

Figure 5-4 Fully Connected Emulated LAN

LAN Emulation Implementation

Cisco implements LECS, LEC, LES, and BUS functionality in the AIP in Cisco IOS Software Release 11.0. These functions are defined on ATM subinterfaces, where one physical interface, such as an OC-3 fiber, is logically divided into up to 255 logical interfaces.

The following are LANE design considerations:

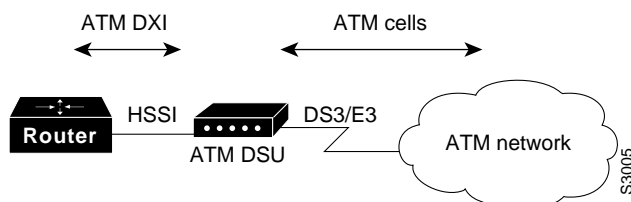
- The AIP is limited to 60K packets per second (pps) bidirectionally.
- One LECS supports all ELANs.
- In each ELAN, there is one LES/BUS pair and some number of LECs.
- The LES and BUS functionality must be defined on the same subinterface and cannot be separated.
- There can be only one LES/BUS pair per subinterface.
- There can be only one LES/BUS pair per ELAN.
- The current LANE Phase 1 standard does not provide for any LES/BUS redundancy.
- The LECS and LES/BUS can be different routers, bridges, or workstations.
- VCCs can be either switched virtual circuits (SVCs) or permanent virtual circuits (PVCs), although PVC design configuration and complexity would probably make anything more than a very small network prohibitively unmanageable and complex.
- When defining VLANs with the Catalyst 5000 switch, each VLAN color should be assigned to a different ELAN. The LES/BUS pair for each ELAN can reside on any of the following:
 - Different subinterfaces on the same AIP
 - Different AIPs in the same router
 - Different AIPs in different routers
- There can be only one LEC per subinterface. If a LEC and a LES/BUS pair share a subinterface, they are (by definition) in the same ELAN.

- If a LEC on a router subinterface is assigned an IP, IPX or AppleTalk address, that protocol is routable over that LEC. If there are multiple LECs present on a router and they are assigned protocol addresses, routing will occur between the ELANs. For routing between ELANs to function correctly, an ELAN should be in only one subnet for a particular protocol.

ATM Data Exchange Interface

To make ATM functionality available as soon as possible, the ATM Forum developed a standard known as the *ATM Data Exchange Interface* (DXI). Network designers can use DXI to provide User-Network Interface (UNI) support between Cisco routers and ATM networks, as shown in Figure 5-5.

Figure 5-5 ATM DXI Topology

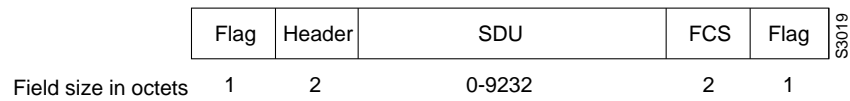


The ATM data service unit (ADSU) receives data from the router in ATM DXI format over a High-Speed Serial Interface (HSSI). The DSU converts the data into ATM cells and transfers them to the ATM network over a DS-3/E3 line.

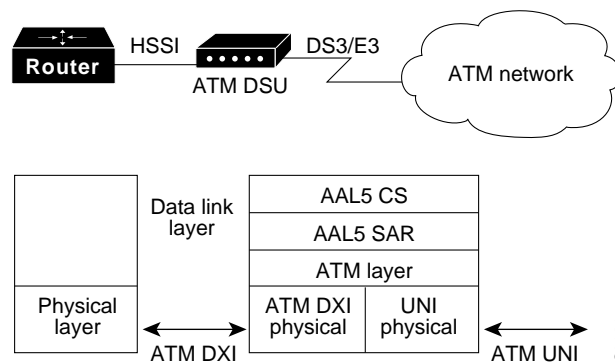
ATM DXI is available in several modes:

- Mode 1a—Supports AAL5 only, a 9232 octet maximum, and a 16-bit FCS, and provides 1023 virtual circuits.
- Mode 1b—Supports AAL3/4 and AAL5, a 9224 octet maximum, and a 16-bit FCS. AAL5 support is the same as Mode 1a. AAL3/4 is supported on one virtual circuit.
- Mode 2—Supports AAL3/4 and AAL5 with 16,777,215 virtual circuits, a 65535 octet maximum, and 32-bit FCS.

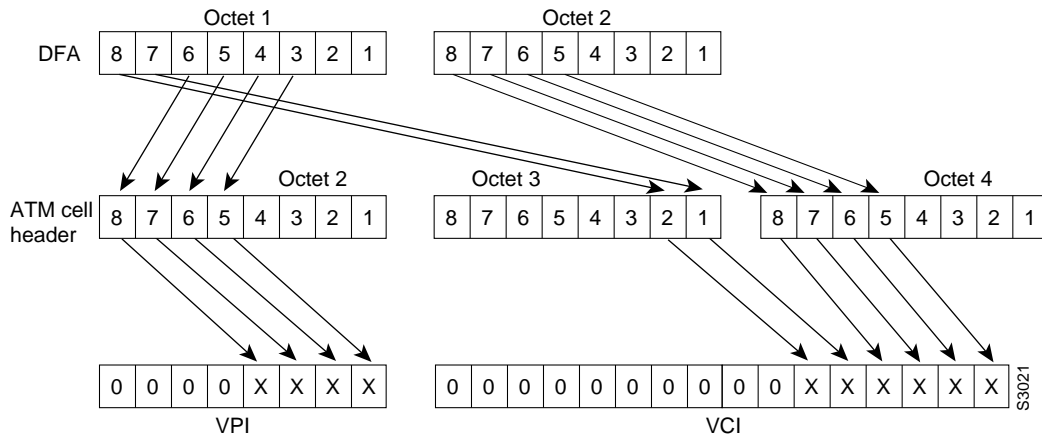
On the router, data from upper-layer protocols is encapsulated into ATM DXI frame format. Figure 5-6 shows the format of a Mode 1a ATM DXI frame.

Figure 5-6 ATM DXI Frame Format

In Figure 5-7, a router configured as a data terminal equipment (DTE) device is connected to an ADSU. The ADSU is configured as a data communications equipment (DCE) device. The router sends ATM DXI frames to the ADSU, which converts the frames to ATM cells by processing them through the AAL5 convergence sublayer (CS) and the segmentation and reassembly sublayer (SAR). The ATM layer attaches the header, and the cells are sent out the ATM UNI interface.

Figure 5-7 ATM DXI Mode 1a and Mode 1b Protocol Architecture for AAL5

ATM DXI addressing consists of a DFA, which is equivalent to a Frame Relay data link connection identifier (DLCI). The DSU maps the DFA into appropriate VPI and VCI values in the ATM cell. Figure 5-8 shows how the DSU performs address mapping.

Figure 5-8 ATM DXI Address Mapping

Note ATM DXI 3.2 is supported in Software Release 9.21 and subsequent software releases. Mode 1a is the only mode supported.

ATM Interface Processor Card

The ATM interface processor (AIP) card supports native ATM in Cisco 7000 and Cisco 7010 routers that are running Cisco IOS Software Release 10.0 or later.

Note Cisco IOS Software Release 10.0 supports AAL5 permanent virtual circuits (PVCs) only.

Cisco IOS Software Release 10.0 supports ATM Forum UNI Specification V3.0, which includes the user-to-network ATM signaling specification. The AIP card uses RFC 1483 (Multiprotocol Encapsulation over AAL5) to transport data through an ATM network. RFC 1483 specifies the use of an LLC/SNAP 8-byte header to identify the encapsulated protocol. It also specifies a null encapsulation (VC Mux) which, instead of headers, creates a separate virtual circuit per protocol.

The AIP supports the following protocols:

- AppleTalk
- Banyan Virtual Network System (VINES)
- Connectionless Network Service (CLNS)
- DECnet
- Internet Protocol (IP)
- Novell Internetwork Packet Exchange (IPX)

The following physical layer interface modules (PLIMs) are available for the AIP:

- TAXI 4B/5B 100-megabits-per-second (Mbps) multimode fiber-optic cable
- SONET/SDH 155-Mbps fiber-optic (STS-3c or STM1) cable
- SONET/SDH 155-Mbps single-mode fiber-optic (STS-3c or STM1) cable
- E3 34-Mbps coaxial cable
- DS-3 45-Mbps cable

The total bandwidth though all the AIPs configured in a router should be limited to 200 Mbps full duplex. For that reason, only the following combinations are supported:

- 2 TAXI interfaces
- 1 SONET and one E3 interface
- 2 SONET interfaces, one of which is lightly used
- 5 E3 interfaces

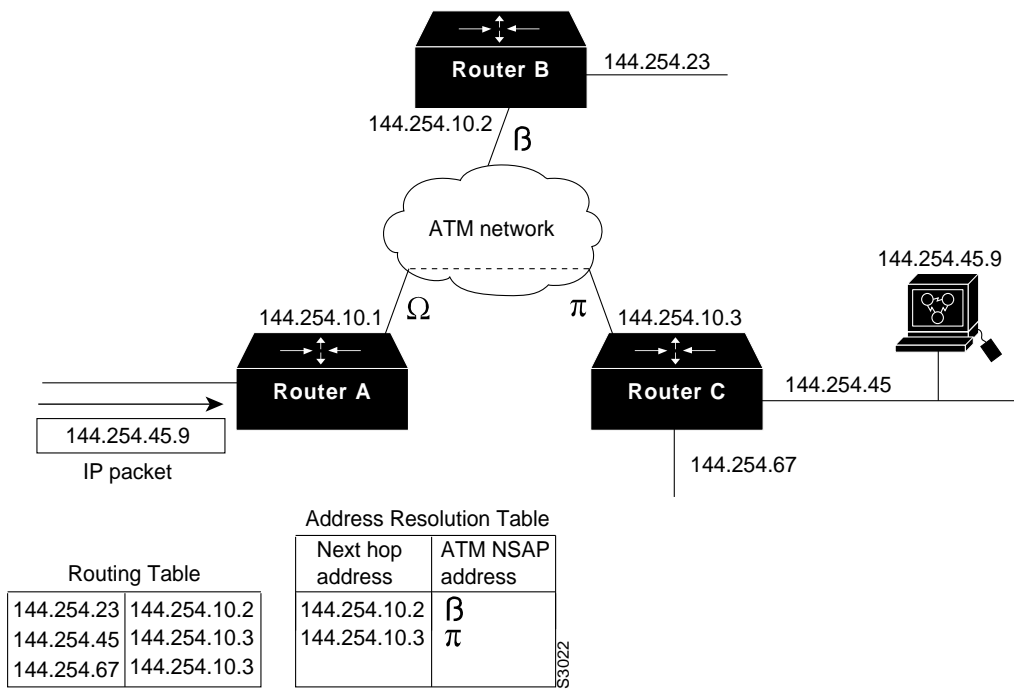
The AIP includes hardware support for various traffic shaping functions. Virtual circuits can be assigned to one of eight rate queues, each of which is programmable for a different peak rate. Each virtual circuit can be assigned an average rate and specific burst size. The signaling request specifies the size of the burst that will be sent at the peak rate, and after that burst, the rest of the data will be sent at the average rate.

Following are the configurable traffic parameters on the AIP:

- Forward peak cell rate
- Backward peak cell rate
- Forward sustainable cell rate
- Backward sustainable cell rate
- Forward maximum burst
- Backward maximum burst

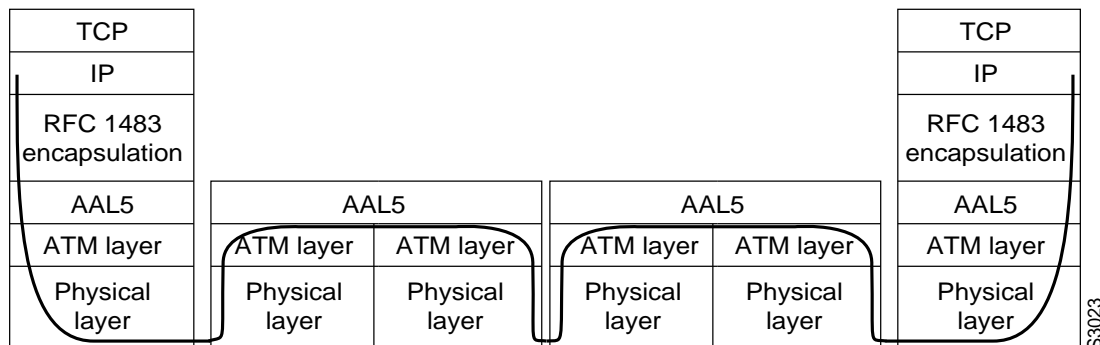
Figure 5-9 shows how the routing table and address resolution table on Router A are used to forward data to a workstation behind Router C.

Figure 5-9 AIP Connects LANs to ATM Fabric



The routing table on Router A performs its usual function of determining the next hop by mapping the network number of the destination (in this case 144.254.45 from the incoming packet) to the IP address of the router to which the destination network is connected (in this case, 144.254.10.3, which is the IP address of Router C). An address resolution table maps the next-hop IP address to an ATM NSAP address (in this case, represented by π). Router A signals Router C over the ATM network to establish a virtual connection, and Router A uses that connection to forward the packet to Router C. Figure 5-10 shows the layers through which the packet travels.

Figure 5-10 Path of an IP Packet over the ATM Fabric



Configuring the AIP for ATM Signaling

The following commands configure an AIP for ATM signaling:

```
interface atm 4/0
ip address 128.24.2.1 255.255.255.0
no keepalive
atm nsap-address AB.CDEF.01.234567.890A.BCDE.F012.3456.7890.1234.12
atm pvc 1 0 5 qsaal
map-group shasta
atm rate-queue 0 155
atm rate-queue 1 45

map-list shasta
ip 144.222.0.0 atm-nsap BB.CDEF.01.234567.890A.BCDE.F012.3456.7890.1234.12
ip 144.3.1.2 atm-nsap BB.CDEF.01.234567.890A.BCDE.F012.3456.7890.1234.12 class QOSclass

map-class QOSclass
atm forward-peak-cell-rate-clp0 15000
atm backward-max-burst-size-clp0 96
```

The following explains relevant portions of the ATM signaling configuration:

- **no keepalive**—Required because Cisco IOS Software Release 10.0 does not support the Interim Local Management Interface (ILMI), an ATM Forum specification.
- **atm nsap-address**—Required for signaling.
- **atm pvc**—Sets up a PVC to carry signaling requests to the switch. In this case, the command sets up a circuit whose VPI value is 0 and whose VCI value is 5, as recommended by the ATM Forum.
- **map-group**—Associates a map list named **shasta** to this interface.
- **atm rate-queue**—Set up two rate queues. Rate queue number 0 is for 155 Mbps transfers, and rate queue number 1 is for 45-Mbps transfers.
- **map-list** and **ip 144.222.0.0**—Set up the static mapping of an IP network number to an ATM NSAP address without any QOS parameters. The **ip 144.3.1.2** command maps an IP host address to an ATM NSAP address with the QOS parameters specified in the map class named **QOSclass**.
- **map-class**, **atm forward-peak-cell-rate-clp0**, and **atm backward-max-burst-size-clp0**—Set up QOS parameters associated with this connection. The connection must support a forward peak cell rate of 15 kbps and a backward burst size of 96 cells.

Interoperability with DXI

When configuring an AIP to communicate with a Cisco router that uses ATM DXI to connect to the ATM network, the AIP requires Network Layer Protocol Identifier (NLPID) encapsulation, which is provided in Cisco IOS Software Release 10.2, or the ATM DXI requires LLC/SNAP encapsulation.

Cisco LightStream 100

The Cisco LightStream 100 is an ATM switch that provides up to 16 modular 155-Mbps ATM interfaces. Each ATM interface is capable of providing 4096 point-to-point connections. The A100 can also support up to 1024 point-to-multipoint connections.

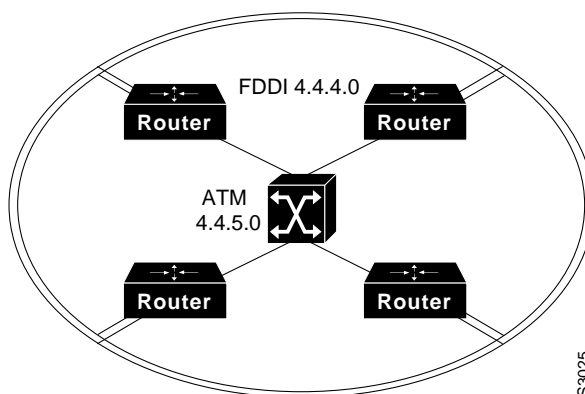
Table 5-1 Cisco LightStream 100 Physical Layer Support

Physical Layer	Data Rate	Media	Connector
STS 3c/STM1	155 Mbps	Multimode fiber	SC
TAXI 4B/5B	100 Mbps	Multimode fiber	MIC (FDDI-style)
STS3c/STM1	155 Mbps	Single-mode fiber	ST
DS-3	45 Mbps	Coaxial cable	BNC
E3	34 Mbps	Coaxial cable	BNC

Single-Switch Designs

Because ATM can use existing multimode fiber networks, Fiber Distributed Data Networks (FDDI) campus backbones can be easily upgraded from 100-Mbps FDDI to 155-Mbps point-to-point ATM. If the network has spare fiber, AIPs can be installed in each router and interconnected with a LightStream 100, as shown in Figure 5-11. In this topology, each router has a 155-Mbps point-to-point connection to every other router on the ring.

Figure 5-11 Parallel FDDI and ATM Backbone



S3025

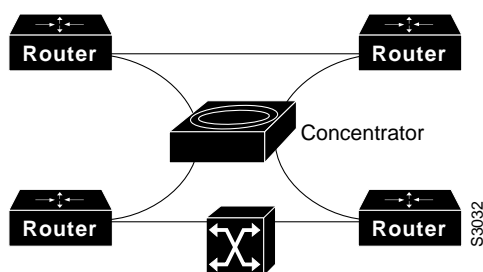
The addition of the ATM switch creates a parallel subnet. During the migration to ATM, a routing protocol, such as the Interior Gateway Routing Protocol (IGRP), can be used to force FDDI routing, as shown by the following commands:

```
interface fddi 1/0
ip address 4.4.4.1 255.255.255.0
interface atm 2/0
ip address 4.4.5.1 255.255.255.0
router igrp 109
network 4.4.0.0
distance 150 4.4.5.0 0.0.0.255
```

The **distance** command causes ATM to appear as a less desirable network and forces routing over FDDI.

If the network does not have spare fiber, a concentrator can be installed. Later, an ATM switch can be installed, as shown in Figure 5-12, which can be used to migrate ATM slowly throughout the network, using FDDI as a backup.

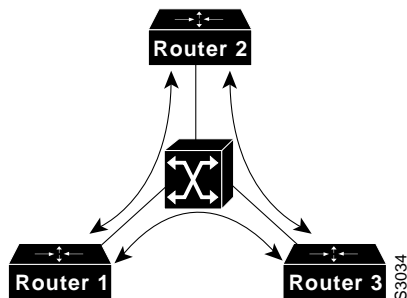
Figure 5-12 FDDI Topology with Concentrator and ATM Switch



Broadcasting in Single-Switch ATM Networks

There are two ways to configure broadcasting in a single-switch ATM network. First, the routers can be configured for “pseudo” broadcasting over point-to-point PVCs, as shown in Figure 5-13.

Figure 5-13 Router-Based “Pseudo” Broadcasting Using Point-to-Point PVCs



The following commands on each router set up a PVC between each router:

```
atm pvc 1 1 1 aal5snap
atm pvc 2 2 1 aal5snap
atm pvc 3 3 1 aal5snap
```

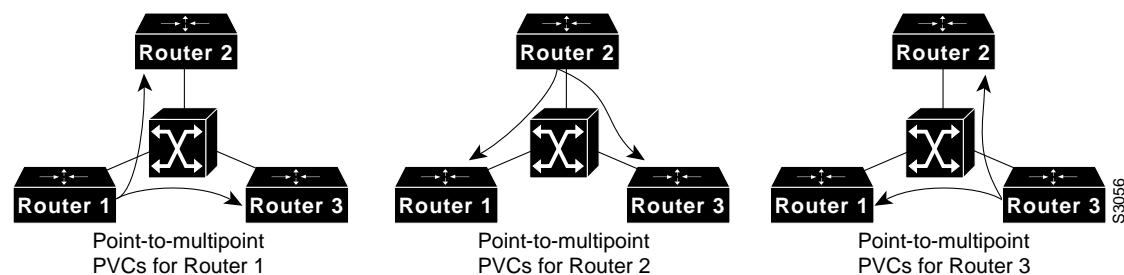
The following commands on each router cause that router to replicate broadcast packets and send them out on each PVC:

```
ip 4.4.5.1 atm-vc 1 broadcast
ip 4.4.5.2 atm-vc 2 broadcast
ip 4.4.5.3 atm-vc 3 broadcast
```

The disadvantage of router-based broadcasting is that it places the burden of replicating packets on the routers instead of on the switch, which has the resources to replicate packets at a lower cost to the network.

The second way to configure broadcasting is to configure the routers for switch-based broadcasting, as shown in Figure 5-14. With switch-based broadcasting, each router sets up a point-to-multipoint PVC to the other routers in the network. When each router maintains a point-to-multipoint PVC to every other router in the network, the broadcast replication burden is transferred to the switch.

Figure 5-14 Switch-Based Broadcasting



The following commands configure a point-to-multipoint PVC on each router:

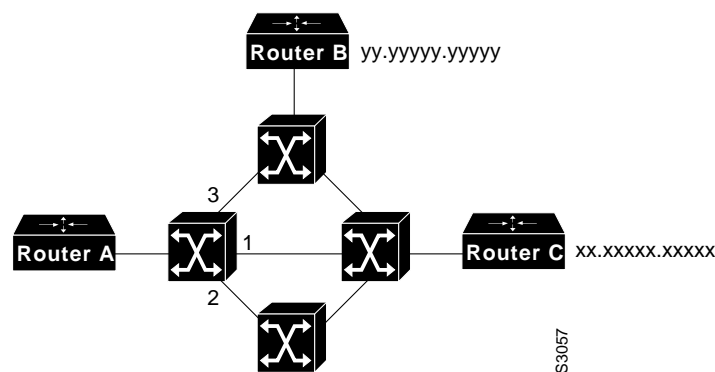
```
ip 4.4.4.1 atm-vc 1
ip 4.4.4.2 atm-vc 2
ip 4.4.4.3 atm-vc 3
ip 4.4.4.0 atm-vc broadcast
```

In Figure 5-14, the routers still have full mesh connectivity to every other router in the network, but the connections are not set up as broadcast PVCs. Instead, each router designates the point-to-multipoint PVC as a broadcast PVC and lets the switch handle replication, which is a function for which the switch is optimized.

Multiple-Switch Designs

The A100 supports the ATM Forum Private Network-Network Interface (PNNI) Phase 0 protocol, which uses static maps to switch around failed links. Figure 5-15 shows the static maps on the switch to which Router A is connected.

Figure 5-15 Example of a Multi-Switch Network That Uses the PNNI Phase 0 Protocol



When a physical link fails, the ATM switch tears down the virtual circuits for that link. When the AIP in Router A detects that a virtual circuit has been torn down, it resignals the network to reestablish the VCC. When the switch receives the new signaling packet and realizes that the primary interface is down, it forwards the request on the alternate interface.

ATM Media

The ATM Forum has defined multiple standards for encoding ATM over various types of media. Table 5-2 lists the framing type and data rates for the various media, including unshielded twisted-pair (UTP) and shielded twisted-pair (STP) cable.

Table 5-2 ATM Physical Rates

Media							
Framing	Data Rate (Mbps)	Multimode Fiber	Single Mode Fiber	Coaxial Cable	UTP-3	UTP-5	STP
DS-1	1.544						
E1	2.048						
DS-3	45						
E3	34						
STS-1	51						
SONET STS3c	155						
SDH STM1							
SONET STS12c	622						
SDH STM4							
TAXI 4B/5B	100						
8B/10B	155						
(Fiber Channel)							

Because the FDDI chipset standard, TAXI 4B/5B, was readily available, the ATM Forum encouraged initial ATM development efforts by endorsing TAXI 4B/5B as one of the first ATM media encoding standards. Today, however, the most common fiber interface is STS3c/STM.

There are two standards for running ATM over copper cable: UTP-3 and UTP-5. The UTP-5 specification supports 155 Mbps with NRZI encoding, while the UTP-3 specification supports 51 Mbps with CAP-16 encoding. CAP-16 is more difficult to implement, so, while it may be cheaper to wire with UTP-3 cable, workstation cards designed for CAP-16 based UTP-3 may be more expensive and will offer less bandwidth.

Because ATM is designed to run over fiber and copper cable, investments in these media today will maintain their value when networks migrate to full ATM implementations as ATM technology matures.

