

Designing ATM Internetworks

This chapter describes current Asynchronous Transfer Mode (ATM) technologies that network designers can use in their networks today. It also makes recommendations for designing non-ATM networks so 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
- Cisco's ATM WAN solutions

ATM Defined

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.

Role of ATM in Internetworks

Today, 90 percent of computing power resides on desktops, and that power is growing exponentially. Distributed applications are increasingly bandwidth-hungry, and the emergence of the Internet is driving most LAN architectures to the limit. Voice communications have increased significantly with increasing reliance on centralized voice mail systems for verbal communications. The internetwork is the critical tool for information flow. Internetworks are being pressured to cost less yet support the emerging applications and higher number of users with increased performance.


To date, local and wide-area communications have remained logically separate. In the LAN, bandwidth is free and connectivity is limited only by hardware and implementation cost. The LAN has carried data only. In the WAN, bandwidth has been the overriding cost, and such delay-sensitive traffic as voice has remained separate from data. New applications and the economics of supporting them, however, are forcing these conventions to change.

The Internet is the first source of multimedia to the desktop and immediately breaks the rules. Such Internet applications as voice and real-time video require better, more predictable LAN and WAN performance. In addition, the Internet also necessitates that the WAN recognize the traffic in the LAN stream, thereby driving LAN/WAN integration.

Multiservice Networks

ATM has emerged as one of the technologies for integrating LANs and WANs. ATM can support any traffic type in separate or mixed streams, delay-sensitive traffic, and nondelay-sensitive traffic, as shown in Figure 8-1.

Figure 8-1 ATM support of various traffic types.



ATM can also scale from low to high speeds. It has been adopted by all the industry’s equipment vendors, from LAN to private branch exchange (PBX). With ATM, network designers can integrate LANs and WANs, support emerging applications with economy in the enterprise, and support legacy protocols with added efficiency.

TDM Network Migration

In addition to using ATM to combine multiple networks into one multiservice network, network designers are deploying ATM technology to migrate from TDM networks for the following reasons:

- To reduce WAN bandwidth cost
- To improve performance
- To reduce downtime

Reduced WAN Bandwidth Cost

The Cisco line of ATM switches provide additional bandwidth through the use of voice compression, silence compression, repetitive pattern suppression, and dynamic bandwidth allocation. The Cisco implementation of ATM combines the strengths of 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.

Improved Performance

ATM offers improved performance through performance guarantees and robust WAN traffic management that support the following capabilities:

- Large buffers that guarantee Quality of Service (QoS) for bursty data traffic and demanding multimedia applications
- Per-virtual circuit (VC) queuing and rate scheduling
- Feedback—congestion notification

Reduced Downtime


ATM offers high reliability, thereby reducing downtime. This high reliability is available because of the following ATM capabilities:

- The capability to support redundant processors, port and trunk interfaces, and power supplies
- The capability to rapidly reroute around failed trunks

Integrated Solutions

The trend in internetworking is to provide network designers greater flexibility in solving multiple internetworking problems without creating multiple networks or writing off existing data communications investments. Routers can provide a reliable, secure network and act as a barrier against inadvertent broadcast storms in the local networks. Switches, which can be divided into two main categories—LAN switches and WAN switches—can be deployed at the workgroup, campus backbone, or WAN level, as shown in Figure 8-2.

Figure 8-2 The role of ATM switches in an internetwork.



Underlying and integrating all Cisco products is the Cisco IOS software. The Cisco IOS software enables disparate groups, diverse devices, and multiple protocols all to be integrated into a highly reliable and scalable network.

Different Types of ATM Switches

Even though all ATM switches perform cell relay, ATM switches differ markedly in the following ways:

- Variety of interfaces and services that are supported
- Redundancy
- Depth of ATM internetworking software
- Sophistication of traffic management mechanism

Just as there are routers and LAN switches available at various price/performance points with different levels of functionality, ATM switches can be segmented into the following four distinct types that reflect the needs of particular applications and markets:

- Workgroup ATM switches
- Campus ATM switches
- Enterprise ATM switches
- Multiservice access switches

As Figure 8-2 shows, Cisco offers a complete range of ATM switches.

Workgroup and Campus ATM Switches

Workgroup ATM switches are characterized by having Ethernet switch ports and an ATM uplink to connect to a campus ATM switch. An example of a workgroup ATM switch is the Cisco Catalyst 5000.

The Catalyst 5500 switch provides high-performance switching between workstations, servers, switches, and routers in wiring closet, workgroup, and campus backbone environments.

The Catalyst 5500 LAN is a 13-slot switch. Slot 1 is reserved for the supervisor engine module, which provides switching, local and remote management, and dual Fast Ethernet uplinks. Slot 2 is available for a second, redundant supervisor engine, or any of the other supported modules. Slots 3–12 support any of the supported modules.

Slot 13 can be populated only with a LightStream 1010 ATM Switch Processor (ASP). If an ASP is present in slot 13, slots 9–12 support any of the standard LightStream 1010 ATM switch port adapter modules (PAMs).

The Catalyst 5500 has a 3.6-Gbps media-independent switch fabric and a 5-Gbps cell-switch fabric. The backplane provides the connection between power supplies, supervisor engine, interface modules, and backbone module. The 3.6-Gbps media-independent fabric supports Ethernet, Fast Ethernet, FDDI/CDDI, ATM LAN Emulation, and RSM modules. The 5-Gbps cell-based fabric supports a LightStream 1010 ASP module and ATM PAMs.

Campus ATM switches are generally used for small-scale ATM backbones (for instance, to link ATM routers or LAN switches). This use of ATM switches can alleviate current backbone congestion while enabling the deployment of such new services as virtual LANs (VLANs). Campus switches need to support a wide variety of both local backbone and WAN types but be price/performance optimized for the local backbone function. In this class of switches, ATM routing capabilities that allow multiple switches to be tied together is very important. Congestion control mechanisms for optimizing backbone performance is also important. The LightStream 1010 family of ATM switches is an example of a campus ATM switch. For more information on deploying workgroup and campus ATM switches in your internetwork, see Chapter 12, “Designing Switched LAN Internetworks.”

Enterprise ATM Switches

Enterprise ATM switches are sophisticated multiservice devices that are designed to form the core backbones of large, enterprise networks. They are intended to complement the role played by today’s high-end multiprotocol routers. Enterprise ATM switches are used to interconnect campus ATM switches. Enterprise-class switches, however, can act not only as ATM backbones but can serve as the single point of integration for all of the disparate services and technology found in enterprise backbones today. By integrating all of these services onto a common platform and a common ATM transport infrastructure, network designers can gain greater manageability and eliminate the need for multiple overlay networks.

Cisco’s BPX/AXIS is a powerful broadband ATM switch designed to meet the demanding, high-traffic needs of a large private enterprise or public service provider. This chapter focuses on this category of ATM switches.

Multiservice Access Switches

Beyond private networks, ATM platforms will also be widely deployed by service providers both as customer premises equipment (CPE) and within public networks. Such equipment will be used to support multiple MAN and WAN services—for instance, Frame Relay switching, LAN

interconnect, or public ATM services—on a common ATM infrastructure. Enterprise ATM switches will often be used in these public network applications because of their emphasis on high availability and redundancy, their support of multiple interfaces, and capability to integrate voice and data.

ATM Overview


This section discusses the following ATM concepts:

- Structure of an ATM Network
- ATM Functional Layers
- ATM Addressing
- ATM Media
- ATM Data Exchange Interface

Structure of an ATM Network

ATM is based on the concept of two end-point devices communicating by means of intermediate switches. As Figure 8-3 shows, an ATM network is made up of a series of switches and end-point devices. The end-point devices can be ATM-attached end stations, ATM-attached servers, or ATM-attached routers.

Figure 8-3 Components of an ATM network.



As Figure 8-3 shows, there are two types of interfaces in an ATM network:

- User-to-Network Interface (UNI)
- Network-to-Network Interface (NNI)

The UNI connection is made up of an end-point device and a private or public ATM switch. The NNI is the connection between two ATM switches. The UNI and NNI connections can be carried by different physical connections.

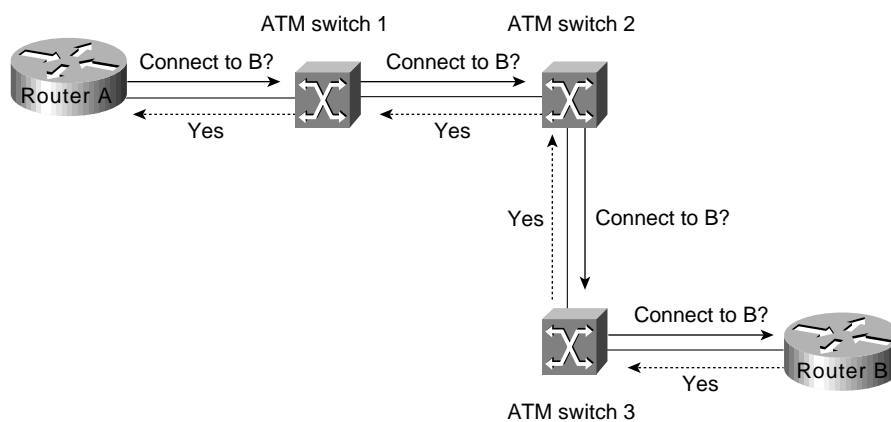
In addition to the UNI and NNI protocols, the ATM Forum has defined a set of LAN Emulation (LANE) standards and a Private Network to Network Interface (PNNI) Phase 0 protocol. LANE is a technology network designers can use to interconnect legacy LANs such as Ethernet and Token Ring with ATM-attached devices. Most LANE networks consist of multiple ATM switches and typically employ the PNNI protocol.

The full PNNI 1.0 specification was released by the ATM Forum in May 1996. It enables extremely scalable, full function, dynamic multi-vendor ATM networks by providing both PNNI routing and PNNI signaling. PNNI is based on UNI 3.0 signaling and static routes. The section "Role of LANE" later in this chapter discusses ATM LANE networks in detail.

General Operation on an ATM Network

Because ATM is connection-oriented, a connection must be established between two end points before any data transfer can occur. This connection is accomplished through a signaling protocol as shown in Figure 8-4.

Figure 8-4 Establishing a connection in an ATM network.



As Figure 8-4 shows, for Router A to connect to Router B the following must occur:

- 1 Router A sends a signaling request packet to its directly connected ATM switch (ATM Switch 1). This request contains the ATM address of the Router B as well as any QoS parameters required for the connection.
- 2 ATM Switch 1 reassembles the signaling packet from Router A, and then examines it.
- 3 If ATM Switch 1 has an entry for Router B's ATM address in its switch table and it can accommodate the QoS requested for the connection, it sets up the virtual connection and forwards the request to the next switch (ATM Switch 2) along the path.
- 4 Every switch along the path to Router B reassembles and examines the signaling packet, and then forwards it to the next switch if the QoS parameters can be supported. Each switch also sets up the virtual connection as the signaling packet is forwarded.

If any switch along the path cannot accommodate the requested QoS parameters, the request is rejected and a rejection message is sent back to Router A.

- 5 When the signaling packet arrives at Router B, Router B reassembles it and evaluates the packet. If Router B can support the requested QoS, it responds with an accept message. As the accept message is propagated back to Router A, the switches set up a virtual circuit.

Note 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.

- 6 Router A receives the accept message from its directly connected ATM switch (ATM Switch 1), as well as the Virtual path identifier (VPI) and Virtual channel identifier (VCI) values that it should use for cells sent to Router B.

Note ATM cells consist of five bytes of header information and 48 bytes of payload data. The VPI and VCI fields in the ATM header are used to route cells through ATM networks. The VPI and VCI fields of the cell header identify the next network segment that a cell needs to transmit on its way to its final destination.

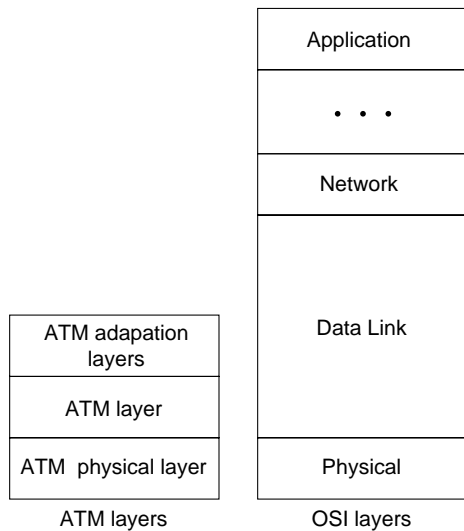
ATM Functional Layers

Just as the *Open System Interconnection* (OSI) reference model describes how two computers communicate over a network, the ATM protocol model describes how two end systems communicate through ATM switches. The ATM protocol model consists of the following three functional layers:

- ATM physical layer
- ATM layer
- ATM adaptation layer

As Figure 8-5 shows, these three layers correspond roughly to Layer 1 and parts of Layer 2 (such as error control and data framing) of the OSI reference model.

Figure 8-5 Relationship of ATM functional layers to the OSI reference model.



Physical Layer

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. The ATM physical layer is divided into two parts:

- Physical medium sublayer
- Transmission convergence sublayer

Physical Medium Sublayer

The physical medium sublayer is responsible for sending and receiving a continuous flow of bits with associated timing information to synchronize transmission and reception. Because it includes only physical-medium-dependent functions, its specification depends on the physical medium used. ATM can use any physical medium capable of carrying ATM cells. Some existing standards that can carry ATM cells are SONET (Synchronous Optical Network)/SDH, DS-3/E3, 100-Mbps local fiber (Fiber Distributed Data Interface [FDDI] physical layer), and 155-Mbps local fiber (Fiber Channel physical layer). Various proposals for use over twisted-pair wire are also under consideration.

Transmission Convergence Sublayer

The transmission convergence sublayer is responsible for the following:

- *Cell delineation*—Maintains ATM cell boundaries.
- *Header error control sequence generation and verification*—Generates and checks the header error control code to ensure valid data.
- *Cell rate decoupling*—Inserts or suppresses idle (unassigned) ATM cells to adapt the rate of valid ATM cells to the payload capacity of the transmission system.
- *Transmission frame adaptation*—Packages ATM cells into frames acceptable to the particular physical-layer implementation.
- *Transmission frame generation and recovery*—Generates and maintains the appropriate physical-layer frame structure.

ATM Layer

The ATM layer establishes virtual connections and passes ATM cells through the ATM network. To do this, it uses the information contained in the header of each ATM cell. The ATM layer is responsible for performing the following four basic functions:

- Multiplexing and demultiplexing the cells of different virtual connections. These connections are identified by their VCI and VPI values.
- Translating the values of the VCI and VPI at the ATM switches or cross connects.
- Extracting and inserting the header before or after the cell is delivered to or from the higher ATM adaptation layer.
- Handling the implementation of a flow control mechanism at the UNI.

ATM Adaptation Layer (AAL)

The AAL translates between the larger service data units (SDUs) (for example, video streams and data packets) of upper-layer processes and ATM cells. Specifically, the AAL receives packets from upper-level protocols (such as AppleTalk, Internet Protocols [IP], and NetWare) and breaks them into the 48-byte segments that form the payload field of an ATM cell. Several ATM adaptation layers are currently specified. Table 8-1 summarizes the characteristics of each AAL.

Table 8-1 ATM Adapter Layers

Characteristics	AAL1	AAL3/4	AAL4	AAL5
Requires timing between source and destination	Yes	No	No	No

Characteristics	AAL1	AAL3/4	AAL4	AAL5
Data rate	Constant	Variable	Variable	Variable
Connection mode	Connection-oriented	Connection-oriented	Connectionless	Connection-oriented
Traffic types	Voice and circuit emulation	Data	Data	Data

AAL1

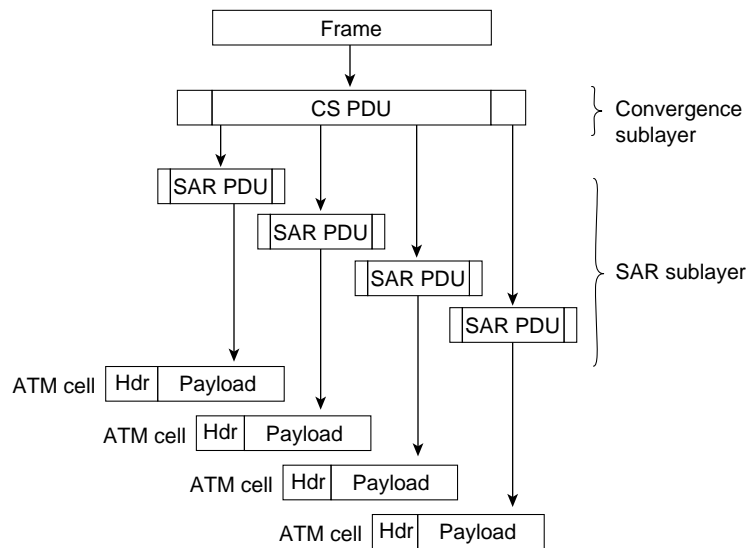
AAL1 prepares a cell for transmission. The payload data consists of a synchronous sample (for example, one byte of data generated at a sampling rate of 125 microseconds). The sequence number field (SN) and sequence number protection (SNP) fields provide the information that the receiving AAL1 needs to verify that it has received the cells in the correct order. The rest of the payload field is filled with enough single bytes to equal 48 bytes.

AAL1 is appropriate for transporting telephone traffic and uncompressed video traffic. It requires timing synchronization between the source and destination and, for that reason, depends on a media that supports clocking, such as SONET. The standards for supporting clock recovery are currently being defined.

AAL3/4

AAL3/4 was designed for network service providers and is closely aligned with Switched Multimegabit Data Service (SMDS). AAL3/4 is used to transmit SMDS packets over an ATM network. The convergence sublayer (CS) creates a protocol data unit (PDU) by prepending a Beginning/End Tag header to the frame and appending a length field as a trailer as shown in Figure 8-6.

Figure 8-6 AAL3/4 cell preparation.



The segmentation and reassembly (SAR) sublayer fragments the PDU and prepends to each PDU fragment a header consisting of the following fields:

- *Type*—Identifies whether the cell is the beginning of a message, continuation of a message, or end of a message.
- *Sequence number*—Identifies the order in which cells should be reassembled.


- *Multiplexing identifier*—Identifies cells from different traffic sources interleaved on the same virtual circuit connection (VCC) so that the correct cells are reassembled at the destination.

The SAR sublayer also appends a CRC-10 trailer to each PDU fragment. The completed SAR PDU becomes the payload field of an ATM cell to which the ATM layer prepends the standard ATM header.

AAL5

AAL5 prepares a cell for transmission as shown in Figure 8-7.

Figure 8-7 AAL5 cell preparation.



First, the convergence sublayer of AAL5 appends a variable-length pad and an 8-byte trailer to a frame. The pad is long enough to ensure that the resulting PDU falls on the 48-byte boundary of the ATM cell. The trailer includes the length of the frame and a 32-bit CRC computed across the entire PDU, which allows AAL5 at the destination to detect bit errors and lost cells or cells that are out of sequence.

Next, the segmentation and reassembly segments the CS PDU into 48-byte blocks. Then the ATM layer places each block into the payload field of an ATM cell. For all cells except the last cell, a bit in the PT field is set to zero to indicate that the cell is not the last cell in a series that represents a single frame. For the last cell, the bit in the PT field is set to one. When the cell arrives at its destination, the ATM layer extracts the payload field from the cell; the SAR sublayer reassembles the CS PDU; and the CS uses the CRC and the length field to verify that the frame has been transmitted and reassembled correctly.

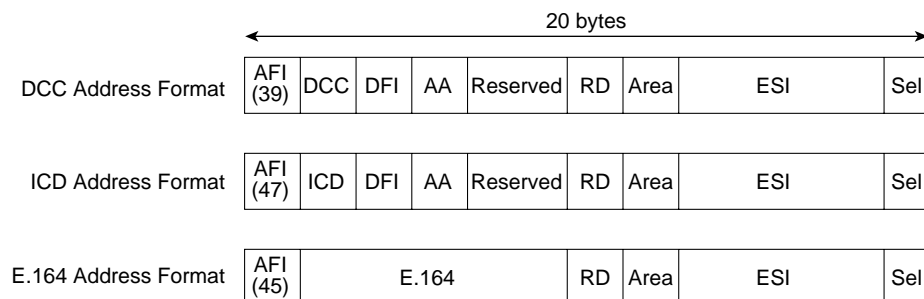
AAL5 is the adaptation layer used to transfer most non-SMDS data, such as classical IP over ATM and local-area network (LAN) emulation.

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 typically use E.164 numbers, which are also used by Narrowband ISDN (N-ISDN) networks.

Figure 8-8 shows the format of private network ATM addresses. The three formats are Data Country Code (DCC), International Code Designator (ICD), and Network Service Access Point (NSAP) encapsulated E.164 addresses.

Figure 8-8 ATM address formats.



Fields of an ATM Address

The fields of an ATM address are as follows:

- *AFI*—One byte of authority and format identifier. The AFI field identifies the type of address. The defined values are 45, 47, and 39 for E.164, ICD, and DCC addresses, respectively.
- *DCC*—Two bytes of data country code.
- *DFI*—One byte of domain specific part (DSP) format identifier.
- *AA*—Three bytes of administrative authority.
- *RD*—Two bytes of routing domain.
- *Area*—Two bytes of area identifier.
- *ESI*—Six bytes of end system identifier, which is an IEEE 802 Media Access Control (MAC) address.
- *Sel*—One byte of Network Service Access Point (NSAP) selector.
- *ICD*—Two bytes of international code designator.
- *E.164*—Eight bytes of Integrated Services Digital Network (ISDN) telephone number.

The ATM address formats are modeled on ISO 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.

ATM Media

The ATM Forum has defined multiple standards for encoding ATM over various types of media. Table 8-2 lists the framing type and data rates for the various media, including unshielded twisted-pair (UTP) and shielded twisted-pair (STP) cable.

Table 8-2 ATM Physical Rates

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 SDH STM1	155						
SONET STS12c SDH STM4	622						
TAXI 4B/5B	100						
8B/10B (Fiber Channel)	155						

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.

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 UNI support between Cisco routers and ATM networks, as shown in Figure 8-9.

Figure 8-9 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 8-10 shows the format of a Mode 1a ATM DXI frame.

Figure 8-10 ATM DXI frame format.




In Figure 8-11, 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 CS and the SAR sublayer. The ATM layer attaches the header, and the cells are sent out the ATM UNI interface.

Figure 8-11 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 8-12 shows how the DSU performs address mapping.

Figure 8-12 ATM DXI address mapping.



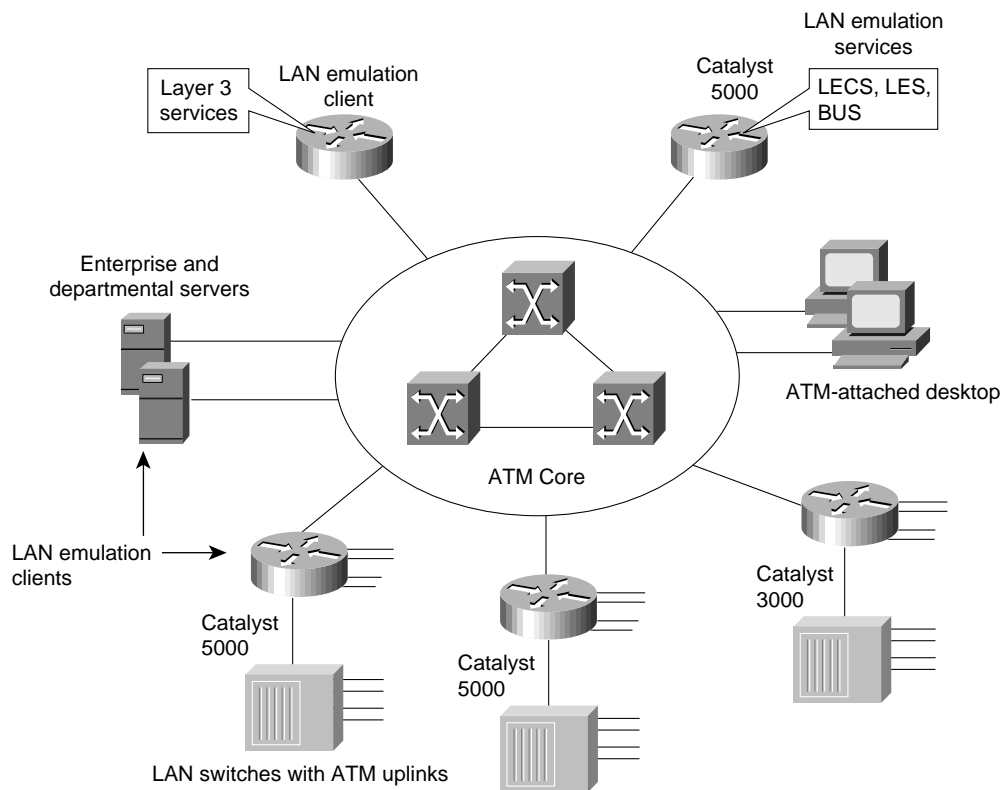
Note ATM DXI 3.2 is supported in the Cisco IOS Software Release 9.21 or later. Mode 1a is the only mode supported.

Role of LANE

The ATM Forum has defined a standard for LANE. LANE is a technology that network designers can deploy to internetwork their legacy LANs (for example, Ethernet and Token Ring LANs), with ATM-attached devices. 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 (ELAN) 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 ELAN, and routers handle inter-ELAN traffic. Figure 8-13 shows an example of an ATM LANE network.

Figure 8-13 Components of an ATM LANE network.



As Figure 8-13 shows, network designers can use the LANE technology to interconnect legacy LANs to any of the following types of ATM-attached devices:

- End stations (for example, ATM-attached servers or ATM-attached workstations)
- Edge devices that bridge the legacy LANs onto an ATM backbone (for example, the Catalyst 5000 or Catalyst 3000 switches that have an ATM uplink)
- ATM-attached routers that are used to route between ELANs

LANE Components

LANE components include the following:

- *LAN emulation client (LEC)*—End systems that support LANE, such as network interface card (NIC)-connected workstations, LAN switches with ATM uplinks (for example, the Catalyst family of switches), and Cisco 7500, 7000, 4500, and 4000 series routers that support ATM attachment, all 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 configuration server (LECS)*—The LECS maintains a database of ELANs and the ATM addresses of the LESs that control the ELANs. It accepts queries from LECs and responds with the ATM address of the LES that serves the appropriate ELAN/VLAN. This database is defined and maintained by the network administrator.

The following is an example of this database.

ELAN Name	LES ATM Address
finance	47.0091.8100.0000.0800.200c.1001. 0800.200c.1001.01
marketing	47.0091.8100.0000.0800.200c.1001. 0800.200c.1001.02

- *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 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.

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.

Note that because LANE is defined at OSI Layer 2, the LES 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 LANE Works

An ELAN provides Layer 2 communication between all users on an ELAN. One or more ELANs can run on the same ATM network. However, each ELAN is independent of the others and users on separate ELANs cannot communicate directly. Communication between ELANs is possible only through routers or bridges.

Because an ELAN provides Layer 2 communication, it can be equated to a broadcast domain. VLANs can also be thought of as broadcast domains. This makes it possible to map an ELAN to a VLAN on Layer 2 switches with different VLAN multiplexing technologies such as Inter-Switch Link (ISL) or 802.10. In addition, IP subnets and IPX networks that are defined on Layer 3-capable devices such as routers frequently map into broadcast domains (barring secondary addressing). This makes it possible to assign an IP subnetwork or an IP network to an ELAN.

An ELAN is controlled by a single LES/BUS pair and the mapping of an ELAN to its LES ATM address is defined in the LES database. ELANs consist of multiple LECs and can be Ethernet or Token Ring but not both at the same time.

In order for ELAN to operate properly, the LECs on that ELAN need to be operational. Each LEC goes through a boot up sequence that is described in the following sections.

LANE Operation

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.

Finding the LECS

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 permanent virtual circuit (PVC) 0/17, a “well-known” PVC.

Contacting the LECS

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. 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.

Configuring the LECS database

You can configure the LECS database in any of the following three ways:

- *Configure ELAN names at the LEC*—In this configuration, all the LECs are configured with an ELAN name that they can embed in their Configure_Requests. This is the most basic form of the LECS database and it needs only to contain the list of ELANs and their corresponding LES ATM addresses. In such a configuration, all LECs that specifically request to join a given ELAN are returned the ATM address of the corresponding LES. A LEC that does not know which ELAN to join can be assigned to a default ELAN if such an ELAN is configured in the LECS database.

The following is an example of LEC-to-ELAN mapping at the LEC:

```
lane database test-1
name finance server-atm-address 47.0091.8100.0000.0800.200c.1001. 0800.200c.1001.01
name marketing server-atm-address 47.0091.8100.0000.0800.200c.1001. 0800.200c.1001.02
default-name finance
```

- *Configure LEC to ELAN assignment in the LECS database*—In this configuration, all the information is centralized in the LECS database. The LECs do not need to be intelligent, and they can simply go to the LECS to determine which ELAN they should join. Although this is a more time-intensive configuration, it provides tighter control over all the ELANs. Consequently, it can be useful when security is important.

With this method, the LECs are identified by their ATM addresses or MAC addresses. Because wildcarding of ATM address prefixes is also supported, it is useful to make such relationships as “Assign any LEC joining with a prefix of A to ELAN X.” The following is an example of LEC-to-ELAN mapping in the LECS database:

```
lane database test-2
name finance server-atm-address 47.0091.8100.0000.0800.200c.1001. 0800.200c.1001.01
name marketing server-atm-address 47.0091.8100.0000.0800.200c.1001. 0800.200c.1001.02
default-name finance

client-atm-address 47.0091.8100.0000.08... name finance
client-atm-address 47.0091.8100.0000.09... name marketing

mac-address 00c0.0000.0100 name finance
mac-address 00c0.1111.2222 name marketing
```


- *Hybrid combination*—You can configure a combination of the preceding two methods.

Joining the LES

After 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 8-14 shows examples of LES connections.

Figure 8-14 LAN emulation server (LES) connections.



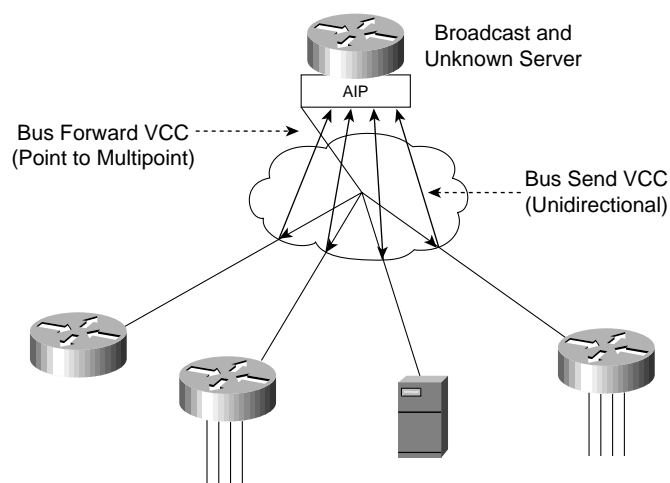
Finding the BUS

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 ELAN. Figure 8-15 shows examples of BUS connections.

Figure 8-15 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 against 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 8-16 shows an example of a fully connected ELAN.

Figure 8-16 Fully connected ELAN.



LANE Implementation

As Table 8-3 indicates, the LANE functionality (the LECS, LEC, LES, and BUS) can be implemented in different Cisco devices.

Table 8-3 Cisco LANE Implementation

Cisco Product	Available LANE Components	Required Software Release
Family of Catalyst 5000 switches	LECS, LES, BUS, LEC	ATM Module Software Version 2.0 or later
Family of Catalyst 3000 switches	LECS, LES, BUS, LEC	ATM Module Software Version 2.1 or later
Family of Cisco 7000 routers	LECS, LES, BUS, LEC	Cisco IOS Software Release 11.0 or later
Family of Cisco 7500 routers	LECS, LES, BUS, LEC	Cisco IOS Software Release 11.1 or later
Family of Cisco 4500 and 4000 routers	LECS, LES, BUS, LEC	Cisco IOS Software Release 11.1 or later

These functions will be defined on ATM physical interfaces and subinterfaces. A subinterface can be defined as a logical interface and is a part of a physical interface such as an Optical Carrier 3 (OC-3) fiber. ATM interfaces on the Cisco routers and the ATM module on the Catalyst 5000 switch can be logically divided into up to 255 logical subinterfaces. On the Catalyst 3000 switch, although the same Cisco IOS Software code is used, the subinterface concept does not apply. The LEC can be configured using the menu-driven interface.

This section examines the implementation of ATM LANE networks and covers the following topics:

- LANE Design Considerations
- LANE Redundancy

LANE Design Considerations

The following are some general LANE design considerations:

- The AIP provides an interface to ATM switching fabrics for transmitting and receiving data at rates of up to 155 Mbps bidirectionally. The actual rate is determined by the Physical layer interface module (PLIM).
- One active 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 active 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 might make anything more than a very small network prohibitively unmanageable and complex.
- When defining VLANs with the Catalyst 5000 switch, each VLAN 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 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.

PNNI in LANE Networks

Network designers can deploy PNNI as a Layer 2 routing protocol for bandwidth management, traffic distribution, and path redundancy for LANE networks. PNNI is an ATM routing protocol used for routing call setups and is implemented in the ATM switches. Most LANE networks consist of multiple ATM switches and typically employ the PNNI protocol.

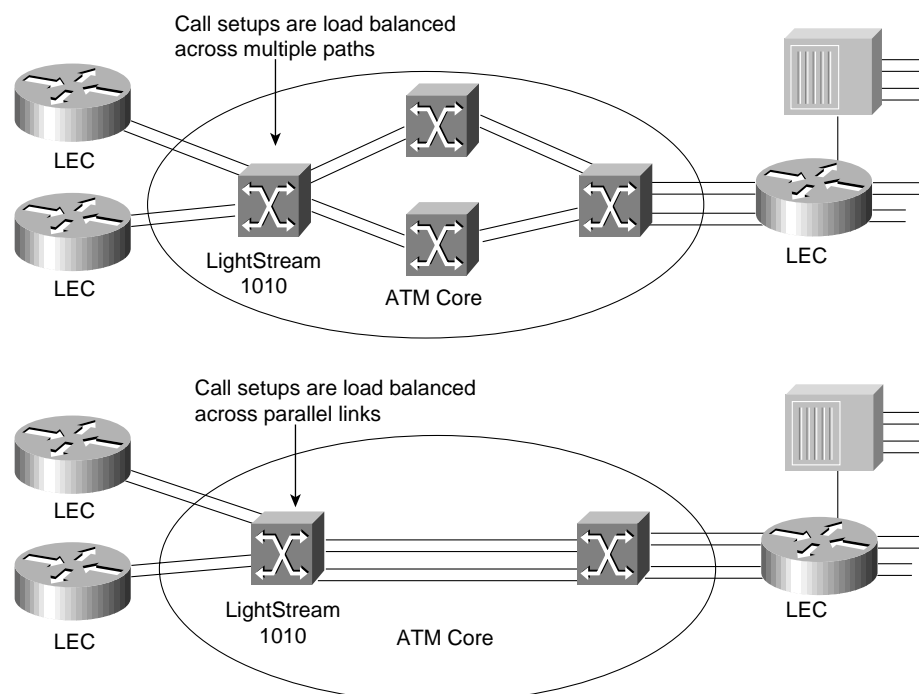
Note Although PNNI is an advanced routing protocol and supports QoS-based routing, this particular aspect of PNNI is not discussed in this chapter because most LANE networks are based on the best-effort traffic category.

The LightStream 1010 ATM switch supports some PNNI-related features that can be useful in scaling LANE networks:

- To load balance call setup requests across multiple paths between two end stations
- To load balance call setups across multiple parallel links
- To support link and path redundancy with fast convergence
- To provide excellent call setup performance across multiple hops using the background routing feature

Figure 8-17 shows how the LightStream 1010 switch supports load balancing.

Figure 8-17 Load balancing calls across multiple paths and multiple links.



As Figure 8-17 shows, load balancing of calls is enabled by default on the LightStream 1010 switch. Background routing, however, is not enabled by default. Background routing can be thought of as routing of call setups using a path from a precomputed route database. The background routing process computes a list of all possible paths to all destinations across all the service categories (for example, constant bit rate [CBR], virtual bit rate-real time [VBR-RT], virtual bit rate and nonreal time [VBR-NRT] and available bit rate-unspecified bit rate [ABR-UBR]).

When a call is placed from Point A to Point B, PNNI picks a cached route from the background route table instead of computing a route on demand. This eases the CPU load and provides a faster rate of processing the call setups.

Background routing can be useful in networks that have a stable topology with respect to QoS. It is, however, not very effective in networks that have rapidly changing topologies (for example, Internet Service Providers [ISP] networks or carrier networks). Campus LANE networks can use this feature effectively because all the SVCs in the network belong to the UBR or ABR category. To enable this feature, use the following command:

```
atm router pnni
  node 1 level 56
  bg-routes
```

The current implementation of PNNI on the LightStream 1010 switch is full, ATM Forum-PNNI Version 1 compliant. The LightStream default PNNI image license supports a single level of hierarchy, where multiple peer groups can be interconnected by IISP or by other switches that support full PNNI hierarchy; extra PNNI image license will support multiple levels of routing hierarchy.

The PNNI protocols have been designed to scale across all sizes of ATM networks, from small campus networks of a handful of switches, to the possible global ATM Internet of millions of switches. This level of scalability is greater than that of any existing routing protocol, and requires very significant complexity in the PNNI protocol. Specifically, such scalability mandates the support of multiple levels of routing hierarchy based upon the use of prefixes of the 20-byte ATM address space. The lowest level of the PNNI routing hierarchy consists of a single peer group within which all switches flood all reachability and QoS metrics to one another. This is analogous, for instance, to a single area in the OSPF protocol.

Subsequently, multiple peer groups at one level of the hierarchy are aggregated into higher-level peer groups, within which each lower-level peer group is represented by a single *peer group leader*, and so on iteratively up the PNNI hierarchy. Each level of the hierarchy is identified by a prefix of the ATM address space, implying that PNNI could theoretically contain over 100 levels of routing hierarchy. However, a handful of levels would be adequate for any conceivable network. The price to be paid for such scalability is the need for highly complex mechanisms for supporting and bringing up the multiple levels of hierarchy and for electing the peer group leaders within each peer group at each level.

Scaling an ELAN—Spanning-Tree Protocol Issues

Spanning-Tree Protocol is implemented in Layer 2 switches/bridges to prevent temporary loops in networks with redundant links. Because a LEC essentially bridges Ethernet/Token Ring traffic over an ATM backbone, the Spanning-Tree Bridge Protocol Data Units (BPDUs) are transmitted over the entire ELAN. The ATM network appears as a shared Ethernet/Token Ring network to the spanning-tree process at the edge of the Layer 2 switches.

The spanning-tree topology of a LANE-based network is substantially simpler than a pure frame-switched network that employs the Spanning-Tree Protocol. It follows that spanning-tree convergence times, which can be a major issue in large frame-switched networks, can be less of an issue in LANE networks. Note that Spanning Tree must reconverge if there are failures at the edge devices or inside the ATM network. If there is a need to tune the convergence time to a lower or higher value, the forward delay parameter can be used.

LANE Redundancy

Although LANE allows network designers to connect their legacy LANs to an ATM network, LANE Version 1.0 does not define mechanisms for building redundancy and fault tolerance into the LANE services. Consequently, this makes the LANE services a single point of failure. Moreover, router redundancy and path/link redundancy are also issues that the network designer needs to consider.

Network designers can use the following techniques to build fault-tolerant and resilient LANE networks:

- Simple Server Replication Protocol (SSRP) for LANE Services redundancy that works with Cisco and any third-party LECs.
- Hot Standby Router Protocol (HSRP) over LANE provides redundancy for the default router configured at IP end stations.

- Dual PHY LANE card on the Catalyst 5000 switch, or multiple ATM uplinks on the Catalyst 3000 switch.
- Spanning-Tree Protocol on the Ethernet-ATM switches.

The following subsections examine these various mechanisms and highlights design rules and issues to consider while implementing redundant LANE networks. It begins with a discussion on SSRP that was developed to provide redundant LANE services.

Although many vendors have implemented redundant LANE services of some fashion, they violate the LANE 1.0 specification and therefore are not interoperable with other third-party implementations. SSRP, however, does not violate the LANE 1.0 specification and is interoperable with third-party LEC implementations, which is important when implementing an interoperable ATM network.

The discussion on SSRP is followed by a description of HSRP over LANE, which provides a mechanism for building router redundancy. Following this is a discussion on the Spanning-Tree Protocol and other product-specific features that can be used to build link and path redundancy into edge devices.

Issues in a LANE 1.0 Network

The main issue with a LANE 1.0 network is that only one set of LANE service components can be accessed by a LEC at any given time. This results in the following limitations:

- Only a single LECS supports all ELANs.
- There can be only one LES/BUS pair per ELAN.

A failure in any of these service components has the following impact on network operation:

- *LECS failure*—A failed LECS impacts all the ELANs under its control because it provides access control for all the ELANs under its control. Although the existing ELANs would continue to work normally (assuming only Cisco LECs), no new LEC can join any ELAN under the control of that LECS. Also, any LEC that needs to rejoin its ELAN or change its membership to another ELAN cannot because the LES cannot verify any LEC trying to join an ELAN.
- *LES/BUS failure*—The LES/BUS pair is needed to maintain an operational ELAN. The LES provides the LE_ARP service for ATM-MAC address mappings and the BUS provides broadcast and unknown services for a given ELAN. Therefore, a failure of either the LES or the BUS immediately affects normal communication on the ELAN. However, a LES/BUS failure impacts only the ELAN served by that pair.

It is clear that these issues can be limiting to networks where resiliency and robustness is a requirement and might even be a deciding factor in your design of whether to implement LANE-based ATM networks. In addition, there are other design considerations such as the placement of the LANE service components within an ATM network that can have implications on the overall robustness of the LANE environment.

Resiliency in LANE 1.0 Networks

Increasing the resiliency of a LANE-based network essentially includes delivering increased robustness in the LANE service components such as the LECS, LES, and BUS. Such robustness is provided by SSRP through a primary-secondary combination of the LANE services. For LECS redundancy, one primary LECS is backed up by multiple secondary LECSs. LES/BUS redundancy is also handled in a similar fashion where one primary LES/BUS pair is backed up by multiple secondaries. Note that the LES/BUS functions are always co-located in a Cisco implementation and the pair is handled as one unit with respect to redundancy.

LECS Redundancy

In the LANE 1.0 specification, the first step for a LEC during initialization is to connect with the LECS to obtain the LES ATM address for the ELAN it wants to join. In order for the LEC to connect to the LECS, multiple mechanisms are defined. The first mechanism that a LEC should use is to query the ATM switch it is attached to for the LECS address. This address discovery process is done using the ILMI protocol on VPI, VCI - 0, 16.

The following is an example of the configuration command to add a LECS address to a LightStream 1010 switch:

```
atm lecs-address <LECS NSAP address> <index>
```

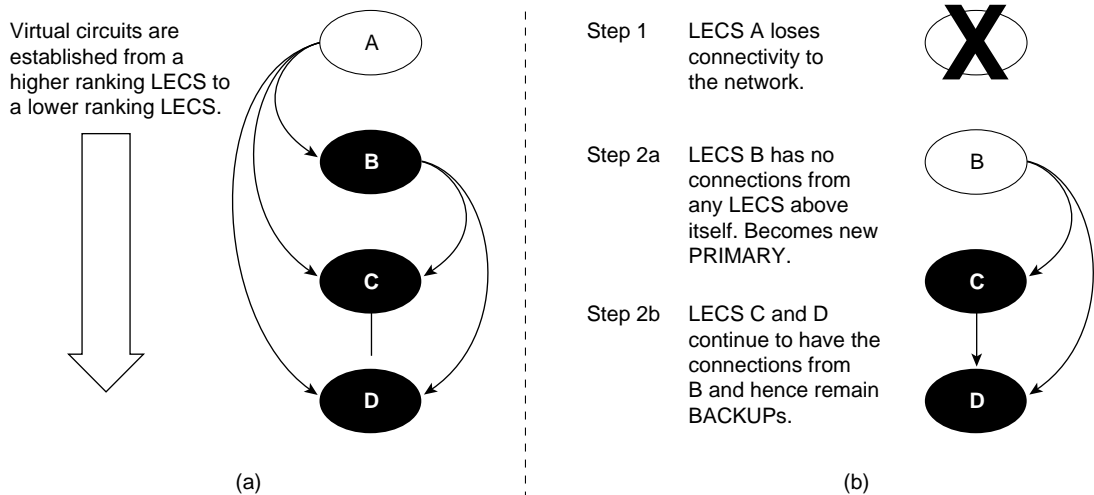
With SSRP, multiple LECS addresses are configured into the ATM switches. An LEC, which requests the LECS address from the ATM switch, gets the entire table of LECS addresses in response. The behavior of the LEC should be to attempt to connect to the highest ranking LECS address. If this fails, it should try the next one in the list and so on until it connects to the LECS.

Whereas the LEC always tries to connect to the highest ranking LECS available, SSRP ensures that there is only a single primary that responds to the Configure Request queries coming from the LEC. The establishment of a primary LECS and placing the others in backup goes to the heart of SSRP. The following describes the mechanism used by SSRP to establish a primary LECS. Upon initialization, a LECS obtains the LECS address table from the switch. The LECS then tries to connect to all the LECSs that are below itself in rank. The rank is derived from the index entry in the LECS address table.

If a LECS has a connection (VCC) from a LECS whose rank is higher than its own, it is in backup mode. The highest ranking LECS does not have any other LECS that connect to it from above and assumes the role of the primary LECS.

Figure 8-18 shows the procedure of a backup taking over in the case of a failed primary LECS. The LANE network shown in Figure 8-18 has four LECS entities (LECS A, B, C, and D). All the ATM switches in the network are configured with the same LECS address table. After startup, LECS A obtains the LECS address table from the ATM switch it is attached to and finds that it has three LECSs below itself and therefore tries to connect to LECS B, C, and D. LECS B connects to LECS C and LECS D, and LECS C connects to LECS D. There is a downward establishment of VCCs. Because LECS A does not have any VCCs from above, it becomes the primary LECS.

Figure 8-18 LECS redundancy.



During normal network operation, LECS A responds to all the configure requests and the backup LECS (LECS B, C and D) do not respond to any queries. If for some reason the primary LECS (LECS A) fails due to such conditions as a box failure, LECS B loses its VCC from LECS A as do the other LECS.

At this point, LECS B does not have any VCCs from above and therefore is now the highest ranking available LECS in the network. LECS B now becomes the primary LECS. LECS C and LECS D still have connections from higher ranking LECSs and therefore continue to operate in backup mode as shown in Step 2b of Figure 8-18.

LES/BUS Redundancy

The LES/BUS redundancy portion of SSRP supports the configuration of multiple LES/BUS pairs that work in a primary-secondary fashion. However, the mechanisms used here are different from those used for the LECS redundancy described in the preceding section.

Multiple LES/BUS pairs for a given ELAN are first configured into the LECS database. Within this database, each LES/BUS pair is assigned a priority. After initialization, each LES/BUS opens a VCC with the primary LECS using the LECS address discovery mechanism. The LES/BUS pair with the highest priority that has an open VCC to the LECS is assigned as the primary LES/BUS by the primary LECS.

SSRP Usage Guidelines

There are no theoretical limits on the number of LECSs that can be configured using SSRP, however a recommended number is two (one primary plus one backup) or three LECSs (one primary plus two backups). Any more redundancy should be implemented only after very careful consideration because it will add a significant amount of complexity to the network. This added complexity can result in a substantial increase in the amount of time required to manage and troubleshoot such networks.

SSRP Configuration Guidelines

To support the LECS redundancy scheme, you must adhere to the following configuration rules. Failure to do so will result in improper operation of SSRP and a malfunctioning network.

- Each LECS must maintain the same database of ELANs. Therefore, you must maintain the same ELAN database across all the LECSs.
- You must configure the LECS addresses in the LECS address table in the same order on each ATM switch in the network.
- When using SSRP with the Well Known Address, do not place two LECSs on the same ATM switch. If you place two LECs on the same ATM switch, only one LECS can register the Well Known Address with the ATM switch (through ILMI) and this can cause problems during initialization.

SSRP Interoperability Notes

SSRP can be used with independent third-party LECs if they use ILMI for LECS address discovery and can appropriately handle multiple LECS addresses returned by the ATM switch. For example, the LEC should step through connecting to the list of LECS addresses returned by the ATM switch. The first LECS that responds to the configuration request is the master LECS.

Behavior of SSRP with the Well Known LECS Address

SSRP also works with LECS Well Known Address (47.0079....) defined in the LANE 1.0 specification. The Cisco LECS can listen on multiple ATM addresses at the same time. Therefore, it can listen on the Well Known Address and the auto-configured ATM address, which can be displayed using the **show lane default** command.

When the LECS is enabled to listen on the Well Known Address, it registers the Well Known Address with the ATM switch so that the ATM switches can advertise routes to the Well Known Address and route any call setups requests to the correct place.

Under SSRP, there are multiple LECSs in the network. If each LECS registers the Well Known Address to the ATM switches that it is connected to, call setups are routed to different places in the network. Consequently, under SSRP you must configure an autoconfigured address so that the negotiation of the master first takes place and then the master registers the Well Known Address with the ATM switch. If the master fails, the Well Known Address moves with the master LECS. The PNNI code on the LightStream 1010 switch takes care of advertising the new route to the Well Known Address when there is a change of LECS mastership. Therefore, third-party LECs that use only the Well Known Address can also interoperate with SSRP. SSRP is the only redundancy scheme that can be used with almost any LEC in the industry.

To implement SSRP with the Well Known Address, use the following steps:

- Step 1** Configure the LECS to listen on the autoconfigured address (or if you want a separate ATM address that you have predetermined). This autoconfigured (or other) address should be programmed into the ATM switches for the LECS address discovery mechanism.
- Step 2** Configure each LECS to listen on the Well Known address using the **lane config fixed-config-atm-address** command. After the master LECS is determined using the LECS redundancy procedure, the master registers the Well Known Address to the ATM switch.

Note SSRP with the Well Known Address does not work properly under certain circumstances (during failover) if two LECS are attached to the same ATM switch. This is due to the possibility of duplicate address registration on the same switch, which ILMI does not allow. Make sure each LECS is on a separate ATM switch.

Behavior of SSRP in Network Partitions

In the event of network partitions where two separate ATM clouds are formed due to an interconnecting link or switch failure, each cloud has its own set of LANE services if SSRP is configured to handle network partitions.

When configuring SSRP, use the following guidelines to accommodate the possibility of network partition:

- Configure each partition with its own LANE services that can become active during a network partition. For example, if you are connecting two sites or campuses across a MAN and you want the same ELANs at both locations, configure each campus/site with its own LANE services.
- Routing behavior should be carefully examined during a network partition in the case where an ELAN maps to a Layer 3 network (for example, an IP subnet or IPX network) because there are now two routes to the same subnet (assuming there are redundant routers in the network). If there are no redundant routers, one of the partitions will be effectively isolated from the rest of the network. Intra-ELAN traffic will continue to behave properly.

HSRP over LANE

HSRP is a protocol that network designers can use to guard against router failures in the network. The HSRP protocol is exchanged between two routers and one of them is elected as the primary router interface (or subinterface) for a given subnet. The other router acts as the *hot standby* router.

In HSRP, a default IP address and a default MAC address are shared between the two routers exchanging the HSRP protocol. This default IP address is used as the default gateway at all IP end stations for them to communicate with end stations outside their immediate subnet. Therefore, when there is a primary router failure, the hot standby router takes over the default gateway address and the MAC address so that the end station can continue communicating with end stations that are not in their immediate subnet.

Because HSRP is a Layer 2 mechanism and needs a MAC address-based Layer 2 network, it is possible to implement HSRP style recovery over LANE. The mechanisms used are the same as for any Ethernet interface and can be configured at a subinterface level.

Redundant ATM Port Card for the Catalyst 5000

Another aspect of addressing the redundancy needs from a physical network perspective is the addition of a redundant PHY portion of an ATM card. The Catalyst 5000 switch employs the dual PHY redundant ATM card. This redundancy is only at a physical level and is useful in cases where the primary link to the ATM switch goes down.

Role of Stratm Technology

Stratm Technology is a new approach to ATM switching technology that incorporates patented standards-based Cisco technology into custom silicon. These application-specific integrated circuits (ASICs) dramatically increase ATM efficiency and scalability and significantly lower the absolute cost of delivering ATM solutions. Stratm Technology can be implemented in switches and routers across LANs, campus networks, and WANs, enabling the delivery of high-performance, end-to-end ATM services to meet a wide range of needs.

Benefits of Stratm Technology

The benefits of Stratm Technology include the following:

- Dramatic improvement in network price/performance scalability
- Increased application *goodput*
- Protection of technology investments
- Increased portability
- Guaranteed infrastructure

Each of these benefits is described in more detail in the following sections.

Improved Network Price/Performance Scalability

Stratm Technology features can dramatically improve network price/performance and scalability as follows:

- Support of up to eight OC-3 (155-Mbps) port interfaces per card slot, and up to 12 digital signal Level 3 T3/E3 (45-Mbps) port interfaces per card slot
- A 30 percent increase in SVC completions to more than 4,000 per second per node

- An increase in connection density per switch by 500 percent
- An increase in the buffering capability of each card to 200,000 cells per card, upgradable to nearly one million cells
- A reduction in the price per port for high-speed connections by up to 50 percent
- The ability to support per-virtual-connection control queuing, rate scheduling, statistics collection, and fair sharing of network resources on an individual connection basis

Increased Application Goodput

Intelligent ATM features are embodied in Stratm Technology. These features are designed to increase application goodput dramatically through advanced features, which are distributed throughout the BXM module in silicon.

- *Distributed ATM functions*—Stratm distributes such ATM services as traffic management, per-VC queuing, class of service (COS) management, SVCs, and multicasting to each card on a silicon chip. Distributed functionality ensures faster, more efficient processing, and it eliminates the possibility of a single point of failure disrupting the entire network.
- *Highest bandwidth efficiency*—Stratm delivers guaranteed bandwidth on demand, QoS, and fair sharing of network resources to each individual connection. With fast, efficient processing and guaranteed bandwidth, application performance is significantly enhanced.
- *Advanced traffic management capabilities*—Stratm incorporates the industry's first commercially available Virtual Source/Virtual Destination (VS/VD) implementation of the full ATM Forum's Traffic Management Specification Version 4.0. This ensures the highest efficiency in bandwidth utilization and provides support for the multicasting capabilities required to successfully deliver multimedia and switched internetworking services.
- *End-to-end intelligence*—With VS/VD implementation, Stratm also represents the industry's first complete LAN-to-WAN ARB implementation. This feature enables ATM services to be delivered to the desktop, ensuring high performance for the most demanding applications.

Industry-Leading Investment Protection

Stratm allows you to protect your current investments by integrating with today's network infrastructures, and providing advanced features and functionality to protect investments far into the future. You can protect your technology investment because of the following Stratm capabilities:

- *Seamlessly integrates with existing switches*—Stratm Technology integrates into Cisco's ATM switching platforms, allowing you to enhance your investment in Cisco technology.
- *Delivers unparalleled performance*—Current ATM switching platforms deliver performance that enables end-to-end delivery of high-quality, high-performance network services.
- *Delivers the future*—Stratm Technology extends the features and functionality of current switches to support next generation requirements. With this technology, you can easily deliver multiple services from a single network infrastructure and ensure the highest QoS possible.


Increases Portability and Guarantees an Infrastructure

With a modular chip set, Stratm increases the portability of standards-based ATM. ATM in silicon stabilizes the transport layer of networks, thereby guaranteeing the necessary infrastructure for efficient, high-performance delivery of emerging multimedia and Internet-based applications.

Cisco ATM WAN Products

As Figure 8-19 shows, Cisco provides end-to-end network ATM solutions for internetworks.

Figure 8-19 End-to-end network solutions.



The Cisco ATM products suited for WAN deployment include the following:

- Cisco/StrataCom IGX switch, which is well suited for deployment in an enterprise WAN environment
- Cisco/StrataCom BPX/AXIS switch, which meets the needs of high-end, enterprise WAN and service provider environments
- Cisco AIP for the Cisco 7500 and 7000 series of routers
- Cisco ATM Network Interface Module (NIM) for the Cisco 4700 and 4500 series of routers
- Cisco edge devices such as the Catalyst 5000 and Catalyst 3000 switches, which connect legacy LANs with an ATM network

Note The LightStream 1010 is a Cisco campus ATM switch that is specifically designed for workgroup and campus backbone deployment. However, it can also meet the needs of a low-end enterprise environment. For more information on the LightStream 1010 switch as a workgroup switch, see Chapter 12, “Designing Switched LAN Internetworks.”

Stratm-Based Cisco WAN Products

Stratm Technology is the basis for a new class of ATM WAN switch products. These products are designed to take users to the next level in building the world’s most efficient and scalable ATM networks. High-speed, high-density products based on Stratm Technology provide advanced features, such as the following:

- Standards-based traffic management
- Fair sharing of bandwidth
- Unmatched port density and switch scalability
- High-performance SVCs
- Multicast capability


Cisco/StrataCom BPX

The Cisco/StrataCom BPX Service Node is a standards-based, multiservice ATM switch designed to deliver the highest levels of network scalability, flexibility, and efficiency. The BPX achieves multiservice functionality, efficient use of bandwidth, high performance for all users, and guaranteed QoS for all traffic types through its advanced traffic management features. These advanced traffic management capabilities are based on the first fully compliant implementation of the ATM Forum’s Traffic Management Specification V. 4.0, as well as the International Telecommunications Union (ITU) Recommendations I.371 and I.35B.

The BPX incorporates Stratm Technology, which is implemented in custom silicon ASICs. Stratm distributes advanced ATM capabilities throughout the switch modules, resulting in unmatched port density, support for hundreds of thousands of connections, and new functionality. Advanced traffic management features, together with an optimized hardware architecture, enable the switch to simultaneously support ATM, Frame Relay, Internet, voice, wireless communication, video, switched internetworking, and circuit emulation services.

The BPX also offers operational ease. With the BPX’s 20-Gbps capacity of high-throughput, low-latency switching, and support for multiple classes of service, service providers can deliver innovative revenue-generating data, voice, and video services. Large enterprises can combine LAN, Systems Network Architecture (SNA), voice, and other types of traffic over a single WAN backbone, as shown in Figure 8-20. The BPX enables organizations to migrate to a new generation of ATM networks and complement existing investments in routers and Frame Relay switches.

Figure 8-20 BXP multiservice platform.



As Table 8-4 indicates, Stratm allows the BPX to deliver high application performance and guaranteed network responsiveness for all users.

Table 8-4 BXP Metrics with Stratm (Continued)

Category	Amount
----------	--------

Ports/Node	
• DS3/E3	144
• OC-3/STM-1	96
• OC-12/STM-4	24
SVCs/Node	
• Active	384,000
• Calls/Second	4,000
Buffers/Node	22,000,000 cells
Buffers/Card	900,000 cells
Nodes/Network	
• Peer group nodes	100
• Number of peer groups	Unlimited

The BPX includes high-density, Broadband Switch Module (BXM) cards that provide standard interfaces for connecting to cell-based customer premises equipment via ATM UNI or to non-Cisco networks via NNI.

Cisco/StrataCom BXM Switch Modules

The Stratm-based BXM cards are a family of highly configurable interface modules that extend today's central crosspoint ATM switch architecture to highly scalable distributed architectures. By integrating Stratm into its ATM platforms, Cisco delivers dramatically increased connections and port density as well as new features and functionality. Equipped with these Stratm-based products, network designers can deploy the most efficient, scalable ATM networks possible.

The major functions of the BXM modules include the following:

- Available Bit Rate Engine (ABRE)
- Serial interface and multicast buffer subsystem (SIMBS)
- Routing, control, monitoring, and policing (RCMP)
- SONET/Synchronous Digital Hierarchy (SDH) UNI (SUNI)

Table 8-5 provides a summary of the BXM switch modules.

Table 8-5 BXM Family of Switch Modules

BXM Switch Module	Description
BXM-T3/E3 Broadband Switch Module	A 6- or 12-port, Digital Signal 3 (DS3) (45 or 34 Mbps) ATM interface card, which supports E3/DS3 native ATM access and trunk ports, for the BPX switch. The interface can be configured for trunk, public, or private UNI applications on a per-port basis to provide a high-density, low-cost, broadband ATM networking solution.

BXM-155 Broadband Switch Module	An OC-3c/STM-1 version of the BXM interface card, which supports OC-3/STM-1 native ATM access and trunk ports for the BPX switch. It operates at the SONET/SDH rate of 155.520 Mbps. The card provides four or eight OC-3/STM-1 ATM ports per card, each of which can be configured for either trunk or access application.
BXM-622 Broadband Switch Module	An OC-12c/STM-4 version of the BXM interface card, which supports OC-12/STM-4 native ATM access and trunk ports for the BPX switch. It operates at the SONET/SDH rate of 622.08 Mbps. One- or two-port versions of the card are available; both can be configured for either trunk or access application.

The BXM cards support ATM-Frame Relay internetworking and service internetworking. They also allow you to configure PVCs or SVCs for the following defined service classes:

- Constant Bit Rate (CBR)
- Variable Bit Rate-Real Time (VBR-RT)
- Variable Bit Rate-Non-Real Time (VBR-NRT)
- Unspecified Bit (UBR)
- Available Bit Rate (ABR)

The BPX with Stratm architecture supports up to 16 independent classes of service, thereby protecting your hardware investment as the industry defines additional traffic types.

AXIS Interface Shelf

The AXIS interface shelf enables the BPX Service Node to support a wide range of user services. AXIS modules adapt incoming data to 53-byte ATM cells using industry-standard ATM adaptation layers (AALs) for transport over the ATM network.

Because the AXIS interface shelf will support a range of services from a single platform, organizations can reduce equipment costs, fully utilize their investments in existing premises equipment, and rapidly deploy new services as required.

Services below 34 Mbps are provisioned on the AXIS shelf, and the following interfaces are supported:

- Frame Relay
- High-speed Frame Relay
- ATM Frame UNI
- SMDS
- T1/E1 ATM UNI
- n x T1/E1 inverse multiplexing for ATM (IMATM) UNI
- Circuit emulation
- ISDN switched access

Each AXIS shelf aggregates traffic from as many as 80 T1 or E1 ports onto a single port of the multiport broadband interface card. This high-port density maximizes use of the BPX high-capacity switch fabric. A compact footprint minimizes the space required within central offices. Each 9-inch, rack-mounted shelf supports more than 2,000 64-Kbps users.

Cisco/StrataCom IGX Family of Switches

For wide-area networking, LAN data flowing between different enterprise sites is aggregated by the router and then mixed with voice and other legacy data streams across the corporate wide-area backbone. Traditionally, these corporate backbones use TDM technology. However, as the use of LAN data has exploded and older TDM equipment has been fully depreciated, newer solutions can be cost justified. Enterprises are increasingly turning to a new public service officers (for example, VPN, Frame Relay, and intranets) and a new generation of Frame Relay/ATM-based enterprise switches to maximize the efficiency and minimize the cost of their networks.

The Cisco/StrataCom IGX family of switches provides the needed linkage to integrate the high-speed LAN data and the lower-speed voice and legacy data across the enterprise backbone in the most cost-effective manner. The IGX family of switches is specifically designed for enterprise integration.

The IGX family of ATM enterprise WAN switches includes the IGX8 (8-slot switch), the IGX16 (16-slot switch), and IGX32 (32-slot switch). The IGX family can provide the following enterprise WAN support:

- *Voice*—UVM and CVM
- *Legacy data*—HDM and LDM
- *ATM*—ALM
- *Frame Relay*—UFM and FRM
- *Trunks*—NTM, BTM, and ALM

Benefits of the IGX

With the IGX switch, you can leverage ATM to save costs as follows:

- Apply utilization rates in your network design to source PVCs
- Combine multiple networks into one multiservice network
- Optimize the transmission network with design tools

For example, you can use StrataView+, a network management tool, for network discovery. You can also use the Configuration Extraction Tool (CET) to populate the design data set with existing facilities. With such design tools, incremental network design is possible.

In addition to lower costs of networks, other major benefits of deploying the IGX in an enterprise internetwork include the following:

- Multiband/multiservice
- Better application performance
- Reliability
- Investment protection

Sample IGX Configuration


This section provides an example of how IGX switches can be deployed in an enterprise internetwork. In this example, a postal service has 180,000 employees and 180 mail-sorting offices. The current network design, which is a TDM network, has 750 LANs, 350 routers, 220 X.25 switches, and 110 PBXs. The network handles approximately 70 million voice minutes of traffic per year.

Currently, the enterprise is confronted with the following problems with the existing network design and network requirements:

- Poor performance on the existing TDM network
- Exponential growth of LAN traffic
- Many new applications
- Inability of the existing WAN to scale up

Figure 8-21 shows an example of how the IGX switches can be deployed throughout the network to address these problems.

Figure 8-21 Example of an IGX deployment.



By deploying the IGX switches throughout the enterprise internetwork, the following benefits are obtained:

- Integration of the voice and data networks
- Improved performance for each type of traffic
- Better response times for new applications
- Reduced downtime
- Higher bandwidth utilization (fivefold increase in traffic using existing trunks)
- Implementation of a scalable network that supports rapid deployment of new services
- Simplification of network design with a reduction in management costs

Cisco ATM LANE Products

Cisco offers a complete ATM LANE solution by providing the following:

- Inter-ATM ELAN communication through routing
- LEC/BUS/LECS/LES on Cisco 7500, 7000, 4500, and 4000 series routers
- LEC/BUS/LECS/LES on the Catalyst 5000 and Catalyst 3000 switches
- Cisco LightStream 1010 ATM switch

Cisco 7500 Router Series

Data center consolidation, client-server architectures using centralized servers, and growth in remote sites all drive the rapidly growing need for WAN bandwidth. High-performance routing provides critical functionality in the high-speed WAN environment.

The Cisco 7500 router series extends the capabilities of the Cisco 7000 family and incorporates distributed switching functions. The distributed switching capability allows network designers to provide the high-performance routing necessary to support networks using ATM, multilayer LAN switching, and VLAN technologies.

The Cisco 7500 family of routers offers a broad support of high-speed ATM and WAN interfaces. The higher port densities supported the Cisco 7500 series easily handles the large number of interfaces that result from more remote site connectivity. The Cisco IOS software adaptive rerouting increases network availability, and its flexible interfaces provide support for multiple services and a migration path to ATM. Network designers can deploy the Cisco 7500 series in the WAN environment to access multiple types of carrier service offerings as they migrate from TDM backbones to ATM backbones. The Cisco 7500 series also provides network security and minimizes the loss of transparency. The Cisco 7500 series running Cisco IOS Software Release 11.0 or later provides tools for network configuration, fault detection, and minimizing unnecessary traffic across expensive wide-area links.

Note The features discussed for the Cisco 7000 series are also applicable to the Cisco 7500 series.

Cisco 7000 Series

With a CxBus card, the AIP can be installed in a Cisco 7000 and Cisco 7500 series routers and is compatible with all the interface processors as well as the Route Processor (RP), the Switch Processor (SP), the Silicon Switch Processor (SSP), and the new Route Switch Processor (RSP). The AIP supports the following features:

- Single, native ATM port with transmission rates up to 155 Mbps over a variety of ATM physical layer interface media modules (PLIM), eliminating the need for an external ATM data service unit (DSU).
- Multiprotocol support over ATM for all the popular network operating systems and the internet protocol: IP, AppleTalk, Novell IPX, DECnet, Banyan VINES, XNS, and OSI CLNS.
- ATM Adaptation Layers (AALs) 3/4 and 5.
- Dual RISC and dual-SAR design for high-speed cell and packet processing.
- Interim Local Management Interface (ILMI) for ATM address acquisition/registration.

Note Cisco IOS Software Release 10.0 supports AAL5 PVCs only.

Cisco IOS Software Release 10.0 and later support 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 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 multimode 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 through all the AIPs configured in a router should be limited to 200 Mbps full duplex. For that reason, only the following combinations are supported:

- Two TAXI interfaces
- One SONET and one E3 interface
- Two SONET interfaces, one of which is lightly used
- Five E3 interfaces

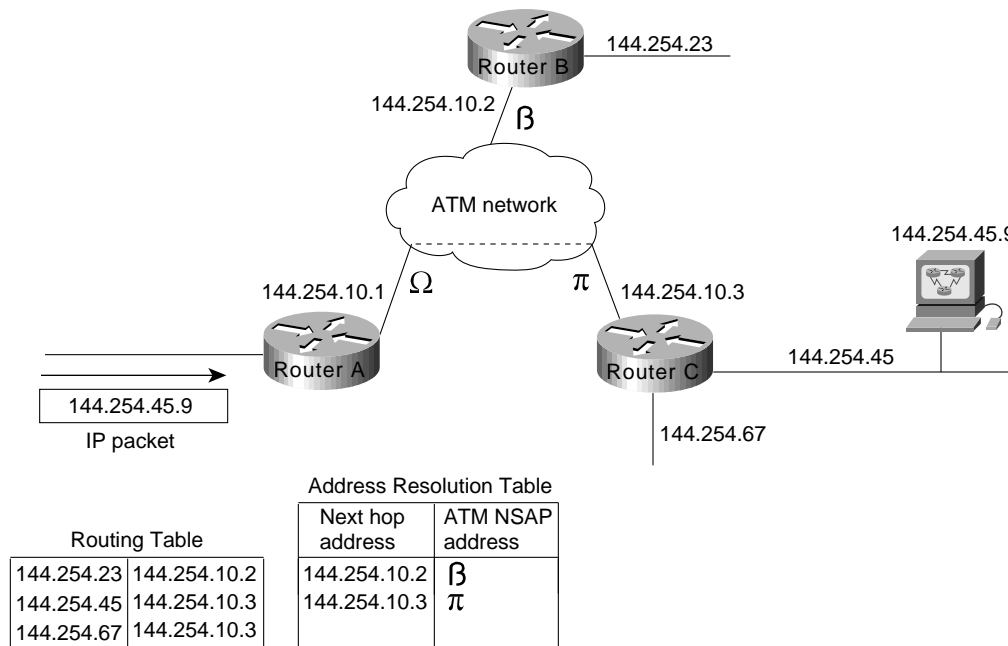
The AIP includes hardware support for various traffic-shaping functions. Virtual circuits can be assigned to one of eight rate queues, each of which is programmable for a different peak rate. Each virtual circuit can be assigned an average rate and specific burst size. The signaling request specifies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

The 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 8-22 shows how the routing table and address resolution table on Router A are used to forward data to a workstation behind Router C.

Figure 8-22 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 8-23 shows the layers through which the packet travels.

Figure 8-23 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 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**—Sets 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**—Sets 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**—Sets 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 4500/4700 ATM NIM

The NIM is the midrange ATM router interface for the Cisco 4500 and Cisco 4700 series of routers. The function of this ATM module is internally much different than that of the AIP module. With the Cisco 4500 and 4700 series router, the packet memory is kept in a 4-Mbps pool that is shared by all of the NIMs. The Cisco IOS software also runs on these routers so the same ATM functionality and commands work on both the Cisco AIP and NIM. The performance of the NIM is actually better in the NIM for process and fast-switched protocols, but the Autonomous/SSE switching available in the Cisco 7000 series, and fast switching available on the Cisco 7500 series remain the fastest in the product family.

In this regard, network designers can deploy the Cisco 4700 series of routers to offer LANE services because the BUS is in the fast switching path. It is important to note that the NIM supports 1024 VCCs and this should be taken into consideration in SVC-intensive LANE networks.

ATM Edge Devices

The Catalyst 5000, 5500, and Catalyst 3000 switches are LAN switches that have ATM uplinks. Consequently, network designers can use these switches as edge devices to interconnect legacy LANs to an ATM network.

Catalyst 5000 as an ATM Edge Device

The Catalyst 5000 ATM LANE Dual PHY modules integrates high-speed, switched LANs across an ATM campus network providing legacy LANs with access to ATM-based services in the backbone. The ATM module supports two (one primary and one secondary) 155-Mbps OC-3c interfaces with a wide range of media options (for example, single-mode fiber, multimode fiber, and unshielded twisted pair [UTP Category 5]).

A maximum of three ATM LANE modules can be supported simultaneously in one Catalyst 5000 switch to provide redundant, fault-tolerant connections. This module delivers redundant LANE services through Cisco's LANE SSRP.

The Catalyst 5000 ATM module is designed to provide Ethernet to ATM functionality by acting as a LANE client. The BUS functionality on the Catalyst 5000 ATM card was designed for very high performance. The data path for the BUS is implemented entirely in firmware/hardware.

Catalyst 3000 as an ATM Edge Device

The Catalyst 3000 switch can also function as an ATM LANE edge device. Like the Catalyst 5000 switch, the Catalyst 3000 switch also supports an ATM LANE module. The Catalyst 3000 ATM module supports a 155-Mbps OC-3c multimode optical interface that is compliant with the ATM Forum UNI 3.0 and UNI 3.1 specifications. The ATM module in conjunction with other Catalyst 3000 modules can also be used to connect Fast Ethernet hubs, switches, and routers to the ATM backbone.

Support for Cisco's VLAN Trunking Protocol (VTP) allows multiple Catalyst 3000 and Catalyst 5000 switches within a network to share ELAN or VLAN configuration information. For example, VTP will automatically map VLANs based upon Fast Ethernet trunks (ISL) to ELANs based upon ATM trunks.


LightStream 1010 ATM Switches

The LightStream 1010 family of switches are modular switches designed for campuses or workgroups depending upon the types of interfaces used. Its central processor is dedicated to a single, field replaceable ATM Switch/Processor module (ASP).

Single-Switch Designs

Because ATM can use existing multimode fiber 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 1010 switch, as shown in Figure 8-24. In this topology, each router has a 155-Mbps point-to-point connection to every other router on the ring.

Figure 8-24 Parallel FDDI and ATM backbone.




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 8-25, which can be used to migrate ATM slowly throughout the network, using FDDI as a backup.


Figure 8-25 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 8-26.

Figure 8-26 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 8-27. 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 8-27 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 8-27, 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.

