

What is the LightStream 1010 ATM Switch?

The LightStream 1010 ATM switch is Cisco System's next-generation Asynchronous Transfer Mode (ATM) switch for workgroup and campus backbone deployment. Incorporating support for the latest ATM Forum specifications and building upon the Cisco System's Internetwork Operating System (IOS™) software, the LightStream 1010 offers the most complete, sophisticated feature set of any ATM switch in its class. It also provides the performance, scalability, and robustness required for production ATM deployment.

The LightStream 1010 uses a five-slot, modular chassis featuring the option of dual, fault-tolerant, load-sharing power supplies. The central slot in the LightStream 1010 is dedicated to a single, field-replaceable ATM switch processor (ASP) module that supports five-Gbps shared memory and the fully nonblocking switch fabric. The ASP also supports the feature card and high-performance reduced instruction set (RISC) processor that provides the central intelligence for the device. The remaining slots support up to four hot-swappable carrier modules (CAMs). Each CAM supports up to two hot-swappable port adapter modules (PAMs) for a maximum of eight PAMs per switch, supporting a wide variety of desktop, backbone, and wide-area interfaces.

The LightStream 1010 ATM switch provides switched ATM connections to individual workstations, servers, LAN segments, or other ATM switches and routers using fiber-optic, unshielded twisted-pair (UTP), and coaxial cable.

The LightStream 1010 ATM switch can accommodate up to 32 OC-3 switched ATM ports in a standard 19-inch (48-centimeter) rack.

Summary of Features

The LightStream 1010 ATM switch provides the following features:

- ATM User Network Interface (UNI) for both network and user sides
- ATM Network-to-Network Interface (NNI) between switches
- Support for the following ATM connection types:
 - Permanent virtual channel connections (PVCCs)
 - Permanent virtual path connections (PVPCs)
 - Soft permanent virtual channel connections (SPVCCs)
 - Soft permanent virtual path connections (SPVPCs)
 - Switched virtual channel connections (SVCCs)
 - Switched virtual path connections (SVPCs)
 - Virtual path (VP) tunneling
 - Point-to-point ATM connections
 - Point-to-multipoint ATM connections
- ATM fault management Operation, Administration, and Maintenance (OAM) (F4 and F5 flow)
- Resource management to support the following categories of service:
 - Constant bit rate (CBR)
 - Variable bit rate (VBR)
 - Available bit rate (ABR)
 - Unspecified bit rate (UBR)
- Interim Local Management Interface (ILMI) as per UNI 3.0/3.1
- ATM UNI signaling protocol for both user and network as per UNI 3.0/3.1

- ATM NNI signaling protocols:
 - Interim Interswitch Signaling Protocol (IISP) as defined by the ATM Forum
 - Private Network-Network Interface (PNNI), a pre standard release of the ATM Forum PNNI that supports a single peer group
- LAN Emulation Client as defined by ATM Forum LAN Emulation over ATM Version 1.0
- In-band device network management using IP over ATM Simple Network Management Protocol ([SNMP] agent V1/V2 and Telnet)
- Out-of-band device network management using Ethernet and console ports
- ATM and other applicable industry Management Information Bases (MIBs)
- Command line interface (CLI) at the RS-232 console and Telnet
- Upload/Download software code image
- Port snooping

The following interfaces are supported:

- 155-Mbps:
 - Synchronous Optical Network (SONET) multimode fiber-optic cable
 - SONET single-mode fiber-optic cable
 - UTP cable
- 45-Mbps DS-3 coaxial cable
- 34-Mbps E-3 coaxial cable
- 622 Mbps SONET OC-12 single-mode fiber-optic cable

ATM Overview

ATM technology now plays a central role in the evolution of current workgroup, campus and enterprise networks. ATM delivers important advantages over existing LAN and WAN technologies, including the promise of scalable bandwidths at unprecedented price and performance points and QOS guarantees. These features allow many new classes of applications, such as multimedia, to use ATM technology.

These benefits, however, come at a price. Contrary to common misconceptions, ATM is a very complex technology, perhaps the most complex ever developed by the networking industry. Whereas the structure of ATM cells and cell switching do facilitate the development of hardware-intensive, high-performance ATM switches, the deployment of ATM networks requires the overlay of a highly complex, software intensive, protocol infrastructure. This infrastructure is required to allow both individual ATM switches to be linked into a network, and for such networks to function with the vast installed base of existing local and wide-area networks.

Connection-Oriented Network

The fact that ATM is connection-oriented implies the need for ATM-specific signaling protocols and addressing structures, as well as protocols to route ATM connection requests across the ATM network. These ATM protocols, in turn, influence the manner in which existing higher layer protocols can operate over ATM networks. The routing which can be done in many different ways, each with its own advantages and characteristics, will be discussed in the section “ATM Routing.”

ATM Network Operation

An ATM network consists of a set of ATM switches interconnected by point-to-point ATM links or interfaces. ATM switches support two kinds of interfaces:

- UNI—Used to connect ATM end systems (hosts, routers, and so on) to an ATM switch. See the section “ATM Signaling.”
- NNI—Used to connect two ATM switches together.

See the section “ATM Routing.”

Slightly different cell formats are defined across UNI and NNI. More precisely, however, an NNI is any physical or logical link across which two ATM switches exchange the NNI protocol. This routing setup will be described in greater detail in the section “ATM Routing.”

ATM networks are fundamentally connection-oriented, which means that a virtual circuit needs to be set up across the ATM network prior to any data transfer. ATM circuits are of two types:


- Virtual path (VP), identified by a virtual path identifier (VPI).
- Virtual Channel (VC), identified by the combination of a VPI and a virtual channel identifier (VCI).

Both are described in the following paragraphs:

Virtual Path Connections

A single virtual path connection (VPC) can be used to route many virtual channels through the ATM network. Figure 1-1 shows a VPC between switch A and switch C through switch B.

Figure 1-1 Example of a Virtual Path Connection



In Figure 1-1, notice that VPI 20 is made up of VCI numbers 501, 502, and 503 and that the VCI numbers did not change as the VP was switched through switch B from VPI 20 to VPI 43. Because a VP simply routes the VC through the network, a cell is guaranteed to have the same VCI when it exits the VPC as when it enters.


Note Concatenation of the VCIs with the VPIs is needed to uniquely identify a virtual connection as it passes through the ATM network.

All VCIs and VPIs, however, have only local significance across a particular link and are remapped, as appropriate, at each switch.

Virtual Channel Connection

A virtual channel connection (VCC) is established as a bidirectional facility to transfer ATM traffic between two ATM layer users. A VCC is established at the time a VC session is activated. Figure 1-2 shows a VCC between ATM users A and D.

Figure 1-2 Example of a Virtual Channel Connection



In each direction, at a given interface, different VCs are multiplexed onto a physical circuit. The VPIs and VCIs identify these multiplexed connections.

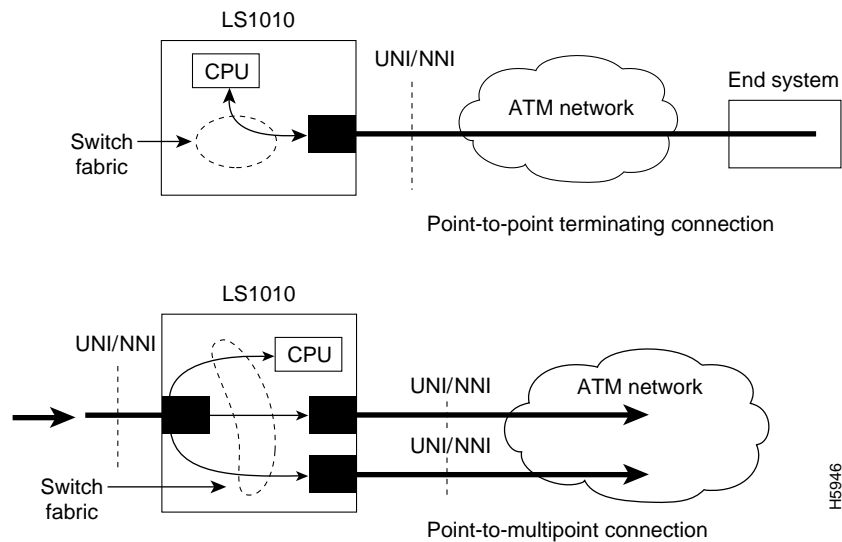
Note The value of the identifiers may change as the traffic is relayed through the ATM network.

Virtual Connection Types

The LightStream 1010 ATM switch supports the following types of VC connections, shown in Figure 1-3:

- Terminating—used to create a datapath to the LightStream 1010 switch CPU for network management.
- Transit—used to allow a single datapath through the switch.
- Point-to-Multipoint—used to allow a single ingress datapath to branch to multiple egress paths.

Figure 1-3 Virtual Connection Types



Basic Switch Operation

The basic operation of an ATM switch requires the following steps:

- Step 1** Receive a cell across a link on a known VCI or VPI value.
- Step 2** Look up the connection value in a local translation table to determine the outgoing port (or ports) of the connection and the new VPI/VCI value of the connection on that link.
- Step 3** Retransmit the cell on that outgoing link with the appropriate connection identifiers.

The switch operation is so simple because external mechanisms set up the local translation tables prior to the transmission of any data. The manner in which these tables are set up determines the two fundamental types of ATM connections:

- Permanent virtual connection (PVC)—a connection set up by some external mechanism, typically network management, in which a set of switches between an ATM source and a destination ATM system are programmed with the appropriate VPI/VCI values. As discussed later, ATM signaling can facilitate the setup of PVCs, but, by definition, PVCs always require some manual configuration, causing their use to be somewhat cumbersome.
- Switched virtual connection (SVC)—a connection that is set up automatically through a signaling protocol. SVCs do not require the manual interaction needed to set up PVCs and, as such, are likely to be much more widely used. All higher-layer protocols operating over ATM primarily use SVCs, and it is these protocols that are primarily used by the LightStream 1010 switch.

ATM signaling is initiated by an ATM end system that desires to set up a connection through an ATM network; signaling packets are sent on a well-known VC, VPI = 0, VCI = 5. Table 1-1 lists other well-known terminating PVCCs.

Table 1-1 Some Well-Known Terminating PVCCs

VPI/VCI	Creation Condition	Purpose
0/5	Create for both UNI/NNI	For signaling
0/16	Create for both UNI/NNI	For ILMI ¹
0/18	Create for an NNI only	For PNNI

1. Interim Local Management Interface

The signaling is routed through the network, from switch to switch, setting up the connection identifiers as it goes, until it reaches the destination end system. The latter can either accept and confirm the connection request, or can reject it, clearing the connection.

Note Because the connection is set up along the path of the connection request, the data also flows along this same path.

The following paragraphs describe the detailed implementation used by the LightStream 1010 to switch cells on an ATM network:

Light Stream 1010 Applications

The LightStream 1010 ATM switch can accommodate up to 32 switched OC-3 ATM ports in a standard 19-inch (48-centimeters) rack. The LightStream 1010 chassis has five slots. The middle slot (number 2) is used for the ASP, which provides Layer 2 switching, with both local and remote management. See the section “Hot-Swapping Feature” in the chapter “LightStream 1010 ATM Switch Hardware” for a detailed description of the switch module and port numbering scheme.

Figure 1-4 shows an example of a network configuration using the LightStream 1010 ATM switch in a high-performance workgroup.

Figure 1-4 **LightStream 1010 Workgroup Configuration Example**




Figure 1-5 shows an example of a network configuration using the LightStream 1010 ATM switch for a campus backbone.

Figure 1-5 LightStream 1010 Backbone Configuration Example

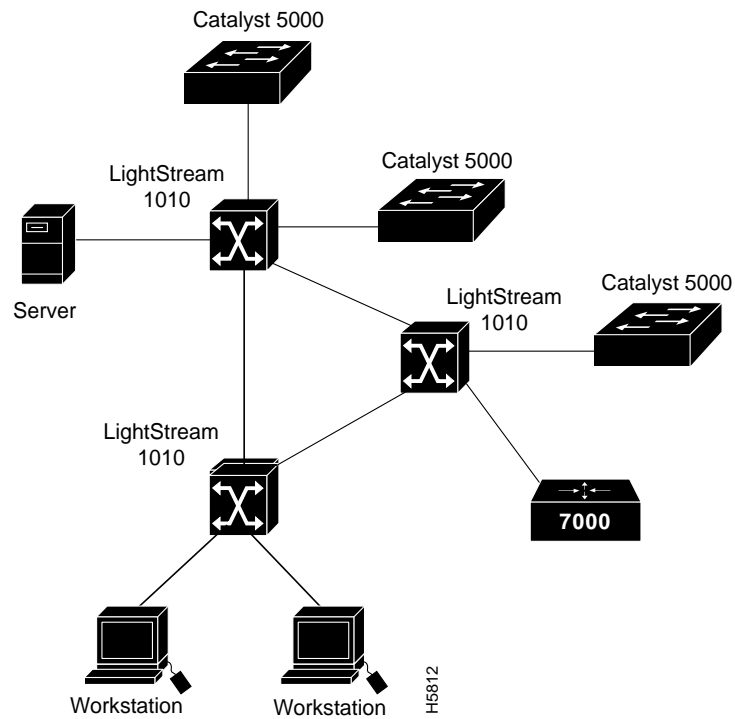


Figure 1-6 shows an example of a network configuration using the LightStream 1010 ATM switch in a multisite network.

Figure 1-6 LightStream 1010 Multisite Configuration Example




Figure 1-7 shows an example of a network configuration using the LightStream 1010 ATM switch in a LANE network.

Figure 1-7 LightStream 1010 LANE Configuration Example





Figure 1-8 shows an example of a network configuration in which private switches form a private network interconnected over permanent VPs. These VPs provide logical trunks tunnelling across a public network.

Figure 1-8 **LightStream 1010 VP Tunneling Example**



LightStream 1010 System Functions

The following sections describe the implementation used by the LightStream 1010 to switch cells on an ATM network:

- ATM Addressing and Plug-and-Play Operation
- ATM Signaling
- ATM Routing
- Advanced Traffic Management
- Operation Administration and Maintenance Support

- LAN Emulation Client
- Classical IP over ATM
- Network Management
- ATM Standards Compliance

ATM Addressing and Plug-and-Play Operation

ATM routing protocols use ATM end-system addresses (AESAs) based on the ISO network service access point (NSAP) format addresses. Prefixes of such addresses identify individual switches and collections of switches within peer groups. Switches also supply these prefixes to attached end systems, using the ILMI protocol. ATM end-system address prefixes can be obtained either from ATM service providers or directly from the various national authorities designated by the International Organization for Standardization (ISO) to allocate the NSAP address space. The mechanisms for administering such addresses are not yet well understood by the industry, however, and there are few ATM service providers from whom customers may obtain their addresses.

In order to overcome these problems, all LightStream 1010 switches are preconfigured with Cisco ATM address prefixes. These prefixes are combined with one of the switch's preconfigured Media Access Control (MAC) address in order to develop a unique node identifier. These identifiers are then used both to configure attached end systems and to automatically bring up the PNNI routing hierarchy, hence offering true plug-and-play operation. Such autoconfigured addresses will suffice for a small peer group of a few dozen switches, permitting small-scale ATM internetworks to be deployed while allowing time for customers to obtain and allocate their own ATM address prefixes.

The LightStream 1010 also automatically recognizes the type of any port modules plugged into a carrier module and, if the new PAM is of the same type as the previously plugged in PAM, will also automatically restore any interface-specific configuration and PVCs saved to memory. Upon a reboot, the switch will also automatically restore all PVCs and any other configuration information stored within its non volatile memory.

The LightStream 1010 also uses the ILMI protocol to automatically recognize the nature of any new ATM interface—for example, whether it is a UNI or NNI, public or private interface—hence eliminating any manual configuration.

Such recognition and restoration mechanisms, when combined with such IP address autoconfiguration mechanisms as BootP, cause the LightStream 1010 to be made fully self-configuring. Through network management applications or the text-based CLI, the network operators, if desired, will also have the capability to:

- Configure and customize all aspects of the operation of the switch.
- Monitor device and network status through sophisticated troubleshooting or debugging procedures.

ATM Signaling

The LightStream 1010 has full, integrated support for ATM Forum-compliant UNI 3.0/3.1 signaling, and UNI 3.1 signaling. All elements of the UNI standards, including both point-to-point and point-to-multipoint signaling capabilities, are supported, as is the ILMI protocol. The ILMI protocol uses SNMP based mechanisms across all interfaces to automatically identify which of its interfaces are UNI, attached to ATM end systems and which are NNI, attached to other systems. Further it can differentiate between private and public network links. Such information is used by the ATM routing protocols to automatically discover and bring up a network of interconnected LightStream 1010 switches.

The ILMI protocol is also used for ATM address registration across ATM UNI in order to program ATM end systems with ATM address prefixes. The ILMI protocol is also used so that the switch can discover the 48-bit MAC addresses of the attached systems. ILMI not only reduces the need for manual configuration of attached end systems, but is also important for the plug-and-play operation of LightStream 1010-based networks—for instance, it is also used for informing LAN Emulation (LANE) clients of the location of the LANE Configuration Server (LECS).

UNI signaling is also used by ATM end systems to inform LightStream 1010s of the desired traffic characteristics and QOS for new connections and the acceptance or rejection of such connection requests. Fully integrating support for ATM signaling onto the ASP module ensures not only high performance, but also high reliability since no external signaling processor is required.

LightStream 1010 System Functions

Normally, when ILMI link auto-determination is enabled on the interface and is successful, the switch will take the UNI-version returned by ILMI. If the ILMI link auto-determination is unsuccessful or if ILMI is disabled, UNI-version defaults to Version 3.0. This can be overridden by using the **atm uni-version** command.

Soft PVCs

Soft PVC support facilitates PVC setup by using signaling protocols across the network.

VP Tunneling

VP tunneling enables signaling across public networks.

Access Lists

The access control list is used by the ATM signaling software to filter setup messages on an interface or sub interface as either destination or source. Access lists can be used to deny connections that are known to be a security risk and permit all other connections, or to permit those connections that are considered acceptable and deny all the rest.

ATM Routing

The LightStream1010 supports the routing of ATM signaling requests across a network of switches using the ATM routing protocols. Two standard routing protocols have been developed by the ATM Forum—IISP and the PNNI Version 1. Both protocols are supported by the LightStream 1010—the default IISP software image supports the IISP protocol only; the optional PNNI image supports the PNNI Version 1 protocol.

IISP

IISP uses a combination of static routing and UNI signaling, and is more suitable for small networks of switches. See the section “ATM Signaling.”

PNNI

The PNNI Version 1 Protocol has been designed to scale to the largest possible ATM networks, encompassing thousands of switches, and is the preferred solution for larger networks. Switches running the PNNI image can also be configured with IISP interfaces in order to permit a smooth evolution to the PNNI Version 1 Protocol. Additional ATM routing capabilities allow for load-balancing across redundant links and increase system reliability.

The PNNI Version 1 Protocol supports sophisticated mechanisms to permit both QOS-based routing of ATM connection requests and scalability. Routing is supported by a link-state routing protocol and the exchange of QOS routing metrics and attributes between all nodes. The metrics and attributes supported by Cisco PNNI implementation include the following:

- Administrative weight
- Peak-to-Peak (maximum) cell delay variation for CBR and VBR connections
- Maximum cell transfer delay for CBR and VBR connections
- Available cell rate per class
- Maximum cell rate per interface
- Maximum Loss Ratio for the Cell Loss Priority (CLP) 0 cells of CBR and VBR connections
- Cell rate margin
- Variance factor

The PNNI Version 1 Protocol specifies the formats and types of information exchanged by the protocol, and PNNI signaling is used to forward signaling requests across the network. However many other aspects of the protocol are left as a matter of implementation. The PNNI implementation on LightStream 1010 incorporates many value-added capabilities to extend beyond the base standards.

Crankback

Both IISF and PNNI Version 1 implementations on the LightStream 1010 support crankback, which allows connections to be rerouted around nodes whose local CAC reject the connection, hence reducing setup latency. Load-balancing mechanisms allow for the support of redundant links, with either first-fit or load-balancing selection algorithms.

Advanced Traffic Management

The LightStream 1010 supports the most advanced traffic management and ATM signaling and routing capabilities. In addition to ATM Forum-compliant ABR congestion control, the LightStream 1010 supports the following mechanisms required to deliver QOS on demand to all ATM traffic classes and ATM adaptation layer (AAL) types:

- Traffic policing
- Multiple levels of priority
- Intelligent packet discard
- Connection admission control
- Traffic pacing

Many of these capabilities are supported on the field-replaceable feature card on the ASP module. The feature card allows easy upgrading as and when newer mechanisms are standardized or required.

The following different mechanisms are required to support the guaranteed QOS classes:

- CBR
- VBR

and the following best-effort classes

- UBR
- ABR

In order to set up a connection with a guaranteed QOS across a network of LightStream 1010 switches, an end system must first inform the switch, attached across a UNI, of the required QOS parameters and the characteristics of the intended traffic flow.

UNI 3.0 and UNI 3.1 do not explicitly specify individual QOS parameters or explicitly request traffic classes. Since only ATM Forum UNI 4.0 signaling does, the LightStream 1010 will map the UNI 3.0/3.1 bearer classes into the appropriate UNI 4.0 service categories and QOS parameters. These characteristics together constitute the traffic descriptor for the connection.

Connection Admission Control

Sophisticated connection admission control (CAC) algorithms, based upon the particular architecture of the LightStream 1010, allow maximization of switch and link utilization, while precluding the possibility of preestablished guarantees being violated. These algorithms, together with the ATM Forum-specific generic CAC (GCAC) algorithm, permit the rapid, accurate determination of source routes across a network of LightStream 1010 ATM switches. User controls allow regulation of strictness or looseness of the CAC algorithm, providing either greater network utilization or stricter control over guarantees.

Once the traffic descriptor is derived, the source switch first performs a CAC function to determine whether local resources can support the requested connection. If the resources can support the requested connection, the LightStream 1010 will use the PNNI Version 1 protocols to determine a source route through the network that its GCAC algorithms suggest will be able to support the requested connection, and will then forward the request using PNNI Version 1 signaling. At each switch, CAC is used to ensure that the requested connection can be supported. If and when the connection is routed through to the final destination and is accepted, then the LightStream 1010 will complete the connection and will inform the requesting end system.

Categories of Service

In order to deliver such guarantees, the LightStream 1010 implements the following four configurable levels of delay priority into which the various traffic classes are mapped:

- CBR—used to support isochronous applications such as high-quality video
- VBR—used to support information retrieval services in which large amounts of video are sent to the user followed by delays as the user examines the information
 - VBR-RT-real-time (VBR-RT)
 - VBR-non real-time (VBR-NRT)

LightStream 1010 System Functions

- ABR—used to support bursty data traffic
- UBR—used to support bursty data traffic

By supporting so many levels of priority, the LightStream 1010 can ensure full separation of the various traffic classes, and hence ensure absolute priority for the guaranteed services over the best-effort services.

The buffering of the LightStream 1010 can be flexibly allocated between each of the traffic classes and output ports in order to ensure that the switch can be tuned for any particular traffic profile or deployment scenario. The following configurable buffer controls are available:

- Per-switch controls on maximum buffer limits for each service class
- Per-service category maximum buffer sizes per port

The maximum buffer controls on the service categories can be modified to adjust the over-subscription factor. This future controls the degree of random multiplexing—and hence the amount of effective buffering—within the switch fabric. Fixed minimum buffer limits preclude the possibility of buffer starvation where high priority traffic consumes all available switch buffers, leading to the possibility of head-of-line blocking.

In addition to the use of CAC, delay priority, and buffer controls in order to deliver QOS guarantees, the LightStream 1010 also uses traffic policing or usage parameter control (UPC) in order to monitor the compliance of ATM end systems to agreed traffic contracts. As per the ATM Forum UNI specification, the LightStream 1010 supports a dual-mode leaky bucket algorithm; highly flexible, this algorithm, that permits the automatic policing of the most important parameters for each of the traffic classes—for instance, the peak cell rate and cell delay variation tolerance of CBR connections, the burst size and sustainable cell rate of VBR connections, or the peak cell rate for UBR/ABR connections. Such traffic policing parameters are extracted automatically from UNI signaling and programmed into the switch fabric.

Cells found to be noncompliant by the UPC function with their stated traffic contracts can either be dropped or, if buffer resources permit, passed but tagged for preferential dropping by having their cell loss priority (CLP) bits set to 1. Correspondingly, the LightStream 1010 also implements two levels of CLP. A configurable threshold can be set on each of the per-service class port buffers; beyond the threshold the switch will accept only cells with the CLP bit not set, hence favoring conforming traffic. Cells with the CLP bit set to 1 will

be preferentially dropped. Such cells, particularly for the best-effort services, will not be dropped arbitrarily, however, but by using the intelligent packet discard mechanisms of the LightStream 1010.

Intelligent Packet Discard

The ATM Forum UNI signaling protocols allow end systems to designate which particular connections carrying packet-based traffic can be eligible for packet drop; the LightStream 1010 interprets such signaling and automatically identifies the connections. These connections will typically be used to carry traffic from protocols such as TCP/IP, which analysis has shown to yield far higher throughput in networks employing packet-based rather than cell-based discard policies. The LightStream 1010 builds upon such analysis with its intelligent packet-discard mechanisms, ensuring that any event that might cause a cell drop—cell-tagging threshold or that of buffer overflows being exceeded—invokes the packet-dropping mechanisms.

When a cell-dropping threshold is exceeded for a cell of a particular packet flow, the LightStream 1010 will put the associated connection into a tail-drop mode, dropping all the other cells that might arrive for that particular packet. The only cell that remains is the final cell in the packet which is forwarded with the CLP bit reset to 0. This setting ensures that the cell reaches the final destination, and so allows completion of any packet re-assembly that may have started. A second, higher threshold is defined for each port buffer that defines the intelligent packet discard (PD) point. Beyond this threshold, the switch begins discarding all packets along a particular connection. Future enhancements will let connections be picked for discard based upon various policies, such as highest volume to increase fairness in throughput.

These mechanisms act together to make the LightStream 1010 essentially emulate a packet switch such as, a LAN switch or router, which also acts upon and drops entire packets. These packet discard mechanisms work closely with higher layer protocols such as TCP, which are designed to rapidly adjust packet flow rates in response to network congestion, as signaled by packet discards, or, conversely, rapidly increase packet throughputs so as to take advantage of network bandwidth availability. Through such mechanisms, the LightStream 1010 can give to ATM networks the same scalability of current packet networks—which, as the global Internet with its millions of nodes attests, are essentially scalable without limit.

The LightStream 1010 can also go beyond current packet-based mechanisms, however, for even greater scalability and faster network convergence through its support of the ATM Forum ABR congestion-control mechanisms.

ABR Support and ATM Congestion Control

ABR specifications define a series of congestion-control mechanisms that switches and ATM end systems use to regulate the amount of traffic sent across ABR connections, and hence, minimize cell loss.

ABR is a very sophisticated specification with multiple possible modes of operation, specifying the behavior of both source and destination end systems and of intermediate switches. Source end systems periodically generate in-line resource management (RM) cells, which are sent intermixed with data cells along all connections. These cells are then received by destination end systems and returned along the backward connection to the source indicating whether or not intermediate switches have experienced congestion. Switches indicate congestion in three ways:

- 1 In explicit forward congestion indication (EFCI) marking mode, switches set a bit in the headers of forward cells to indicate congestion; these bits are then turned around at the destination and sent back to the source.
- 2 In relative rate marking mode, switches set a bit within either forward *or* backward or forward *and* backward RM cells to indicate congestion.
- 3 In explicit rate marking mode, switches indicate within either forward *or* backward or forward *and* backward RM cells exactly which rate they are willing to receive along a particular path.

The source end system uses specified algorithms to control the allowed cell rate (ACR) depending upon the type of feedback received. In general, this list of mechanisms is ranked in order of both sophistication and complexity. In general terms, EFCI marking—though widely available, particularly on wide-area switches—provides only minor improvements in performance because of the latency of turning around the forward congestion indication. As with any feedback mechanism, congestion control schemes operate optimally when the latency of the feedback path is minimized—indeed, excessive latencies can be counterproductive, since they might cause sources to slow down unnecessarily once the network congestion has already eased.

For this reason, the relative rate mode delivers much greater performance than the EFCI mode, since the ability to use the backward RM cells to indicate congestion rather than relying upon the destination end system to turn it around, greatly reduces latency.

The LightStream 1010 allows individual thresholds to be set for EFCI or relative rate marking, and a higher threshold for early packet and CLP discard. Any occasion in which a cell is dropped, however, including UPC, will trigger the tail packet discard mechanism. This intelligent packet discard, in which no cell is ever arbitrarily dropped is a unique, patented capability of the LightStream 1010.

Traffic Pacing

In addition to traffic policing and congestion-control mechanisms, the LightStream 1010 also supports a unique traffic pacing mechanism. This mechanism allows the rate at which traffic is sent across any port to be limited to any of a wide variety of peak traffic rates. This capability is particularly important when connecting across a public UNI to a public network, since many such networks will base their tariffs to a particular aggregate bandwidth. Future enhancements to the LightStream 1010, implemented on the next version of the feature card, will allow even greater traffic-shaping capabilities—such as, on a per-connection or VP basis.

Operation Administration and Maintenance Support

The LightStream 1010 has full support for ATM OAM cell flows—F4 flows used within VPs and F5 flows used within VCs. The LightStream 1010 can be configured either for end-to-end or segment loopback F4 and F5 flows; cells can be sent either on demand or periodically in order to verify link and connection integrity. Alarm indication signal (AIS) and remote defect indication (RDI) functions are also supported on both F4 and F5 flows.

In addition to standard OAM functions, the LightStream 1010 also has the unique value-added capability to send OAM pings (OAM cells containing ATM node prefixes or IP addresses) of intermediate switches, allowing network administrators to determine the integrity of a chosen connection at any intermediate point from any other point along the connection. This capability greatly eases network debugging and troubleshooting.

Port Snooping

The LightStream 1010 offers a unique port-snooping and connection-steering mechanism that lets all connections in any given direction on a selected port be transparently mirrored to a particular monitoring port to which an external ATM analyzer is attached. This capability is critical for the monitoring and troubleshooting of ATM switching systems since, unlike current shared-medium LANs, network traffic flows cannot easily be monitored with external devices. Future enhancements will further increase the transparent monitoring capabilities of the LightStream 1010, since such monitoring is critical to network managers confidently deploying switching technologies.

LAN Emulation Client

The LightStream 1010 provides an interface to switched LANs across an ATM network, providing LAN users with access to ATM-based services. LAN Emulation (LANE) extends virtual LAN (VLAN) throughout the network by establishing ATM connections between clients and servers. Figure 1-7 shows an example of an ATM LANE configuration.

A LANE client can be configured for the CPU for in-band management (Telnet and Simple Network Management Protocol [SNMP]).

Classical IP over ATM

Cisco implements classical IP and Address Resolution Protocol (ARP) over ATM as described in RFC 1577. RFC 1577 defines an application of classical IP and ARP in an ATM environment configured as a logical IP subnetwork (LIS). It also describes the functions of an ATM ARP server and ATM ARP clients in requesting and providing destination IP addresses and ATM addresses in situations when one or both are unknown. Our switches can be configured to act as either ARP client or ARP server.

The ATM ARP mechanism is applicable to networks that use SVCs in a pure PVC environment. It requires a network administrator to configure only the device's own ATM address and that of a single or multiple ATM ARP server into each client device. When the client makes a connection to the ATM ARP server, the server sends ATM inverse ARP requests to learn the IP network address and ATM address of the client on the network. It uses the addresses to resolve future ATM ARP requests from clients. Static configuration of the server is required.

In Cisco's implementation, the ATM ARP client tries to maintain a connection to the ATM ARP server. The ATM ARP server can tear down the connection, but the client attempts once each minute to bring the connection back up. No error messages are generated for a failed connection, but the client will not send packets until the ATM ARP server is connected and translates IP network addresses.

For each packet with an unknown IP address, the client sends an ATM ARP request to the server. Until that address is resolved, any IP packet sent to the ATM interface will cause the client to send another ATM ARP request. When the ARP server responds, the client opens a connection to the new destination so that any additional packets can be sent to it.

Network Management

The LightStream 1010 ATM switch can be managed through the administrative interface. The administrative interface connects directly to a console terminal, or through a modem that connects to the EIA/TIA-232 interface on the ASP. See the chapter “Configuring Terminal Lines and Modem Support” in the *LightStream 1010 ATM Switch Software Configuration Guide* publication. Alternatively, the administrative interface can be accessed using SNMP, Telnet, Serial Line Internet Protocol (SLIP), or Point-to-Point Protocol (PPP).

Simple Network Management Protocol

Today, SNMP is the most popular protocol for managing diverse commercial, university, and research internetworks. SNMP is an application-layer protocol designed to facilitate the exchange of management information between network devices. The SNMP system consists of three parts: SNMP manager, SNMP agent, and MIB.

Instead of defining a large set of commands, SNMP places all operations in a *get-request*, *get-next-request*, and *set-request* format. For example, an SNMP manager can get a value from an SNMP agent or store a value into that SNMP agent. The SNMP manager can be part of a network management system (NMS), and the SNMP agent can reside on a networking device such as a switch. The SNMP agent can respond to MIB-related queries being sent by the NMS.

Basic functions supported by SNMP agents include:

- Accessing a MIB variable using the *get-request* or *get-next-request* format—initiated by the SNMP agent as a result of a request for the value of a MIB variable from a network management station. The SNMP agent gets the value of a MIB variable by accessing information stored in the MIB and then responding to it.
- Setting a MIB variable—initiated by the SNMP agent as a result of a message from a network management station. The SNMP agent requests that the value of a MIB variable be changed.
- SNMP trap—notifies a network management station that an extraordinary event has occurred at an agent. When a trap condition occurs, the SNMP agent sends an SNMP agent trap message to each of the network management stations as specified in the trap receiver table.

CiscoView for the LightStream 1010

The CiscoView application for the Cisco LightStream 1010 ATM switch is a graphical user interface (GUI)-based device management application that provides dynamic status, statistics, and configuration information for the LightStream 1010. This application is a subset of the CiscoView application that offers similar features for the other switched internetworking products such as Catalyst LAN switches and Cisco routers.

The CiscoView application is also included within CiscoWorks, Cisco's enterprise network management application suite. It displays a graphical view of any Cisco device and shows the real-time LED in interface status. The CiscoView application provides comprehensive monitoring functions and simplifies basic troubleshooting tasks.

With the CiscoView application for the LightStream 1010, users can more easily understand the complex management data and configuration parameters available for ATM switches. It organizes this information into a graphical representation in a clear, consistent format.

The CiscoView application for the LightStream 1010 can be integrated with several of the leading network management platforms, providing management application integration. CiscoView can be run on UNIX workstations as a fully functional, independent LightStream 1010 management application. In addition, it can be launched from the AtmDirector topology map by simply double-clicking on a LightStream icon.

ATM Standards Compliance

ATM is defined by a large number of cross-referenced specifications from a variety of standards bodies including the ATM Forum, the American National Standards Institute (ANSI) T1S1 Committee, Bellcore, the European Telecommunications Standards Institute (ETSI), and the International Telecommunication Union Telecommunication Standardization Sector (ITU-T). The Forum specifications are regarded as preeminent since they build upon, and refer to, specifications from all other ATM standards bodies; hence, compliance with ATM Forum specifications implies compliance also with all referenced specifications. The LightStream 1010 supports the following standards:

- ATM Forum UNI 3.0 /3.1
- ATM Forum ILMI
- ATM Forum PNNI Version 1
- ATM Forum IISP
- ATM Forum TM 4.0 (relative rate and EFCI marking modes)
- ATM Forum LAN Emulation Release 1.0 Protocol

Internet Protocols

The LightStream 1010 ATM switch uses the following standard internet protocols:

- RFC 1483: multiprotocol encapsulation over ATM AAL 5
- RFC 1577: classical IP and ARP over ATM AAL 5
- ARP—used to determine the destination MAC address of a host using its known IP address
- BOOTP—uses connectionless transport layer User Datagram Protocol (UDP). BOOTP allows the switch (BOOTP client) to get its IP address from a BOOTP server.
- Internet Control Message Protocol (ICMP)—allows hosts to send error or control messages to other hosts. ICMP is a required part of IP. For example, the **ping** command uses ICMP echo requests to test if a destination is alive and reachable
- Internet Protocol (IP) over ATM—suite used to send IP datagram packets between nodes on the Internet

MIBs Supported

- TCP—a reliable, full-duplex, connection-oriented end-to-end transport protocol running on top of IP. For example, the Telnet protocol uses the TCP/IP suite
- Packet Internet groper (ping)—used to test the accessibility of a remote site by sending it an ICMP echo request and waiting for a reply
- TFTP—used for downloading network software updates and configuration files (flashcode) to workgroup switch products
- RARP—used to determine an IP address when only a MAC address is known; for example, BOOTP and RARP broadcast requests are used to get IP addresses from a BOOTP or RARP server
- SLIP—a version of IP that runs over serial links, allowing IP communications over the administrative interface
- PPP—provides host-to-network and switch-to-switch connections over synchronous and asynchronous circuits
- SNMP—agents that process requests for network management stations and report exception conditions when they occur; this setup requires access to information stored in a MIB (for more information, refer to the following section, “MIBs Supported”)
- Telnet—a terminal emulation protocol that allows remote access to the administrative interface of a switch over the network (in-band)
- UDP—enables an application (such as an SNMP agent) on one system to send a datagram to an application (a network management station using SNMP) on another system. UDP uses IP to deliver datagrams. UDP/IP protocol suites are used by Trivial File Transfer Protocol (TFTP)

MIBs Supported

The LightStream 1010 ATM switch supports standard and enterprise-specific MIBs. The following MIBs are supported:

- RFC 1573 (interfaces MIB)
- Cisco proprietary MIBs. (Refer to the appendix, “Workgroup MIB Reference” of the *LightStream 1010 ATM Switch Software Configuration Guide* for a description of the MIBs offered.)

- MIB II (system, IP, ICMP, TCP, UDP, SNMP) (RFC 1213)
- AToM MIB (RFC 1695)
 - ATM interface configuration
 - ATM interface DS3 physical layer convergence procedure (PLCP)
 - ATM interface TC sublayer
 - ATM interface virtual path link/virtual channel link (VPL/VCL) configuration
 - ATM VP/VC cross-connect
 - AAL5 connection performance statistics
- LAN Emulation Client (LEC) MIB (draft from AToM MIB)
- Signaling MIB (draft from AToM MIB or AToM WG)
- SONET MIB (subset of RFC 1595)
- DS-3 MIB (subset of RFC 1407)
- Cisco proprietary MIBs. (Refer to the appendix “Workgroup MIB Reference” of the *LightStream 1010 ATM Switch Software Configuration Guide* for a description of the MIBs offered.)
 - Per-VC AAL5 protocol data unit (PDU) counters
 - OAM MIB
 - PNNI Version 1 routing MIB (subset of MIB is used for IISP)
 - Misc (Address of LECS, per-VC cell counters)
 - Object corresponding to all other CLI commands
 - LightStream 1010 specific MIB extensions

MIBs Supported
