ATM and Token Ring LANE

The Catalyst 3900 and the Catalyst 5000 have ATM expansion modules that provide high-speed connectivity between the switch and an ATM backbone network.

This chapter provides the following information:

- Understanding ATM
- Understanding LAN Emulation
- Recommended Environments

Understanding ATM

ATM is a cell-switching and multiplexing technology that combines the benefits of circuit switching (constant transmission delay and guaranteed capacity) with those of packet switching (flexibility and efficiency for intermittent traffic). Like X.25 and Frame Relay, ATM defines the interface between the user equipment (such was workstations and routers) and the network (referred to as the User-Network Interface [UNI]). This definition supports the use of ATM switches (and ATM switching techniques) within both public and private networks.

Because it is an asynchronous mechanism, ATM differs from synchronous transfer mode methods, where time-division multiplexing (TDM) techniques are employed to preassign users to time slots. ATM time slots are made available on demand, with information identifying the source of the transmission contained in the header of each ATM cell. TDM is inefficient relative to ATM because if a station has nothing to transmit when its time slot comes up, that time slot is wasted. The reverse situation, where one station has large amounts of information to transmit, is also less efficient. In this case, that station can only transmit when its turn comes up, even though all the other time slots are empty. With ATM, a station can send cells whenever necessary.

Another critical ATM design characteristic is its star topology. The ATM switch acts as a hub in the ATM network, with all devices attached directly. This provides all the traditional benefits of star-topology networks, including easier troubleshooting and support for network configuration changes and additions.

Furthermore, ATM's switching fabric provides additive bandwidth. As long as the switch can handle the aggregate cell transfer rate, additional connections to the switch can be made. The total bandwidth of the system increases accordingly. If a switch can pass cells among all its interfaces at the full rate of all interfaces, it is described as nonblocking. For example, an ATM switch with 16 ports set at 155 Mbps would require about 2.5 Gbps aggregate throughput to be nonblocking.

ATM switches transmit data in small units called *cells*. The latency in a cell switch is very small because of the short cell size. Short cells have a tiny store-and-forward delay. In the absence of port contention and buffering, cells are switched quickly in hardware. For information about the format of an ATM cell, see the “Frame Formats”appendix.

In addition to the low latency, ATM is beneficial to large networks because it:

- Is a multiplexing and switching technology designed for flexibility and performance.
- Supports Quality of Service (QOS) options for flexibility and high bandwidth options (up to gigabits per second) for performance.
- Offers both permanent virtual circuits (PVCs) that are set up for static connections and switched virtual circuits (SVCs) that are automatically set up and torn down when data needs to be transferred.
- Supports environments where applications with different performance requirements need to be executed on the same computer, multiplexer, router, switch and network. The flexibility of ATM means that voice, video, data, and future payloads can be transported.
- Has worldwide support. The ATM Forum, an industry forum made up of many companies (including Cisco), works with formal standards bodies to specify ATM.

PVC

A PVC is a non-switched connection that is established beforehand (manually pre-provisioned) to satisfy a standing need for network services. It is a logical (not a physical) connection between two communicating ATM peers. This type of connection is typically established by a network administrator.

User applications that require an on-going, specific level of transmission bandwidth typically use PVCs for interconnectivity. The network bandwidth required in this type of application tends to be more predictable and constant, enabling the physical transmission medium to be tailored to an expected volume of traffic, and vice versa.

With a PVC, everything is statically configured and no signaling is involved. The PVC is mapped to a network in a subinterface point-to-point configuration. The logical data link layer can use SNAP encapsulation (as defined in RFC 1483). This allows multiple protocols to be multiplexed over one PVC. Alternatively, the logical data link layer can use LANE Version 1 over PVC.

The PVC is statically mapped at each ATM node. The path of the PVC is identified at each switch by an incoming virtual channel identifier/virtual path identifier (VCI/VPI) and an outgoing VCI/VPI.

Note: The Catalyst 3900 ATM expansion module does not support PVC configuration.

SVC

An SVC is a switched connection that is established by a defined and standardized ATM signaling protocol. This type of connection is set up dynamically (on demand) across the network, as required by the user’s communications applications. An SVC is established and torn down using a flexible connection setup protocol that supports various connection types.

The transfer of information between network users by means of SVCs typically occurs through shared network facilities, rather than through dedicated transmission lines or owned physical facilities.

Establishing an ATM SVC involves an agreement between the end nodes and all the switches in between. Each end node has a special signaling channel to the connected switch called the UNI. Switches have a signaling channel between them called the Network-to-Network Interface (NNI). Cells that arrive on the signaling channel are reassembled into frames in the reliable Service Specific Connection Oriented Protocol (SSCOP). The signaling information follows the Q.2931 standard.

Establishing an SVC potentially involves signaling between the following:

- Router and a private ATM switch (private UNI)
- Router and a public ATM switch (public UNI)
- Private ATM switch and a public ATM switch (public UNI)

The UNI is defined by the ATM Forum UNI specification.

Interfaces to public ATM networks are identified by an E.164 address. Interfaces to private ATM networks are identified by a network service access point (NSAP) address. These addresses are contained in different fields of the same 20-octet address.

Once an SVC is established, it functions like a PVC. SVCs can be used in point-to-point subinterface configuration or point-to-multipoint nonbroadcast multiaccess (NBMA) configuration.

ATM Adaptation Layers

The purpose of the ATM adaptation layer (AAL) is to receive the data from the various sources or applications and convert or adapt it to 48-byte segments that will fit into the payload of an ATM cell. The AAL segments upper-layer user information into ATM cells at the transmitting end of a virtual connection and reassembles the cells into a user-compatible format at the receiving end of the connection. These complimentary functions occur between communicating peers in the network at the same level of the ATM architectural model.

The AAL is not a network process. Rather, AAL functions are performed by the user's network terminating equipment on the user side of the UNI. Consequently, the AAL frees the network from concerns about different traffic types.

How AAL processes are carried out depends on the type of traffic to be transmitted. Different types of AALs handle different types of traffic, but all traffic is ultimately packaged by the AAL into 48-byte segments for placement into ATM cell payloads. Consequently, several different AALs have been defined for different types of services.

Table 4-1 lists these AALs.

Table 4-1 ATM Adaption Layers

Traffic Class	Timing Relationship	Connection Mode	Bit Rate	Traffic Description
Class A (AAL1)	Synchronous	Connection- oriented	Constant	This type of traffic typically consists of constant bit rate (CBR) analog signals. Hence, synchronous timing relationships exist between the senders and receivers of this traffic. This type of traffic over an ATM network is often referred to as circuit emulation service, an example of which is fixed bit rate, uncompressed voice, or video data.
Class B (AAL2)	Synchronous	Connection- oriented	Variable	As with Class A traffic, synchronous timing relationships exist between the senders and receivers of Class B traffic. However, Class B relates to variable bit rate (VBR) traffic, typical examples of which are compressed voice and video traffic. Such traffic is typically "bursty" in nature.

Table 4-1 ATM Adaption Layers (Continued)

Traffic Class	Timing Relationship	Connection Mode	Bit Rate	Traffic Description
Class C (AAL3/4)	Asynchronous	Connection- oriented	Variable	No timing relationships exist between the senders and receivers of data. Hence, such traffic is asynchronous. Class C handles VBR connection-oriented traffic. This class provides point-to-point or point-to-multipoint ATM cell relay services over connections established "on the fly" through signaling and routing messages exchanged between data senders and receivers. This service handles multiple traffic types (data, voice, and video) in which user data is arranged into ATM cells for efficient transport through the network. Class C of traffic contains sequencing bits that allows the cells to take different paths and still be reassembled in the correct order at the receiving station. This type of traffic is sensitive to cell loss, but not to cell transport delay (or latency). Latency is the delay between the time a device receives a cell on its input port and the time the cell is forwarded through its output port.
Class D (AAL5)	Asynchronous	Connectionless	Variable	Class D handles unspecified bit rate (UBR) traffic in a connectionless, asynchronous manner.

Because ATM is inherently a connection-oriented transport mechanism and because the current applications of ATM are heavily oriented toward LAN traffic, many of the current ATM products, including the Catalyst 3900 and the Catalyst 5000, support the Class D adaptation layer with AAL5.

AAL Sublayers

The AAL performs two main functions in service-specific sublayers of the AAL:

- A convergence function in the convergence sublayer (CS)

  The CS provides appropriate traffic services to higher-layer protocols. Once a connection is established between communicating ATM entities with an appropriate QOS, the CS accepts higher-layer traffic for transmission through the network. Depending on the traffic type, certain header or trailer fields are added to the user data payload and formed into information packets called CS protocol data units (CS-PDUs).
- A cell segmentation and reassembly function in the segmentation and reassembly (SAR) sublayer.

  The SAR segments each CS-PDU received from the CS into smaller units and adds a header or trailer field, depending on the traffic type, to form 48-byte payloads called SAR sublayer protocol data units (SAR-PDUs). Once the user data is arranged into SAR-PDUs by the AAL layer, they are passed to the ATM layer, which packages the data into 53-byte ATM cells, making them suitable for transport as outgoing ATM cells by the physical layer.

  Upon receipt of incoming ATM cells from the physical layer (that is, cells delivered from a peer physical layer elsewhere in the network), the AAL removes any AAL-specific information from each cell payload and reassembles the packet for presentation to higher layer protocols in a form expected by the user application.

AAL5 Traffic Processing

AAL5 has been designed to process traffic typical of today's LANs. Originally, AAL3/4 was designed to process this kind of traffic. However, the inefficiency of AAL3/4 for handling LAN traffic led to the use of AAL5 for such traffic.

AAL5 provides a streamlined data transport service that functions with less overhead than AAL3/4. AAL5 is typically associated with UBR traffic.

Another AAL5 attribute contributing to its efficiency is that is uses only 5 bytes of header. None of the payload is used for header information. Also, AAL5 calculates a 32-bit cyclic redundancy check (CRC) over the entire AAL5 protocol data unit to detect cell loss and the incorrect ordering or incorrect insertion of cells.

For purposes of AAL5 traffic processing, the CS is divided into the following parts:

- Common part convergence sublayer (CPCS)—Provides the capability to transfer the CPCS protocol data units (CPCS-PDU payloads) from one AAL5 user to another AAL5 peer in the network. The AAL5 traffic type supports both a message-mode service and a streaming mode service.
- Service specific convergence sublayer (SSCS)—Allows different SSCS protocols to be defined to support specific AAL user services or groups of services. The SSCS may also be null because it provides only for the mapping of equivalent AAL primitives to the CPCS, and vice versa.

Components of an ATM Network

The building blocks of an ATM internetwork may consist of the following:

- Routers with ATM interfaces
- Computers with a native ATM NIC
- LightStream 1010 or other ATM switches
- ATM physical layer, supporting Synchronous Optical Network (SONET) OC-3 with single or multimode fiber, Transparent Asynchronous Transmitter/Receiver Interface (TAXI) with multimode fiber, or DS3/E3 with coaxial cable
- LAN switches with ATM interfaces

ATM and VLANs

The Catalyst 3900 ATM expansion module supports up to 63 VLANs (or ELANs). Each ELAN corresponds to a TrCRF. Each association between the ATM expansion module and a TrCRF creates a virtual ATM port. A virtual ATM port is the equivalent of an LAN Emulation Client (LEC).

Understanding LAN Emulation

LANs can use connectionless service. However, ATM is always a connection-oriented service. Devices first use a signaling process to establish a path with an ATM destination. Devices can send cell-based traffic only after the devices have identifiers pointing to the connection path.

LANE uses point-to-multipoint connections to service the connectionless broadcast service that is required by LAN protocols.

Cisco’s Token Ring implementation of LANE makes an ATM interface look like one or more Token Ring interfaces. Setting up LECs allows the Catalyst 3900 or Catalyst 5000 Token Ring module to operate in an ATM LAN environment containing Cisco 7000 or Cisco 4500 series routers with ATM Interface Processor (AIP) connected to a LightStream 1010 ATM switch.

Figure 4-1 illustrates the physical layout of an ATM network that uses LANE.

Figure 4-1 Physical View of LANE

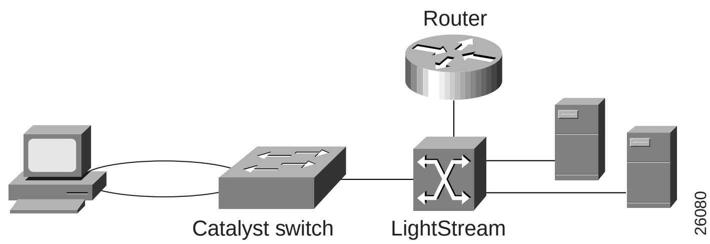

Figure 4-2 illustrates the logical view of the LANE network.

Figure 4-2 Logical View of LANE

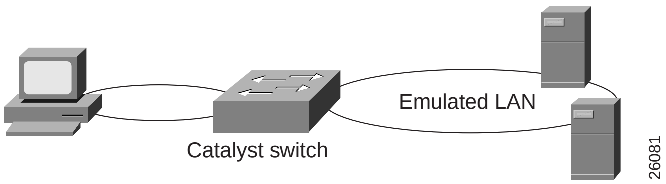

LANE is an ATM service defined by the ATM Forum specification *LAN Emulation over ATM* (ATM_FORUM 94-0035). This service emulates the following LAN-specific characteristics:

- Connectionless services
- Multicast services
- LAN MAC driver services

LANE service provides connectivity between ATM-attached devices and LAN-attached devices. This includes connectivity between ATM-attached stations and LAN-attached stations as well as connectivity between LAN-attached stations across an ATM network.

Because LANE connectivity is defined at the MAC layer, upper protocol-layer functions of LAN applications can continue unchanged when the devices join ELANs. This feature protects corporate investments in legacy LAN applications.

An ATM network can support multiple independent ELANs. Membership of an end system in any of the ELANs is independent of the physical location of the end system. The end systems can move easily from one ELAN to another, regardless of whether or not the hardware is moved.

Components of LANE

A Catalyst 3900 or Catalyst 5000 ATM module can participate in up to 63 of these ELANs.

LANE is defined on a client-server LAN model, as follows:

- LEC
  An LEC emulates a LAN interface to higher layer protocols and applications. It forwards data to other LANE components and performs LANE address resolution functions.
  Each LEC is a member of only one ELAN. However, a router or a Catalyst ATM module can include LECs for multiple ELANs: one LEC for each ELAN of which it is a member.
  If a router has clients for multiple ELANs, the router can route traffic between the ELANs.

Note: If the Catalyst 3900 has multiple ATM modules and each has a client that is active for the same ELAN, the Catalyst 3900 will not bridge between the ELANs on the different modules. The Catalyst 3900 acts as an edge device on an ATM cloud (that is, there are no LANE services in the Catalyst 3900).

- LANE Server (LES)
  The LANE server for an ELAN is the control center. It provides joining, address resolution, and address registration services to the LANE clients in that ELAN. Clients can register destination unicast and multicast MAC addresses with the LANE server. The LANE server also handles LANE ARP (LE_ARP) requests and responses.
  The current configuration is one LES per ELAN.

- LANE Broadcast and Unknown Server (BUS)

  The LANE BUS sequences and distributes multicast and broadcast packets and handles unicast flooding.

  One combined LES and BUS is required per ELAN.

- LANE Configuration Server (LECS)

  The LECS contains the database that determines which ELAN a device belongs to (each configuration server can have a different named database). Each LEC contacts the LECS once, when it joins an ELAN, to determine which ELAN it should join. The LECS returns the ATM address of the LES for that ELAN.

  One LECS is required per ATM LANE switch cloud.

  The LECS database can have the following four types of entries:

  – ELAN name, ATM address of LANE server pairs

  – LANE client MAC address, ELAN name pairs

  – LANE client ATM template, ELAN name pairs

  – Default ELAN name

Note: ELAN names must be unique on an interface. If two interfaces participate in LANE, the second interface may be in a different switch cloud.

- Simple Server Redundancy Protocol (SSRP)

  The LANE simple server redundancy feature creates fault tolerance using standard LANE protocols and mechanisms. If a failure occurs on the LANE configuration server or on the LES/BUS, the ELAN can continue to operate using the services of a backup LANE server.

The Catalyst 3900 ATM module currently supports only the LEC function. A Catalyst 5000 or a Cisco 7000, Cisco 7200, Cisco 7500, RSP 7000, Cisco 4500, or Cisco 4700 with an AIP can supply all LANE functions.

LANE Operation and Communication

Communication among LANE components is typically handled by several types of VCCs. Some VCCs are unidirectional; others are bidirectional. Some are point-to-point and others are point-to-multipoint. Figure 4-3 illustrates the various types of VCCs.

Figure 4-3 LANE VCC Types

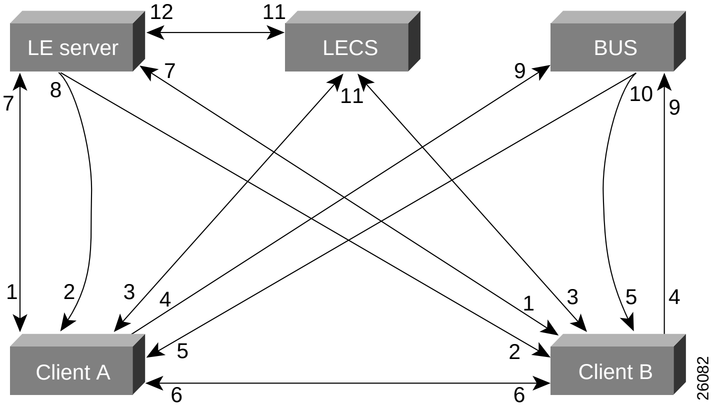

1–7	Control direct	4–9	Multicast send
2–8	Control distribute	5–10	Multicast forward
3–11	Configure direct (client)	6–6	Data direct
		11–12	Configure direct (server)

The following section describes the processes involved with a client requesting to join an ELAN.

Join Process

The following process (illustrated in Figure 4-3) normally occurs after an LEC has been enabled on the ATM module:

Step 1. The client requests to join an ELAN. The client sets up a connection to the LECS to find the ATM address of the LANE server for its ELAN. See the bidirectional, point-to-point link (link 1–7 in Figure 4-3).

An LEC finds the LECS using the following methods in the listed order:

- Locally configured ATM address
- Interim Local Management Interface (ILMI)
- Fixed address defined by the ATM Forum

Step 2. The LECS identifies the LES. Using the same VCC, the LECS returns the ATM address and the name of the LES for the client's ELAN.

Step 3. The client tears down the configure direct VCC.

Step 4. The client contacts the server for its LAN. The client sets up a connection to the LES for its ELAN (bidirectional, point-to-point control direct VCC [link 1–7 in Figure 4-3]) to exchange control traffic. Once a control direct VCC is established between an LEC and LES, it remains up.

Step 5. The LES verifies that the client is allowed to join the ELAN. The server for the ELAN sets up a connection to the LECS to verify that the client is allowed to join the ELAN (bidirectional, point-to-point server configure VCC [link 11–12 in Figure 4-3]).

The server's configuration request contains the client's MAC address, its ATM address, and the name of the ELAN. The LECS checks its database to determine whether the client can join that LAN; then it uses the same VCC to inform the server whether or not the client is allowed to join.

Step 6. The LES allows or disallows the client to join the ELAN. If allowed, the LES adds the LEC to the unidirectional, point-to-multipoint control distribute VCC (link 2–8 in Figure 4-3) and confirms the join over the bidirectional, point-to-point control direct VCC (link 1–7 in Figure 4-3). If disallowed, the LES rejects the join over the bidirectional, point-to-point control direct VCC (link 1–7 in Figure 4-3).

Step 7. The LEC sends LE_ARP packets for the broadcast address, which is all 1s. Sending LE_ARP packets for the broadcast address returns the ATM address of the BUS. Then the client sets up the multicast send VCC (link 4–9 in Figure 4-3) and the BUS adds the client to the multicast forward VCC (link 5–10 in Figure 4-3) to and from the BUS.

Step 8. The LEC registers the ring numbers of all other TrCRFs within its TrBRF that contain active ports on the local switch.

Addressing

On a LAN, packets are addressed by the MAC-layer address of the destination and the source stations. To provide similar functionality for LANE, MAC-layer addressing must be supported. Every LEC must have a MAC address. In addition, every LANE component (server, client, BUS, and configuration server) must have a unique ATM address.

All LECs on the same interface have a different, automatically assigned MAC address. That MAC address is also used as the end-system identifier part of the ATM address as explained in the following sections.

LANE ATM Addresses

A LANE ATM address has the same syntax as an NSAP, but it is not a network-level address. It consists of the following:

- A 13-byte prefix that includes the following fields defined by the ATM Forum: authority and format identifier (AFI) field (1 byte); Data Country Code (DCC) or International Code Designator (ICD) field (2 bytes); Domain Specific Part Format Identifier (DFI) field (1 byte); Administrative Authority (3 bytes); Reserved (2 bytes); Routing Domain (2 bytes); and Area (2 bytes).
- A 6-byte end-system identifier.
- A 1-byte selector field.

ILMI Address Registration

The Catalyst 3900 and Catalyst 5000 ATM modules use ILMI registration to build their ATM addresses and to register the addresses with the ATM switch. To build its ATM address, each module obtains its ATM address prefix from the ATM switch. Then it combines the ATM address prefix with its own MAC address and the selector value of 0 (zero). Once the ATM module has determined its ATM address, it uses ILMI to register this address with the ATM switch.

Address Resolution

As communication occurs on the ELAN, each client dynamically builds a local LE_ARP table. The LE_ARP table maps ELAN MAC addresses (Layer 2) to ATM addresses (also Layer 2). A client's LE_ARP table can also have static, preconfigured entries. The LE_ARP table maps MAC addresses to ATM addresses.

When a client first joins an ELAN, its LE_ARP table has no dynamic entries and the client has no information about destinations on or beyond its ELAN.

To learn about a destination when a packet is to be sent, the client begins the following process to find the ATM address corresponding to the known MAC address:

Step 1. The client sends an LE_ARP request to the LANE server for this ELAN (point-to-point control direct VCC [link 1–7 in Figure 4-3]).

Step 2. If the MAC address is registered with the server, it returns the corresponding ATM address. If not, the LES forwards the LE_ARP request to all clients on the ELAN (point-to-multipoint control distribute VCC [link 2–8 in Figure 4-3]).

Step 3. Any client that recognizes the MAC address responds with its ATM address (point-to-point control direct VCC [link 1–7 in Figure 4-3]).

Step 4. The LES forwards the response (point-to-multipoint control distribute VCC [link 2–8 in Figure 4-3]).

Step 5. The client adds the MAC address-ATM address pair to its LE_ARP cache.

Step 6. Now the client can establish a VCC to the desired destination and proceed to transmit packets to that ATM address (bidirectional, point-to-point data direct VCC [link 6–6 in Figure 4-3]).

For unknown destinations, the client sends a packet to the BUS, which forwards the packet to all clients. The BUS floods the packet because the destination might be behind a bridge that has not yet learned this particular address.

Traffic Handling

The Catalyst 3900 allows you to define up to 64 traffic profiles. These profiles can be used to define the maximum rates for each traffic type.

For each VLAN (or ELAN), a traffic profile must be mapped to each DD-VCC. The process of mapping depends on whether the traffic is incoming or outgoing:

- For incoming calls, the LEC tries to find a traffic profile that best matches the traffic parameters in the call. You can define the maximum discrepancy between the specified parameters and actual values on a per ELAN basis.
- For outgoing calls, you can define up to 10 profiles to use. The destination ATM address is ANDed with the address mask. The resulting ATM address is compared with the ATM address in the mapping table. If there is a match, each defined profile (0 through 9) is used in sequence until a call SETUP is accepted by the destination node.

Multicast Traffic

When an LEC has broadcast or multicast traffic, or unicast traffic with an unknown address to send, the following process occurs:

Step 1. The client sends the packet to the BUS (unidirectional, point-to-point multicast send VCC [link 4–9 in Figure 4-3]).

Step 2. The BUS forwards (floods) the packet to all clients (unidirectional, point-to-multipoint multicast forward VCC [link 5–10 in Figure 4-3]).

This VCC branches at each ATM switch. The switch forwards these packets to multiple outputs. (The switch does not examine the MAC addresses; it simply forwards all packets it receives.)

Recommended Environments

This section describes some scenarios for using the Catalyst Token Ring switches with the ATM expansion module. Figure 4-4 shows how the ATM expansion module can be used to connect Catalyst Token Ring switches through an ATM switch.

Figure 4-4   Token Ring Connection Over an ATM Backbone

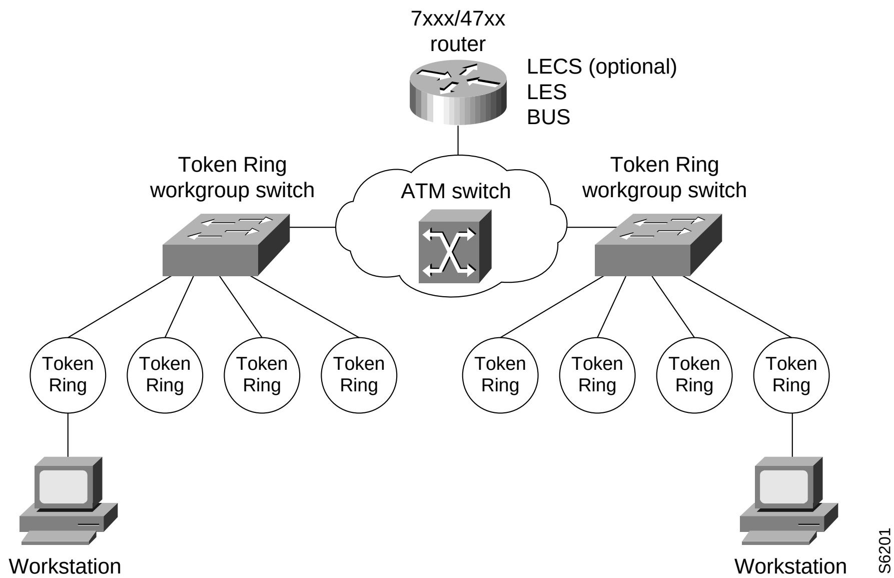

The ATM expansion module is well suited for the following environments:

- ATM backbone for legacy LANs

  A Catalyst Token Ring switch with an ATM expansion module provides a seamless, switched network between legacy LANs communicating over ATM. Therefore, as a first step in migrating from legacy LANs, many users deploy ATM in the backbone or as the WAN technology to connect geographically dispersed legacy LANs.

- LAN-to-ATM interoperability

  A Catalyst Token Ring switch with one or more ATM expansion modules can help protect your legacy LAN investment by providing transparent LAN-to-ATM switching. Therefore, as the next step in a legacy LAN-to-ATM migration, users place high-speed or frequently accessed servers, or both, on the ATM network to take advantage of ATM's scalability. By using LANE over ATM, ATM-based workstations are able to use existing legacy LAN applications with minimal or no upgrade costs for moving to ATM.

Using Multiple ATM Expansion Modules

Because one ATM expansion module can support as many LECs as there are VLANs in a Catalyst Token Ring switch, the question arises: Why use two ATM expansion modules?

The first reason is to provide a backup LEC. By enabling two LECs on two different ATM expansion modules to be members of both the same VLAN and the same ELAN, the Catalyst Token Ring switch's spanning-tree operation will automatically use one LEC for forwarding frames and the other LEC for blocking frames (active standby). For the backup LEC configuration to work, spanning tree *must* be enabled for the related LAN switch domain.

Note: The STP (802.1d or IBM) used depends on whether the bridging mode is configured as SRB or SRT. If the first LEC fails, the second LEC will automatically take over.

The second reason to use two ATM expansion modules is to increase system resources. In some environments the resources associated with a single ATM expansion module might affect individual LEC performance. By moving one or more LECs to a second ATM expansion module, you can significantly increase the resources available to each LEC.

Note: If the Catalyst 3900 has multiple ATM modules and each has clients active for the same ELAN, the Catalyst 3900 will not bridge between the ELANs. The Catalyst 3900 acts as an edge device on an ATM cloud.

Usage Recommendations and Restrictions

Before installing and configuring the ATM module, be aware of the following:

- Because the ATM module cannot internally bridge between TrCRFs, you can configure only one LEC per TrBRF per ATM module. Therefore, while a TrBRF can include only one LEC from a single ATM expansion module, it can contain multiple LECs when each LEC exists on a separate ATM module.
- Configuring a TrBRF to contain an ATM port restricts the MTU of that TrBRF to 4472 or less.
  The ATM module does not support MTUs greater than 4472. Configuring an MTU size larger than 4472 for a TrBRF that contains an ATM port will result in the ATM port being removed from the TrBRF. Additionally, if a TrBRF is configured for an MTU size larger than 4472, none of the TrCRFs assigned to the TrBRF can be assigned to an ATM port.