Campus ATM Network Design

Keerli Melkote, WBU Technical Marketing
kmelkote@cisco.com

Table of Contents

Cisco Systems

Introduction

As switched internetworking is taking over as the new wave in campus networks, Cisco Systems has gained a leadership position in implementing switched campus networks using leading-edge products and technology. Cisco's switched internetworking solutions portfolio includes industry-leading LAN switching and Asynchronous Transfer Mode (ATM) switching with corresponding virtual LAN (VLAN) solutions appropriate for each technology. These VLAN solutions include VLAN multiplexing technologies such as LAN Emulation (LANE) over ATM networks, Inter-Switch Link (ISL) over Fast Ethernet, and 802.10 over Fiber Distributed Data Interface (FDDI). The leading choices for implementing a VLAN-based campus architecture are ISL over Fast Ethernet and ATM LANE. Making the choice is confusing, especially when the solution that better solves the customer requirements is unclear. This paper gives a clearer understanding of the ATM LANE solution set and its benefits.

Scalable bandwidth, quality of service (QoS) guarantees, and traffic management have all been named as factors that point to the need for an ATM infrastructure for today's evolving corporate networking needs. As deployment of ATM networks continues, network designs involving LANE and Private Network-to-Network Interface (PNNI) have become central to most campus implementations of ATM. These standards will continue to play a very important role in evolving ATM networks to accommodate upcoming standards such as Multiprotocol Over ATM (MPOA). Hence an in-depth understanding of how these components work is a requirement to successfully building and operating an ATM network.

LANE is a key enabling technology in migrating legacy networks to ATM, and it has been successfully deployed in production workgroup and campus networks. With LANE, customers can run existing LAN-based applications and broadcast-oriented LAN protocols over ATM. PNNI is very important in scaling the ATM switch network, and the Cisco implementation of PNNI has been used to successfully build production networks consisting of more than 90 ATM switches.

This Design and Implementation guide provides design guidelines and configuration tips for building scalable LANE networks. It will discuss in detail, Cisco's products and their implementation of the above standards as well as provide guidelines for using these products in the form of network designs. Where applicable, configuration guidelines and sample configurations are also explained. Sections 2-6 explore the design and scalability issues for LANE networks from the workgroup and campus network perspectives. Section 7 discusses some of the configuration details for building LANE networks using Cisco products while Section 8 discusses the LANE redundancy schemes like Simple Server Redundancy Protocol (SSRP) and Hot Standby Routing Protocol (HSRP) over LANE .

The Move to Switching

The demands for increased bandwidth, easier administration, and reduced cost of ownership are causing customers to migrate to switched networks. As customers articulate their requirements for these switched networks, it is becoming clear that the requirements are common to numerous customers. These requirements can be broadly classified as:

- Increased network bandwidth—The main factor that drives the need for higher bandwidth in campus networks is the growth in the number of users and newer applications. Newer applications drive the need for more powerful servers and desktops that, in turn, require higher network bandwidth. Moreover, information sharing and faster information dissemination are becoming more of a necessity in the corporation. Intranet technologies and applications are meeting these requirements and are taking over as the next wave in campus networks. Thus the need for network bandwidth is increasing.
- Server location—Although the distributed server architecture, that is, locating the server physically close to the users, was a very popular concept, the high cost of ownership associated with this architecture is driving the need to physically centralize the servers in one location, or the data center. This scenario, in turn, means that most of the client/server traffic has to traverse the trunks or the backbone to get to the central location instead of being localized to a wiring closet. As will be shown, this challenge is possibly the single biggest one in scaling a campus network, and it will have a great impact on network design and implementation.
- Adds, moves, and changes—As user communities grow or reorganize, the need for cost-effective adds, moves, and changes has been and will be one of the most important requirements for many customers. From a networking perspective, this requirement translates itself primarily into a network address and management issue, and the network architecture must be designed to lower the cost of adds, moves, and changes.
- Security—Although security is less of an issue in the campus than in a wide-area network (WAN), the need to protect certain types of data (such as payroll information) from unauthorized users still exists in campus networks. Another example is the necessity to form secure user groups.

Many different solutions, ranging from ATM switching with LANE to pure frame switching solutions such as switched Ethernet, Fast Ethernet, and FDDI, can be proposed to solve these requirements. The following sections discuss the ATM solution with LANE and PNNI in the context of a campus network model. Where applicable, the solution is compared with a corresponding frame switched solution. In particular, scalability issues are considered.

Benefits of LANE with PNNI

LANE Version 1.0 defines the standards for internetworking legacy LANs such as Ethernet and Token Ring with ATM-attached devices. Such devices include end stations (for example, ATM-attached servers), edge devices that bridge legacy LANs onto an ATM backbone (for example, Catalyst™ 5000 switches), and ATM-attached routers to route between emulated LANs (ELANs). LANE is defined as a Media Access Control (MAC) encapsulation (Layer 2), since this approach supports the largest number of existing Layer 3 protocols. The end result is that all devices attached to an ELAN appear to be on one bridged segment. In this way, IP, IPX, AppleTalk, and other protocols should have similar performance characteristics as in a traditional bridged environment.

Figure 1 illustrates a reference LANE network and includes the components needed to implement a LANE network.

Figure 1 Components in an ATM Network Running LANE

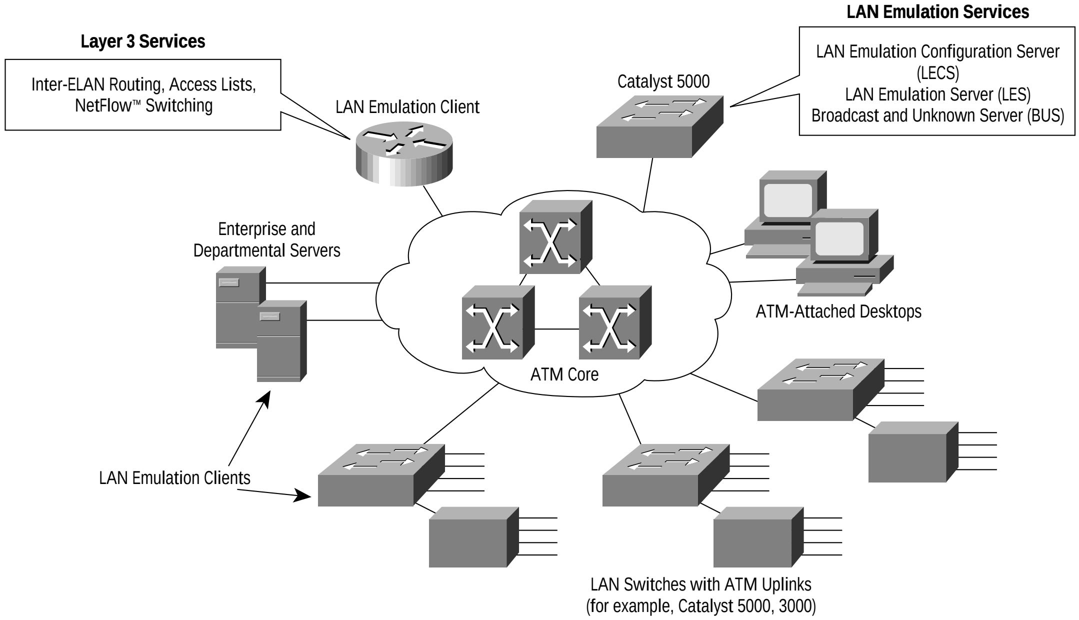

The challenges laid out in the previous section can be addressed using multiple technologies. The following section gives the main benefits of using an ATM solution in the campus.

The VLAN Solution—Mobility, Broadcast Management, and Security

VLANs provide the ability to logically segment the network into multiple groups with similar user populations. For example, all the Marketing users on the network can be thought of as being part of a single, logical workgroup with similar requirements and services. This ability lends itself nicely to implementing a solution that eases mobility (moves and changes). A port-based VLAN solution[1] also provides broadcast separation from other VLANs in the network, not only addressing a security concern but also achieving better scalability through broadcast containment. Broadcast domain design issues are covered in section 5.1.3.

Edge devices such as the Catalyst 5000 enable the user to bridge legacy LANs such as Ethernet onto an ATM backbone and emulate not just one, but multiple ELANs on the same physical interface, in turn, enabling the user to multiplex traffic from multiple LAN segments or broadcast domains onto the ATM backbone. This ability lends itself very well to implementing a VLAN architecture. Other VLAN architectures are based on frame multiplexing technologies such as ISL for use over Fast Ethernet and 802.10 over FDDI.

A VLAN based on ISL, 802.10, or LANE is essentially a logical broadcast domain, as opposed to what was earlier a physically constrained broadcast domain. Therefore, the scalability properties of VLANs based on any of these VLAN multiplexing technologies are very similar. Questions such as how large a VLAN or ELAN can be are answered in terms of scaling a broadcast domain, and the answers are very similar across all the multiplexing technologies.

Scaling the Backbone—PNNI versus Spanning Tree

The main difference between ISL and 802.10-based VLAN networks over a LANE network is in the core or the backbone. While frame multiplexing uses frame switches in the core of the network, LANE uses an ATM switch core. Frame switch networks are primarily bridged networks, and their scalability is limited by the Spanning-Tree Protocol (STP). An ATM core, on the other

---

[1] In port-based VLANs, each user port on a LAN switch can belong to a single VLAN. Trunk ports that are used to aggregate multiple VLANs belong to multiple VLANs. VLANs can be port based, MAC address based, Layer 3 based, and so on. All current Cisco solutions utilize the port-based VLAN approach.

hand, is based on PNNI, which is an ATM routing and signaling protocol and can be thought of as a Layer 2 routing protocol in the context of the Open System Interconnection (OSI) seven-layer model. The PNNI routing protocol is a link-state protocol like Open Shortest Path First (OSPF) except that PNNI is QoS aware. In other words, it has the ability to form topologies for different traffic classes (constant bit rate [CBR], variable bit rate [VBR], and so on) and route calls (virtual circuits [VCs]) across the network based on the traffic class.

Using PNNI in the core enables distribution of the VCs across multiple paths and links, thus providing redundancy and load distribution across the backbone. Network reconvergence times after failures are much faster than in a traditional network based on the Spanning-Tree Protocol. The details are discussed in the section on scaling the campus network.

LANE networks have frame switches at the edges and, therefore, the scalability of the Spanning-Tree Protocol also needs to be considered. However, this limitation is more of a VLAN scalability limitation (how many frame switches a single VLAN can touch before becoming too large), as opposed to a backbone scalability limitation.

The MPOA Factor—Solving the Server Location Issue

In current switched VLAN network architectures, routing is handled in the traditional way in that all the inter-VLAN traffic has to go through a routing process (typically a router). The only difference from non-VLAN architectures is that the router also understands the VLAN multiplexing technology, such as ISL or LANE, and there can participate in multiple VLANs via a single physical interface.[2]

Because of the limited bandwidth available on traditional routers, standard network design recommendations followed the 80 20 rule; that is, keep 80 percent of the traffic local and 20 percent across the router for inter-VLAN traffic. In terms of client/server applications, this setup meant that the server had to be placed on the same segment as the client for optimal performance. Initially this setup translated into a requirement to distribute the servers to where the clients are located since the concept of VLANs did not exist. With VLANs, it was possible to physically centralize servers in data centers and still retain the 0 router hop design in a logical fashion by keeping the clients and servers in the same VLAN. Although this setup solved a major cost issue of managing the servers centrally, it does not address the scalability issue. In other words, VLANs are still limited in size and, for servers that need to be accessed by large user populations (Web servers, e-mail, and so on), there is no optimal path. It is possible to multihome servers to different VLANs, but even that scenario does not scale to large networks. Therefore, the current design recommendation is to keep these servers one hop away since this setup minimizes the server configuration complexity and provides equal access to all users of the network.

The ATM Forum is working on a new standard known as Multiprotocol Over ATM (MPOA), which provides the ability to switch inter-VLAN traffic across the ATM backbone without having to traverse the router. This technique, known as inter-VLAN cut-through routing, provides much better performance over traditional routing, and it is the appropriate technology to use for Web servers and other services that are accessed by large user groups.

MPOA will continue to utilize LANE as the intra-VLAN or bridging standard. Therefore, implementing a LANE network is a prerequisite to MPOA; and LANE, together with ATM-attached routers, will be used to provide routing until MPOA achieves standardization and implementations achieve the stability necessary to be deployed in production networks.

Traffic Engineering—Class of Service Provisioning in the Campus

While class of service (CoS) provisioning has been a big point of discussion when considering the benefits of ATM, it has not been possible thus far to implement such provisioning in the campus network because of implementation limitations on end stations as well as the limitations of current standards. As an example, LANE 1.0 supports only unspecified bit rate (UBR), or best effort calls, and therefore, it is not possible to take advantage of the CoS support provided in the ATM switches and the PNNI protocol. This scenario typically translates itself into servers or end stations that are not able to request a certain CoS for a given traffic flow. To summarize, implementing CoS thus far in the campus needed not only an evolution in the existing LANE standard but also support in the servers.

---

[2] Subinterfaces are used in multiplexing multiple VLANs or logical interfaces over a single physical interface. Each VLAN or ELAN is then attached to a subinterface.

However, customers are requesting the ability to implement CoS guarantees for certain server flows on a network-wide basis without having to upgrade the software in their existing end stations or servers. Cisco's approach to solving this issue is to enable specialized traffic engineering schemes on the ATM switches, thus enabling the implementation of CoS guarantees for server flows or connections. The end stations can continue to exist on current LANE technology and take advantage of the schemes described previously.

A Logical Model for a Campus Network

A campus network can be thought of as multiple workgroups that can be logically separated from each other according to some administrative requirement specified by the network administrator. In a corporation, these workgroups can be divided into functional units such as marketing, engineering, finance, and so on, or by geographical boundaries such as a floor or a wiring closet. Each functional unit or workgroup has its own unique applications and services, which are commonly provided by their own servers. Therefore, Engineering has a certain set of users with their own needs for a file server. Network access to these local [3] servers needs to be very fast and without any bottlenecks since any problems accessing these services adversely impacts the day-to-day functioning of these workgroups.

Figure 2 Campus Network—Logical Model

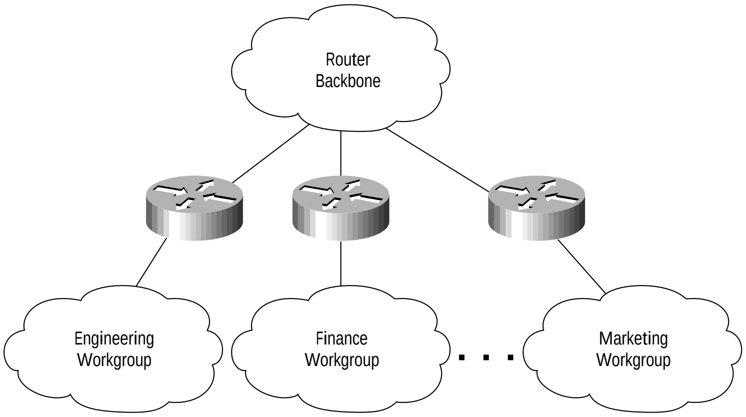

In addition, there are also corporate servers that are common across all the workgroups and need equal access to all workgroups. If these workgroups are completely independent of each other and require no communication between each other, they can be divided into separate logical sections of the network that do not talk to each other. However, this scenario is rare, and most corporations do require some communication across workgroups. Also, some corporations require some amount of security for cross-workgroup communication. An example for such a requirement is the need to allow access for e-mail and Web-based communication but not allow access to any data on the file servers.

This model employs the use of a mix of Layer 2/Layer 3 communication for intraworkgroup traffic while using a pure Layer 3 approach for interworkgroup communication. The rationale for this model arises from the need for fast communication for intraworkgroup traffic, which is typically provided by Layer 2 switches. The need for Layer 3 within the workgroup typically arises when the workgroup is too large to be a one-Layer 3 subnetwork. Security and firewalling is well suited for hierarchical addressing schemes typically employed at Layer 3 rather than a flat, MAC-based addressing scheme typically employed at Layer 2 [4]. The illustration in Figure 2 represents a logical separation of the campus network using routers. The routers may be physically connected, as shown, or they may be connected as 'routers on a stick.'

---

[3] The term "local" is used to mean "local to the workgroup" and not a global or central server.

[4] ATM is a notable exception to the norm of flat Layer 2 addressing schemes. ATM has a hierarchical addressing scheme and a routing protocol (PNNI) that can take advantage of this fact by summarizing addresses in routing updates.

Users within a workgroup do not have to be geographically together; they may be located throughout the campus. Therefore, when users move from one location to another within the campus, they need to maintain similar access to servers and services if part of the same workgroup.

This model is hierarchical in that multiple workgroups are interconnected using routers, while the basic functional unit, which is the workgroup network, is replicated multiple times. Therefore, the scalability of these workgroup networks is addressed in the following section. In the discussion, the workgroup networks are separated by routers at the core.

Designing the Workgroup Network

A workgroup network can be thought of as consisting of a set of users with common servers where shared workgroup data resides. Intraworkgroup communication is mostly client/server, with some amount of peer-to-peer communication, that is, client/client (for example, whiteboarding) and server/server (for example, backups). Although we will try to be fair to all these different communication modes, we shall attempt to optimize client-server communication where it is not possible to be fair to all the modes.

A Flat, Switched Network

The simplest approach to implementing a workgroup network is to use a single broadcast domain or VLAN for the entire workgroup. This network consists of multiple switches that participate in a single spanning-tree domain, as shown in Figure 3.

Figure 3 Flat Workgroup Network

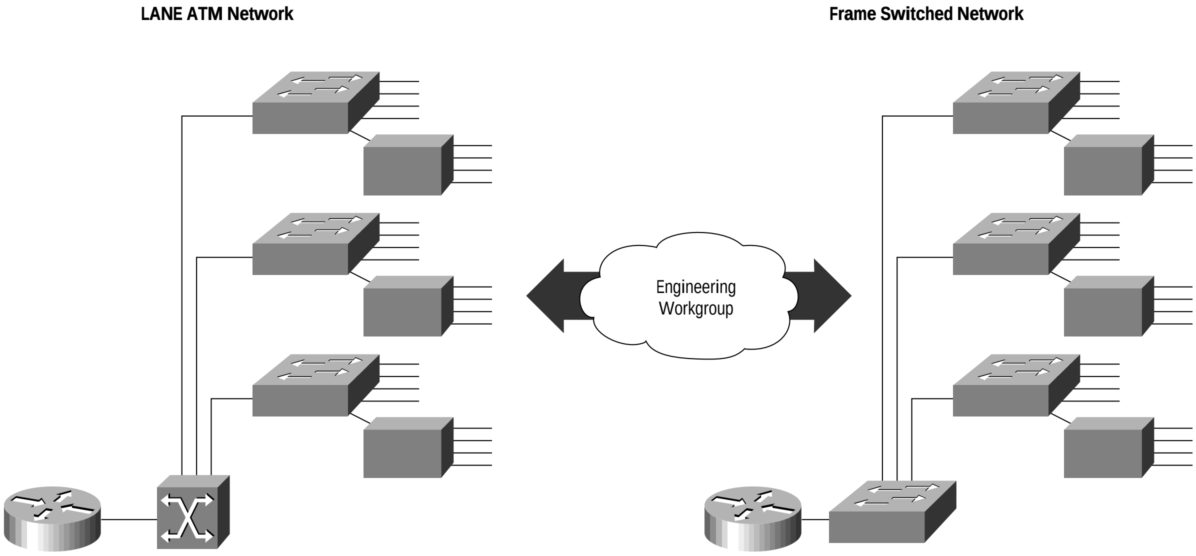

How Do VLANs Fit In?

A flat network has excellent mobility characteristics since every user belongs to a single broadcast domain and, as long as the port that the users are attached to, belongs to this broadcast domain, access to their workgroup remains the same no matter where the users are physically located. The implementation issue comes in when the need arises to extend a broadcast domain to a new user location. Installation of new switch just for the sake of a single user is cost-prohibitive. Therefore, using a switched VLAN architecture with some form of a multiplexing technology to extend the VLANs to the different locations is much more cost-effective.

A flat, switched network has scalability limitations, which are discussed in the following section.

Addressing Limitations—Practical Issues

Most IP networks are either class B with a class C (eight-bit) subnet mask or multiple class C networks. This translates into a limitation of 254 end stations per subnet[5]. A subnet also maps into a broadcast domain, and hence the subnet mask becomes the addressing limit for extending a broadcast domain. Although configurations where multiple subnets are implemented on a single broadcast domain exist, the administrator has to add a route to the other networks on the broadcast domain at every end station in order to make the configuration work effectively. This scenario obviously is an administrative nightmare; a more scalable solution is to implement a single subnet per broadcast domain and use routing for intersubnet communication.

Although this configuration may appear to be an artificial limitation imposed by the addressing scheme used, it is not. For example, moving the subnet mask boundary two bits to the left (that is, a ten-bit subnet mask) results in 1024 end stations per subnet, but the number of subnets is vastly reduced to 64. Moreover, not all workgroups are the same size, and 1024 end stations per subnet is excessive for most workgroups. It can be argued that a variable-length subnet mask (VLSM) strategy can be employed to accommodate the variability in the workgroup sizes, but this strategy not only creates an administrative nightmare, but it also requires changing of the routing protocol from Routing Information Protocol (RIP) or Interior Gateway Routing Protocol (IGRP®) to Enhanced IGRP or OSPF. These changes may seem excessive for a simple benefit of having more users in a single broadcast domain, especially if other solutions can be used.

Another, much more practical, problem is that any change in the addressing strategy requires touching ALL the end stations since their subnet masks need to be changed, a task easier said than done.

For Internetwork Packet Exchange (IPX) and AppleTalk networks, the addressing limitation is not so severe since these networks take advantage of dynamic addressing. However, IPX and AppleTalk networks are much more broadcast intensive, representing a much more serious issue in scaling a broadcast domain.

Broadcast Radiation

The term broadcast domain comes from the fact that broadcasts and multicasts such as RIP updates, Service Advertisement Protocol (SAP) updates, and so on go throughout the domain. These broadcast packets must be processed by all end stations and then thrown away if they have no need for these packets. This scenario is often referred to as broadcast radiation, and it steals CPU cycles from the end stations. Hence, minimizing the impact of broadcast radiation on the ELAN is important. Therefore, the size of an ELAN really depends on the amount of broadcast radiation and the broadcast processing capabilities of the end stations. The broadcast radiation, in turn, depends on the protocols and type of applications that are running in the network.

Studies indicate that there are upper limits to the size of an ELAN; they are quantified in Table 1.

Table 1 Scaling a Broadcast Domain—Number of End Stations per Broadcast Domain

Protocol Type	Number of End Stations
IP	500 (254 is a practical limit if using a class C subnet mask)
IPX	300
AppleTalk	200
Mixed	200

In addition to these issues, there are other scalability issues, such as the Spanning-Tree Protocol, that also play a role in determining the size of the broadcast domain. But this issue has much more significance, as discussed in section on scaling the campus network.

The scalability limitations of broadcast domains require implementation of a workgroup network across multiple broadcast domains.

---

[5] Class C networks are limited to 254 end-station addresses per network address, hence the need for multiples of these addresses for medium (1000 to 2000 nodes) to large (>2000 nodes) campus networks.

Scaling the Workgroup Network

The addressing and broadcast radiation limitation can be solved by extending the workgroup network across multiple broadcast domains. This setup can be accomplished by implementing a separate ELAN or VLAN for each subnet, but it presents some other challenges for the network administrator.

Server Location—Multihoming

In the flat network model for a workgroup, the server can be colocated (0 router hops) with the users so that access to the server is across a Layer 2 path. When a workgroup network needs to be extended across multiple broadcast domains, the server is the optimal (Layer 2) path for some users while it is one hop away for other users. This can be solved by multihoming the servers onto each broadcast domain that the workgroup uses.

Multihoming can be achieved either logically or physically. Physical multihoming requires the server to have one network interface card (NIC) for each subnet that the server needs to be a part of. This approach can be useful when there are a few (two to four) subnets and each subnet requires the throughput of a single NIC. Logical multihoming implies that the server can use a single NIC but participate in multiple subnets. This concept is very similar to that of a router on a stick, where a single physical interface consists of multiple subinterfaces.

Figure 4 Physical versus Logical Multihoming of Servers

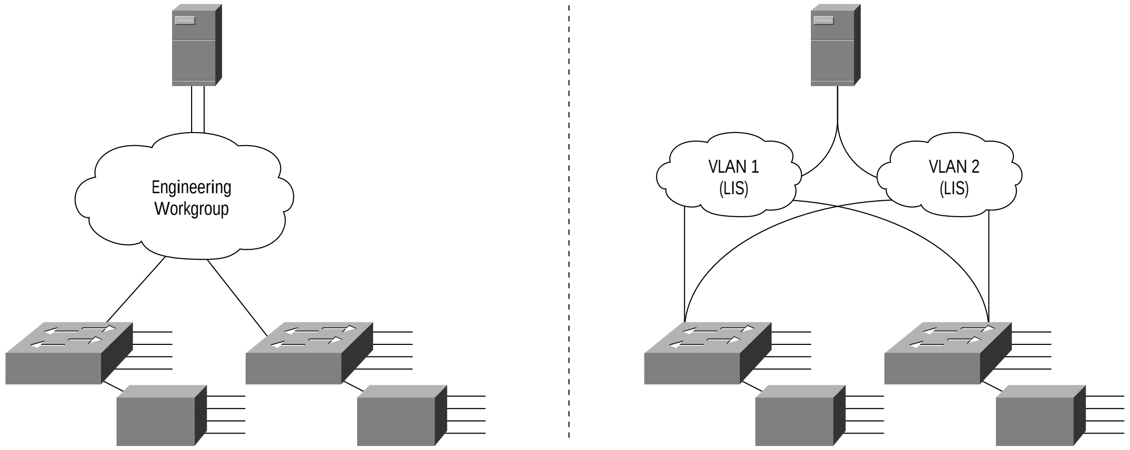

Logical multihoming requires that the NIC understand the VLAN multiplexing technology being used in the backbone. Therefore, in the case of ATM, the NIC should be LANE-compliant, while in the case of ISL, the NIC should be ISL-capable. In terms of support for a diversity of platforms, there are many more options for LANE than for ISL. Table 2 gives a partial list of the platforms supported for LANE and ISL.

Although most NICs allow multiplexing between 8 and 256 ELANs or VLANs in their marketing literature, 4 to 162 ELANs are often more realistic numbers, and they are also more than enough for workgroup environments. The exact number depends on the NIC, the server type, and the applications that run on the server. Although the scalability of servers and NICs themselves is beyond the scope of this document, the following are issues to address when designing with mutihomed servers:

- Server memory vs. onboard NIC memory—A distinguishing feature between NICs is how much onboard memory is available to store packets. Most NICs have enough memory to store a maximum-sized packet that is received on the NIC and then immediately copied (direct memory access [DMA]) to host memory. Some NICs have extra memory to store multiple packets before they are copied (by DMA) to host memory. The number of DMA requests is directly proportional to the number of interrupts sent to the CPU, which, in turn, impacts the CPU utilization. Therefore, having slightly more onboard memory, say two or three packets, which are all simultaneously copied to host memory, results in far fewer (one-half or one-third) interrupts, thus lowering the impact on the CPU. The CPU can be used for other important functions such as disk input/output (I/O), and so on.

- CPU and system type—Obviously the CPU and system type impact the kind of I/O processing the server can handle. For example, a SPARC server from Sun can do much better I/O processing than a PCI-based Pentium server.
- Switched virtual circuit (SVC) counts—NICs can support only a limited number of SVCs. Typical marketing numbers are in the range of 1024 to 2048 virtual channel connection (VCCs), but lower numbers (512 VCCs) are more realistic. Again, there is no definitive rule, but it is important to be sure that you are not running out of SVC capacity on the NICs.

Intersubnet Routing

Since multiple subnets are needed to implement the workgroup network, it is important to design in sufficient routing horsepower to accommodate the bandwidth requirements for intersubnet traffic. As mentioned earlier, in this case the client/server traffic can be optimized by providing local access to servers (by multihoming them to different ELANs) while having client/client communication go through the router. The amount of client/client traffic varies according to the amount of peer/peer applications in the workgroup, and the amount of routing horsepower can be determined by comparing different routers in the performance table.

Table 2 Unidirectional Routing Performance

Router	Cisco 7500 with ATM Interface Processor (AIP)	Cisco 7000 with AIP	Cisco 4700	Cisco 7500 with ATM Lite Port Adapter (PA)	Cisco 7200 with ATM Lite PA
Switching Mode					
Optimum Switching	47.5 kpps	—	—	74 kpps	76.2 kpps
Flow Switching	47.5 kpps	45.8 kpps	—	74 kpps	71.9 kpps
Fast Switching	47.5 kpps	13.4 kpps	34 kpps	76.1 kpps	76.3 kpps
Process Switching	2.7 kpps	1.1 kpps	4.5 kpps	2.9 kpps	4.7 kpps

Figure 5 Unidirectional router LANE throughput

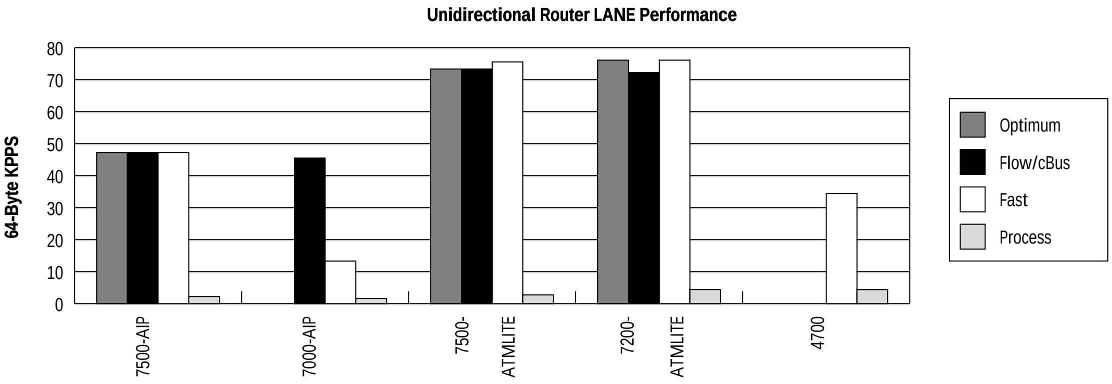

Table 3 Bi-directional Router LANE Performance

Router	7500 with AIP	7000 with AIP	4700	7500 with ATM Lite Port Adapter	7200 with ATM Lite Port Adapter
Switching Mode	—	—	—	—	—

Table 3 Bi-directional Router LANE Performance

Router	7500 with AIP	7000 with AIP	4700	7500 with ATM Lite Port Adapter	7200 with ATM Lite Port Adapter
Optimum switching	47.6 KPPS	—	—	81.6 KPPS	76.0 KPPS
Flow switching	47.1 KPPS	45.8 KPPS	—	81.6 KPPS	76.4 KPPS
Fast switching	47.6 KPPS	14.0 KPPS	34.4 KPPS	82.0 KPPS	76.4 KPPS
Process switching	3.0 KPPS	1.1 KPPS	4.6 KPPS	3.0 KPPS	4.9 KPPS

The performance levels needed here must be added to the amount of routing bandwidth needed to implement interworkgroup communication. Interworkgroup communication is discussed in a later section.

Scaling the Campus Network

The campus network consists of multiple workgroup networks interconnected by routers. Although this description is functionally sound, there are many options when it comes to implementing and scaling the campus network. The architecture of the network really depends on the requirements placed on the network. In the following sections some of these requirements are examined and a campus LANE network that can address these requirements is built.

A LANE network, as shown in Figure 1, consists not only of LAN Emulation Clients (LECs) and LANE services but also ATM switches possibly running PNNI, edge devices running the Spanning-Tree Protocol, and so on. Each of these mechanisms contributes (positively or negatively) toward scalability of a network. Hence, the scalability of a LANE network should be explored in conjunction with other components that are part of a LANE network but not necessarily a part of the LANE standard.

Router-Based Campus Architecture

Under this architecture, the workgroups in the campus network are logically and physically separated by routers. The servers are multihomed to their respective workgroup subnets and can be either centralized or distributed. This architecture is very similar to earlier routed networks with shared technology. The key difference is the jump in bandwidth provided by using LAN switching and ATM. Such a network is illustrated in Figure 7.

Figure 6 Bi-directional Router LANE Throughput

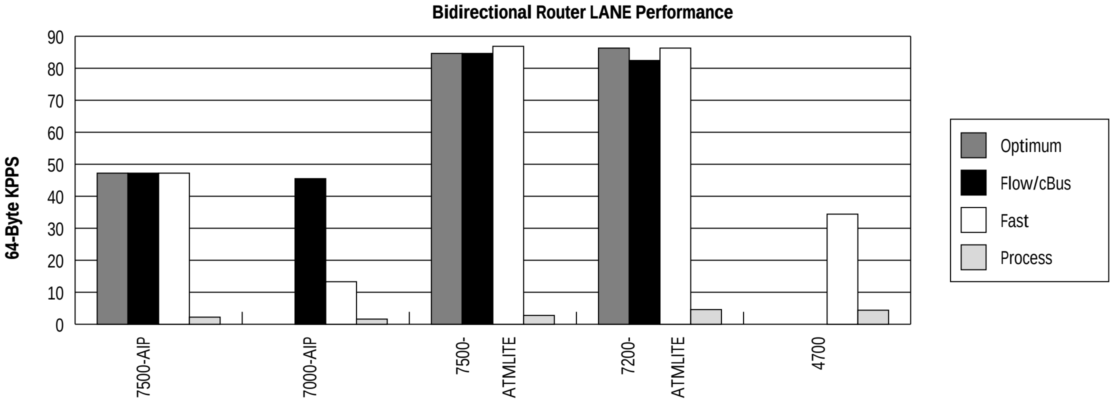

Figure 7 Traditional Campus Network Architecture

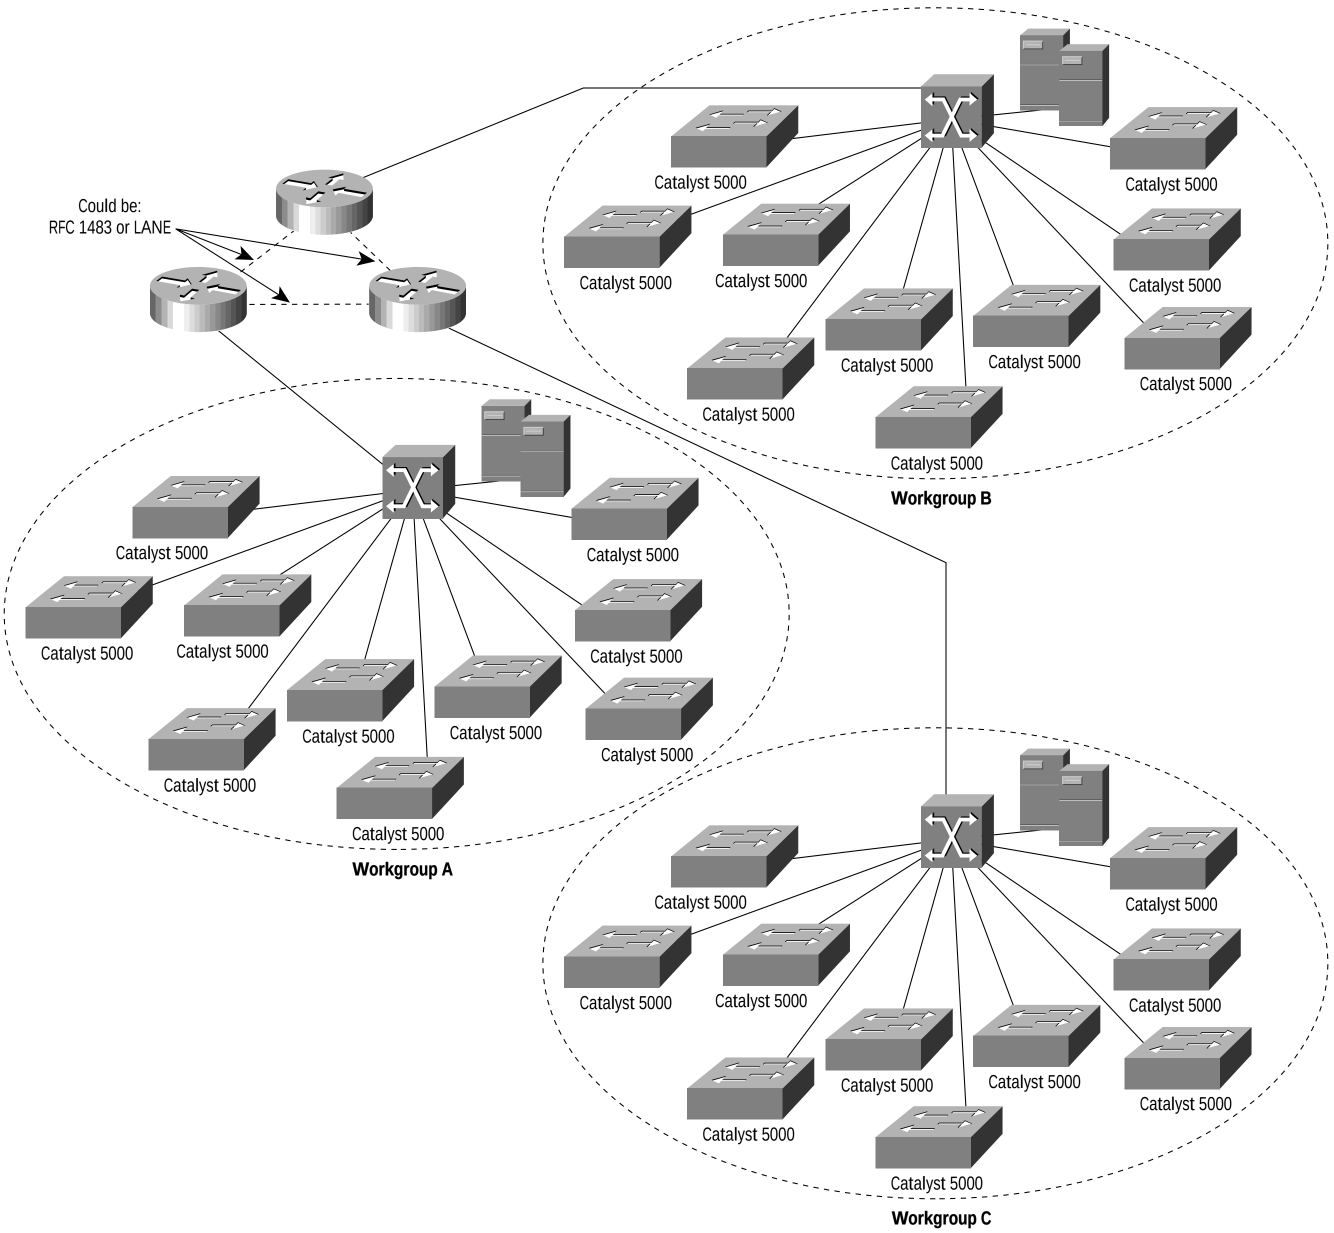

Although only a single ATM switch supports each workgroup in Figure 7, it could be multiple ATM switches. The point is, however, that the switches from different workgroups are not connected to each other. Therefore, the network consists of multiple, switched clouds that are logically and physically separated from each other. This network has good scalability properties since each workgroup network is a self-contained unit and is constrained in size to pose any serious scalability issues. The separation achieved by the routers turns this network into a routed network, and hence it achieves the same scalability and reliability that traditional routed networks have achieved. Another benefit of this architecture is its simplicity, which is critical for troubleshooting. However, these networks have the same issues as pure routed networks; the issues can be broadly summarized as follows:

- Cost—Since the workgroups are physically separated from each other, there are as many backbones as there are workgroups. Therefore, if multiple workgroups are located within a given building, then each one of these workgroup backbones needs to appear in that building. If the workgroups are geographically constrained to be part of a single building, then this scenario does not pose that much of a cost issue, but, if the workgroups are geographically dispersed throughout the campus, then cost does become an issue.
- Mobility—This issue is very closely linked to the cost issue in that, if a user needs to be located in a building where that workgroup does not currently exist, then the workgroup backbone needs to be extended to that building. Therefore, a move not only involves moving the physical workstation but adding switches and possibly fiber to accommodate the workgroup in the new location. These problems can be solved by using a campus-wide VLAN architecture, which is discussed in the next section.
- Centralized servers—Access to centralized servers that are accessed across all workgroups has to be through the routers. Examples of such servers are e-mail, meeting maker, Web servers, and so on. Depending on the amount of such traffic, this scenario may or may not pose a problem, in terms of both router processing capabilities and cost.

Campus-Wide Layer 2 Backbone

This architecture solves the cost and mobility issues of the earlier campus network architecture. It provides a campus-wide Layer 2 backbone that essentially provides a Layer 2 path between any two points in the network. Thus any workgroup VLAN can appear in any wiring closet, making the adds, moves, and changes issue much more manageable.

Figure 8 Campus-Wide VLANs

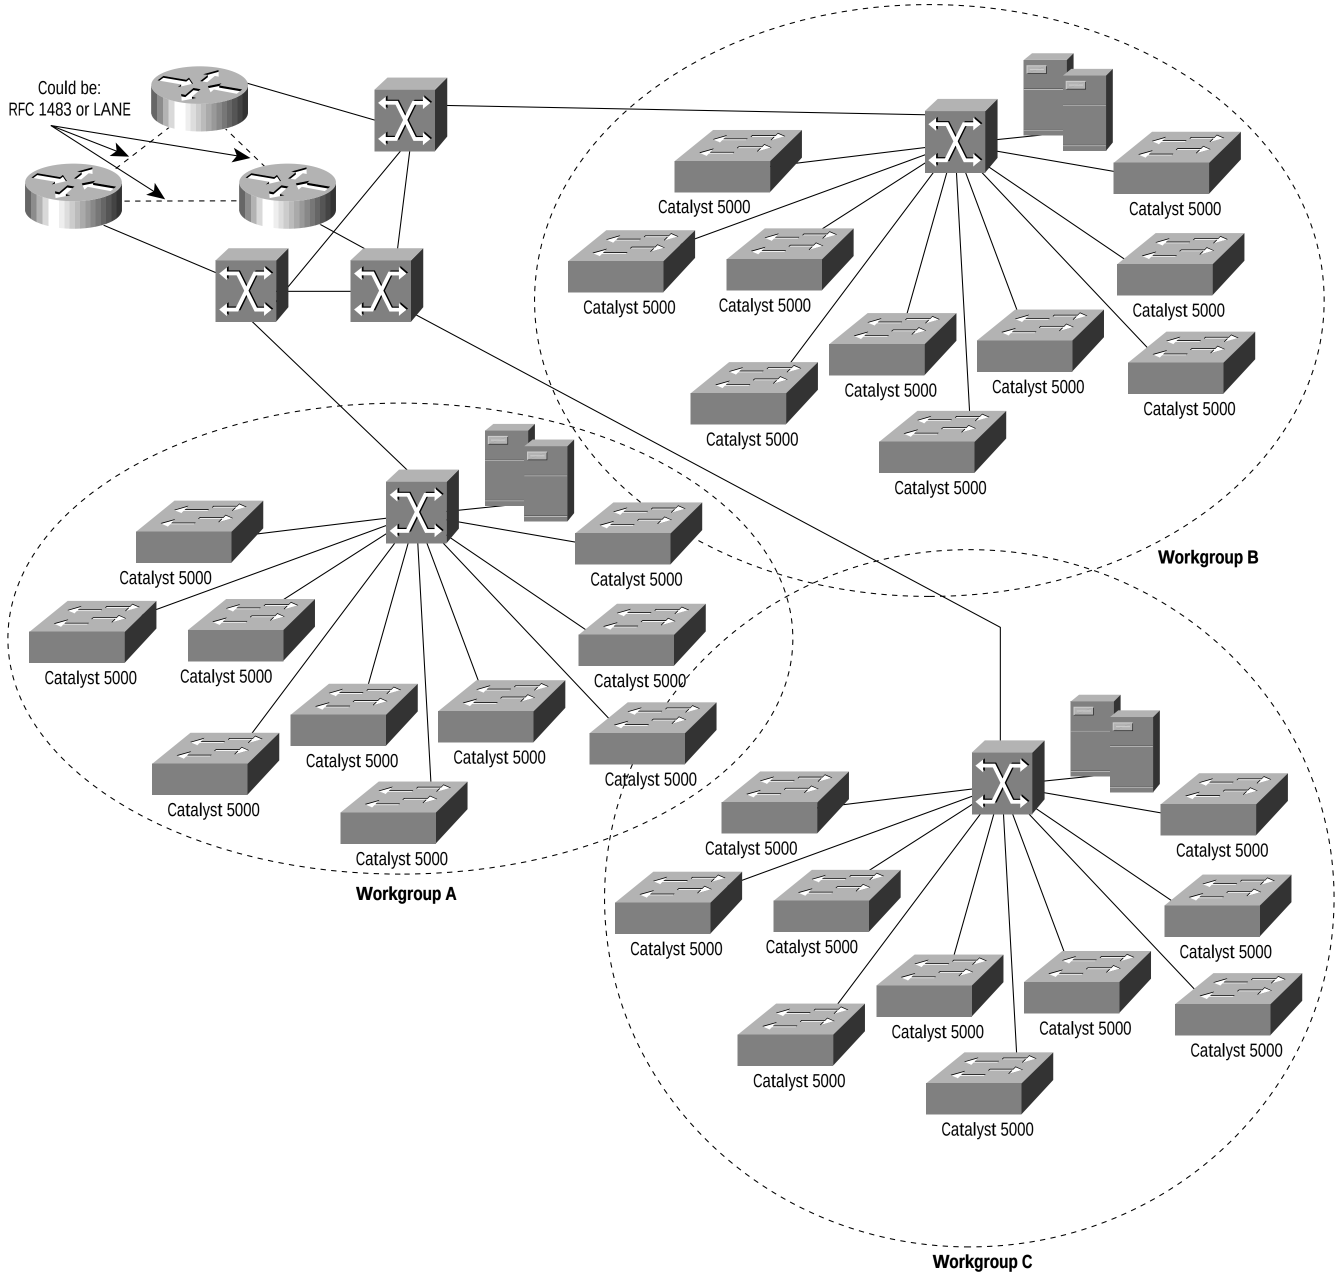

The scalability of such a campus-wide Layer 2 backbone is determined by the technologies used to construct the backbone. It is possible to construct an ATM backbone as well as a frame switch-based backbone with each solution having very different scalability properties.

A campus-wide ATM Layer 2 architecture typically employs edge devices such as the Catalyst 5000s in the wiring closets and ATM switching in the core. The ATM core uses PNNI as the Layer 2 routing protocol, providing a much better way to scale the network than using a frame switch backbone whose scalability is limited by the Spanning-Tree Protocol.

The following sections explore, in detail, the advantages of using the PNNI protocol as a Layer 2 routing protocol and the issues surrounding SVC budgets while also evaluating CPU budgets, BUS performance, and Spanning-Tree Protocol in the context of LANE.

PNNI is an ATM routing protocol used for routing call setups[6] and is implemented in the ATM switches. Most LANE networks consist of multiple ATM switches and typically employ the PNNI protocol. Although PNNI is a very advanced routing protocol and supports QoS-based routing, this particular aspect of PNNI is not discussed here since most LANE networks are based on the best effort traffic category. However, the LightStream® 1010 has some useful PNNI-related features that can be used very effectively to scale LANE networks. This section deals with LANE-related PNNI subjects but does not address the intricacies of PNNI routing.

PNNI as implemented on the LightStream 1010 gives the ability to:

- Load balance call setup requests across multiple, equal cost paths between two end stations
- Load balance call setups across multiple parallel links
- Provide link and path redundancy with fast convergence
- Provide excellent call setup performance across multiple hops using the background routing feature

Figure 9 Load-Balancing Calls across Multiple Paths and Multiple Links

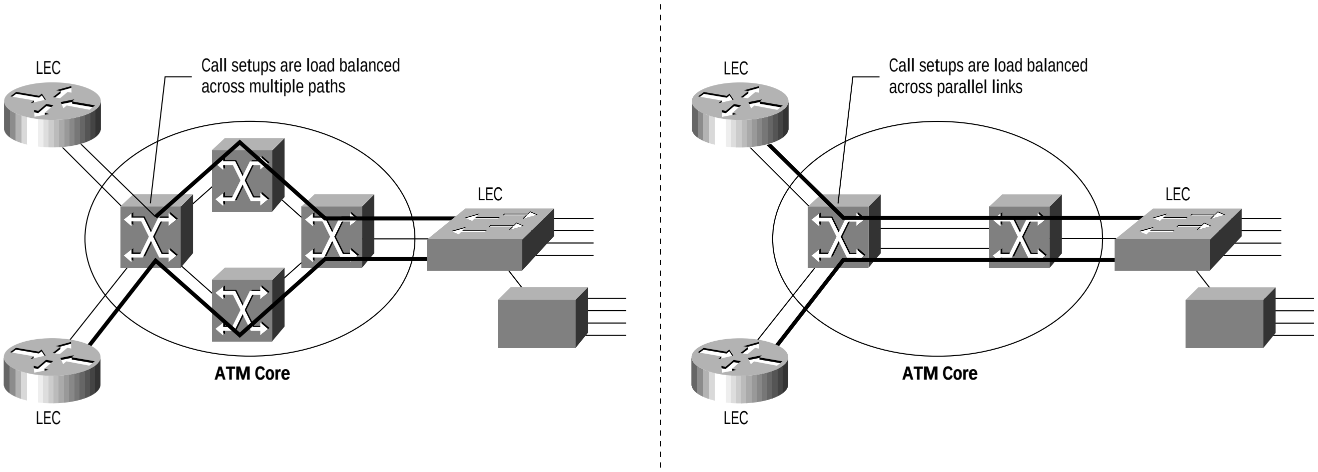

Load balancing of calls, as illustrated in Figure 9, is enabled by default on the LightStream 1010s. However, background routing is not enabled by default. Background routing can be thought of as routing of call setups using a path from a precomputed route database. The background routing process computes a list of all possible paths to all destinations across all the service categories (CBR, VBR-real time [RT], VBR-non-real time [NRT] and available bit rate [ABR]-UBR). When a call is placed from point A to point B, PNNI chooses a cache routed from the background route table instead of computing a route on demand. This scenario eases the load on the CPU and provides for a faster rate of processing the call setups. Background routing can be a very useful feature to employ in networks where the topology with respect to QoS is stable. It is, however, not very useful in networks with rapidly changing topologies, such as Internet service providers (ISPs) or carrier networks. Campus LANE networks can use this feature very effectively since all the SVCs in the network belong to the UBR or ABR category.

In order to enable this feature, use the following command:

```
atm router pnni
background-routes
```

---

[6] Edge devices such as the Catalyst 5000 and ATM end stations are connected to ATM switches across a User-Network Interface (UNI; 3.0/3.1/4.0) and only originate call setups. They do not route the call setups themselves.

The current implementation of PNNI on the LightStream 1010s does not support hierarchy and thus is expected to scale up to 100 to 200 nodes (ATM switches), depending on the interconnections. There are some real networks with more than 90 switches in production. PNNI with hierarchy theoretically supports an unlimited number of nodes. The practical limits are unknown at this time.

Scaling an ELAN—Spanning-Tree Protocol Issues

The Spanning-Tree Protocol is implemented in Layer 2 switches/bridges such as the Catalyst 5000 in order to prevent temporary loops in networks with redundant links. Therefore, it is frequently used as a mechanism to provide redundancy in. Since 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 Layer 2 switches. This setup, illustrated in Figure 10, is contrasted with an all-frame switched network.

As can be seen, the Spanning Tree Protocol (STP) topology of a LANE-based network is substantially simpler than a pure frame switched network employing the STP. However, STP convergence times, which can be a major issue in large frame switched networks, continue to be an issue in LANE networks. The gain in using LANE with PNNI, however, is the scalability of the Layer 2 cloud itself rather than the scalability of the ELAN itself.

Experiments have shown that convergence times lie in the range of 35 to 40 seconds with the default timer settings. Note that spanning tree needs to reconverge if there are failures at the edge devices or inside the ATM network.

Figure 10 Spanning-Tree Topologies for a LANE Network and a Pure Frame Switched Network

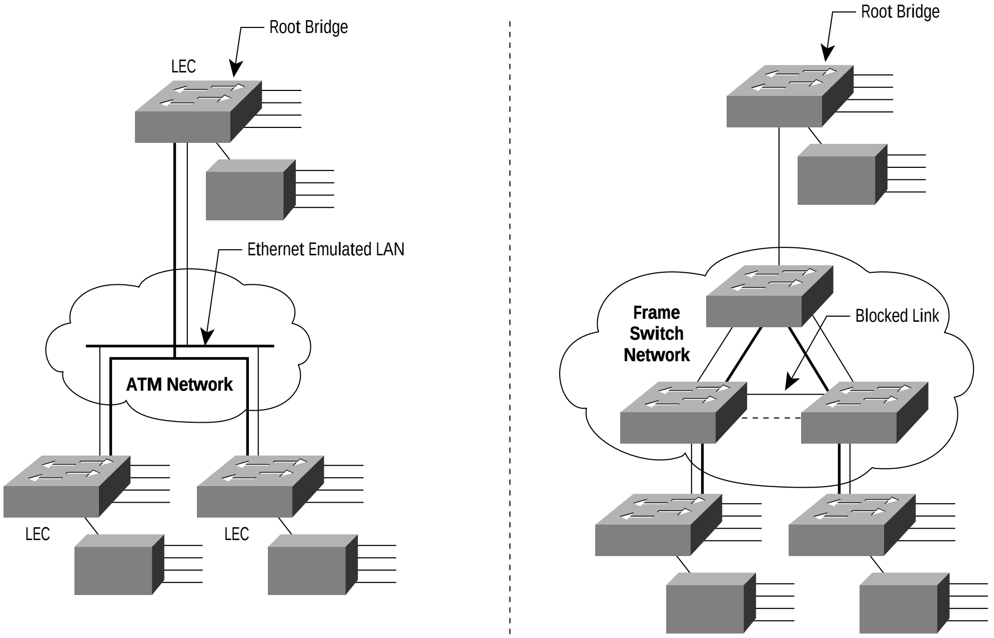

It is strongly advised to retain the default spanning-tree timers since they are known to work well for most networks.

However, if there is a need to tune the convergence time to something lower (or higher), the forward delay parameter can be used. The forward delay parameter controls the amount of time a switch places a port in the listening and learning states of STP when the port is transitioned into or out of the blocking state. The default value of the forward delay parameter is 15 seconds, which results in approximately 30- to 45-second recovery times. On the Catalyst 5000, the minimum forward delay setting is four seconds, bringing the recovery times down to about 15 to 20 seconds. The LANE client reestablishment times have typically been observed in the 0- to 5-second range. The variability in the convergence times comes from the type of network configuration and the number of ELANs that need to be recovered.

The command for adjusting the forward delay parameter on the Catalyst 5000s follows:

```
C5000(enable)> set spantree fwddelay 4 <vlan id>
```

BUS Bandwidth

In each ELAN, the BUS forwards all broadcasts, multicasts, and unknown unicast traffic. Hence, the packet processing capability of the BUS should be able to accommodate the total broadcast/multicast traffic in the ELAN. Additionally, the behavior of the BUS should be predictable when there is a broadcast storm that can overwhelm the BUS processing capability of the hardware/software that implements the BUS function.

Since there can be multiple ELANs in the network, it follows that the cumulative BUS processing capacity should be greater than the cumulative broadcast traffic in a given network.

The following table summarizes the BUS processing capabilities of each Cisco device implementing the BUS function.

Figure 11 BUS Performance

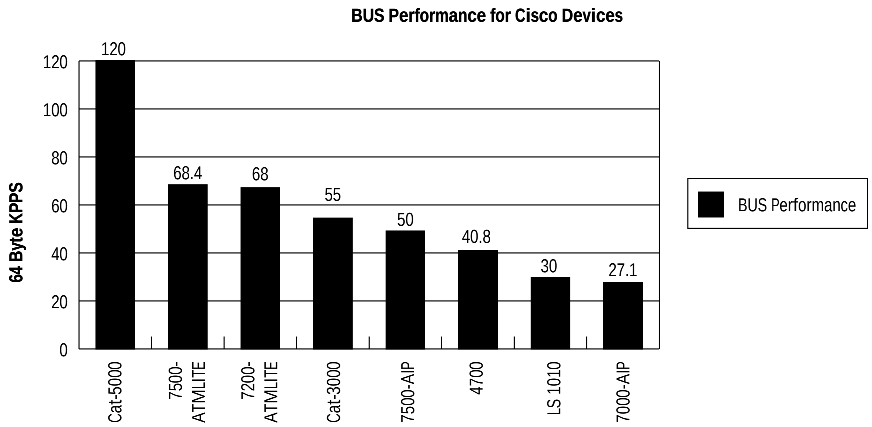

As can be seen from it is clear that the Catalyst-5000 platform has the best BUS performance while having minimal impact on the CPU utilization. All the other devices except the Catalyst 3000, take a hit on the CPU utilization. This is because the hardware on the LANE module for the Catalyst 5000 and 3000 has been optimized for the BUS and all packets destined to the BUS are handled in hardware. All the other platforms use the CPU for BUS traffic. Therefore, while implementing the BUS on these platforms you should keep the CPU budgets in mind.

A frequently asked question concerns the trade-off between implementing the BUS (or LANE services) in the ATM switch versus an edge device such as the Catalyst 5000 or a router. Although the box itself (that is, the ATM switch or an edge device) has little relevance to this discussion, what is more relevant is whether or not the hardware has been optimized to handle the LANE services. Most ATM switches, including the LightStream 1010, do not have specialized hardware for the LANE services. Hence, implementing LANE services on these switches means budgeting the CPU not only for PNNI, signaling, Interim Local Management Interface (ILMI), and so on, but also for LANE services. CPU budgeting is not an exact science, and the best one can do is to make an approximation. Therefore, implementing the LANE services on any platform that does not support it in hardware exposes the network to some amount of unpredictability during broadcast storms, and so on. To avoid any issues with respect to LANE services, Cisco recommends the Catalyst 5000 for use as a platform for LANE services.

LANE Services Placement

Placement of the LANE services in this way will impact the overall performance of the network. While the previous section discussed the issues surrounding the hardware platform chosen to run the LANE services, this section discusses the issues surrounding the centralization or distribution of LANE services.

Centralization of LANE services provides a convenient method of managing the LANE network since all the LANE services across all the ELANs are located on a single device. However, this setup is not very desirable from a fault-tolerance standpoint. Additionally, during failure recovery, all the control SVCs for LANE need to be established to a single ATM UNI port, thus impacting the amount of time necessary for the network to recover from a failure. Therefore, complete centralization is not recommended.

Distributing the LANE services across multiple devices provides for a reduced failure-caused impact on the network. It also increases the signaling bandwidth available for LANE services since they are now distributed across multiple UNIs. However, distributing the LANE services across the entire network poses a management burden and is not very desirable from a troubleshooting complexity standpoint.

Therefore, a hybrid of the centralized and distributed versions is the recommended strategy to employ in large networks. The hybrid solution consists of implementing the LANE services on a few Catalyst 5000s dedicated to the LANE services function. A set of four Catalyst 5000s can serve a large network in a redundant fashion (see section on LANE redundancy).

SVC Budgeting

A practical issue in scaling an ATM network is that of SVC capacities. An ATM network consisting of n edge devices needs a total of $n*(n-1)/2$ SVCs to achieve complete connectivity. This case is the worst case, and it is achieved when each device is talking to every other device at the same time. Fortunately, in real networks such a case rarely occurs and most conversations happen in a time-bound fashion such that SVCs are set up and torn down on demand. However, LANE networks need a minimum number of SVCs to function properly, and computing this number gives an idea of how many SVCs are available for on-demand data-direct SVCs.

In computing the SVC numbers, it is important to total not only the SVCs needed in the network but also the number of SVCs needed at each edge device. For example, the network shown in Figure 12 has ten edge devices attached to the ATM switch, and they are all communicating on a single ELAN.

Figure 12 SVC Requirements for a Single ELAN

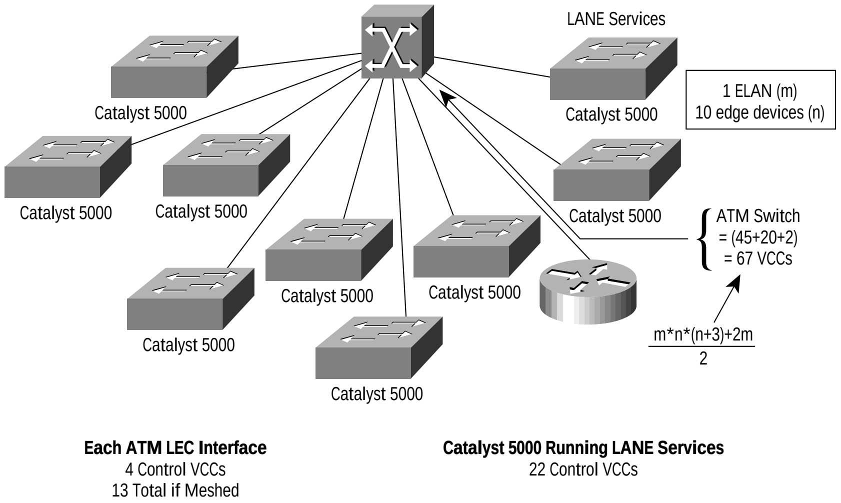

A minimum of 22 SVCs are needed for LANE services, while each edge device (LEC) needs 4 SVCs for being a part of the ELAN. The worst-case requirements are 67 SVCs in the ATM network. If the same example is modified for ten ELANs, the requirements for the number of SVCs increase tenfold, so the number of SVCs needed in the worst case is 670 for the ATM network as a whole.

The formula for computing the SVC requirements in the network is given as follows:

Assumptions:

- Let m be the total number of ELANs in the network.
- Let n be the total number of edge devices in the network. Furthermore, assume that the LANE services are located on a separate edge device and that there are no LECs on the devices implementing the LANE services. Also, assume that all the edge devices need all the ELANs[7].
- The number of SVCs needed at each edge device = $4m + m*(n-1)$
- The number of SVCs needed for LANE services = $m*(2n+2)$
- The total number of unicast SVCs needed in the ATM network = $m* n*(n-1)/2 + m*2n = m*n*(n+3)/2$
- The total number of multicast SVCs needed in the ATM network = $2m$

The figures for the total number of SVCs needed in the ATM network give an idea of how many SVCs are needed for entire networks across multiple ATM switches; they should not be mistaken as requirements for each ATM switch in the network. Moreover, this exercise should be viewed more from the point of view of determining the minimum SVC counts, not the maximum. In other words, this example should serve as a guide to understanding the number of SVCs needed for LANE-control SVCs and the number that are available for data-direct SVCs.

These formulae do not reflect the Catalyst 3000 product line because of the difference in the way the ATM module on these switches assigns SVCs. Refer to (3) for a detailed understanding of SVC budgeting in the Catalyst 3000.

Cisco LANE Configuration

Cisco LANE Networks

The following features characterize Cisco LANE Networks:

- One active LECS supports all ELANs in a single LANE administrative domain.
- 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; they cannot be separated.
- There can be only one LES/BUS pair per subinterface.
- There can be only one active LES/BUS pair per ELAN.
- The LECS and LES/BUS can be different routers/bridges/workstations.
- 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.
- When defining VLANs with the Catalyst 5000, each VLAN should be assigned to a different ELAN.

Cisco's LANE Implementation

The LECS, LEC, LES, and BUS functionality can be implemented at different Cisco devices; a complete list is given in Table 4.

Table 4 LANE Functionality by Product

Product	LANE Components Available	Software Release
Catalyst 5000	LECS, LES, BUS, LEC	ATM module software 2.0 and higher
Catalyst 3000	LECS, LES, BUS, LEC	• Software Version 1.2 for LEC only • Version 2.1 for LECS, LES, and BUS
Cisco 7000	LECS, LES, BUS, LEC	Cisco IOS Release 11.0 and higher
Cisco 75xx, 4xxx	LECS, LES, BUS, LEC	Cisco IOS Release 11.1 and higher
LightStream 1010	LECS, LES, BUS, LEC	LightStream 1010 Software Version 11.2

---

[7] Although this assumption is unrealistic for most networks, it is used to compute the worst-case SVC requirements. The accurate SVC requirements can be determined by adding the counts at each edge device.

These functions are defined on ATM physical interfaces and subinterfaces. A subinterface, which can be defined as a logical interface, is a part of a physical interface such as an OC-3. ATM interfaces on the Cisco routers and the ATM module on the Catalyst 5000 can be logically divided into up to 255 subinterfaces.

Although the same Cisco Internetwork Operating System (Cisco IOS™) code is used on the Catalyst 3000, the subinterface concept does not apply. The LEC can be configured using its menu-driven interface.

Simple LANE Network Configuration

Figure 13 Simple LANE Network with Associated Components

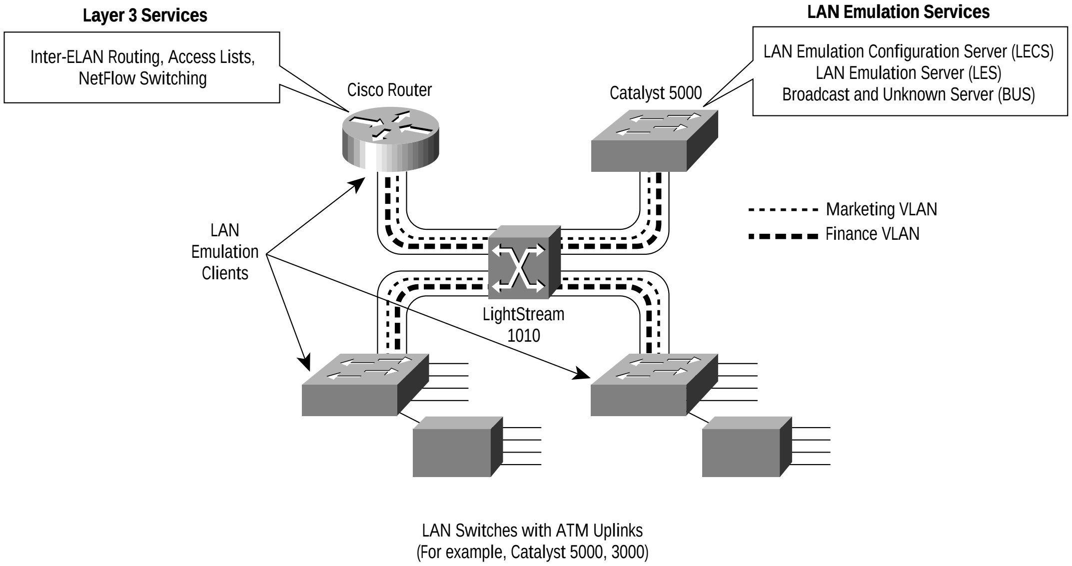

The network requirement for building a LANE network is to support two VLANs (ELANs) and have the ability to route between the VLANs. The first choice must be where to run the LANE services. For this example, the Catalyst 5000 runs the LANE services.

Cisco LANE Automatic Address Assignment

Each of the LANE components, namely the LECS, LES, BUS, and LEC, requires an ATM address to be uniquely identified at the end station. Since the automatic address assignment is used extensively, its understanding is very important [8]. The structure of the ATM address is shown in Figure 14.

Figure 14 ATM Address Structure

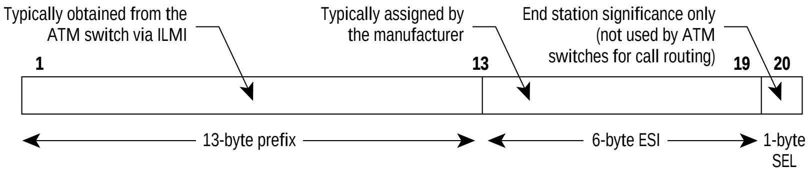

---

[8] LANE can be configured using an ATM addressing structure defined by the network administrator. It is important to completely understand the ramifications of addressing for a campus network before embarking on such a mission. Automatic address assignment is much more convenient and suitable for most networks.

Cisco routers and the Catalyst 5000 come preconfigured with MAC addresses that are used as the end system identifier (ESI) portion of the 20-byte ATM addresses. These MAC addresses are programmed on the chassis backplane and hence are preserved across board swaps, Route Switch Processor (RSP) swaps (for Catalyst 7500s) and supervisor swaps (for Catalyst 5000s). They can be displayed using the command `show lane default` when the ATM uplink is in shutdown state.

```
C50000-ATM> show lane default
interface ATM0:
LANE Client:        …00000C511570.**
LANE Server:        …00000C511571.**
LANE Bus:           …00000C511572.**
LANE Config Server: …00000C511573.00
note: ** is the subinterface number byte in hex
```

As can be seen, each LANE component has a unique ESI portion. Only a single LECS can be configured on a single module, and hence the selector byte for the LECS is 0x00. The selector byte for the remaining components is used to identify different subinterfaces. Since there can be a single ELAN per subinterface, it follows that this selector byte uniquely identifies the ELAN that is configured on the subinterface. This fact can be very useful in troubleshooting since subinterfaces are only an internal Cisco representation and do not mean anything to external ATM analyzers.

Once the ATM module is plugged into an ATM switch, it obtains the 13-byte prefix from the ATM switch using ILMI. In order to do so, the ILMI permanent virtual circuit (PVC) (virtual path identifier/virtual channel identifier [VPI/VCI]—0/16) must be configured on the ATM module. In addition, the signaling PVC (VPI/VCI—0/5) is also needed.

```
interface atm 0
 atm pvc 0 5 qsaal
 atm pvc 0 16 ilmi
```

These lines need to be configured on the router ATM modules while they are preconfigured on the Catalyst 5000.

Once the lines are in the configuration, the show lane default command shows the 19-byte ATM address (provided the module is plugged into the ATM switch and the line protocol is up). The selector byte is used when the LES/BUS and the LEC components for different ELANs are configured.

```
C50000-ATM> show lane default
interface ATM0:
LANE Client:    47.0091.8100.0000.0800.200c.1001.00000C511570.**
LANE Server:    47.0091.8100.0000.0800.200c.1001.00000C511571.**
LANE Bus:       47.0091.8100.0000.0800.200c.1001.00000C511572.**
LANE Config     47.0091.8100.0000.0800.200c.1001.00000C511573.00
Server:
note: ** is the subinterface number byte in hex
```

The main configuration steps include:

- Identify the box that will run the LECS and get its automatic address using show lane default.
- Identify the box and the subinterfaces that will run the LES/BUSes for the different ELANs. Based on the subinterface number, obtain the ATM address of each LES using show lane default on that box. The hex value of the subinterface number becomes the selector byte. Subinterface 1 represents selector byte 01, subinterface 2 represents 02, and subinterface 10 represents 0A.
- Configure the ATM switch(es) with LECS ATM address.
- Configure the LECS database with the ELAN names to their LES ATM address mappings. Start up the LECS.
- Start up the LES/BUS pairs at the different subinterfaces.
- Start up the LECs at all the edge devices and make ELAN/VLAN associations where needed. A decision on whether to use Virtual Trunking Protocol (VTP) needs to be made at this point as well.

In this example, all the LANE services are configured on a single Catalyst 5000.

LightStream 1010 Configuration

Each LEC queries the ATM switch for the LECS ATM address via ILMI. To enable such queries to get a successful response from the ATM switch, the ATM address of the LECS should be configured into the ATM switch (LightStream 1010). An example of the command to set the ATM address of the LECS on the LightStream 1010 follows:

```
LS1010(config)# atm lecs-address 47.0091.8100.0000.0800.200c.1001.00000C511573.00
 <NSAP of first LECS address>
```

Note that the ATM address used is the one obtained from the show lane default command on the Catalyst 5000 that runs the LECS. At this point, the switch is able to supply the clients with the LECS address, which will point at the Catalyst 5000.

Catalyst 5000 (LANE Services) Configuration

- LECS configuration: The LECS database identifies two LANE servers, finance and marketing, and their ATM addresses.

  The ATM addresses for these LESs are again obtained using the show lane default command on the Catalyst 5000. The 19-byte address for the LES is appended with the subinterface number (in hex), and this 20-byte address is used to define the ELAN name to its LES ATM address mapping. In this example, subinterfaces 1 and 2 are used for ELANs finance and marketing, respectively. These LES addresses are the addresses that the LECS give to LE clients trying to join a given ELAN. The default ELAN is set up for clients that did not match any of the assignment statements. In a very secure environment, specific statements map each client specifically to an ELAN; this decision is made by the network administrator. Refer to the LANE configuration guide for more details.

```
lane database example
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
```

After the LECS database is defined, the primary interface on the Catalyst 5000 must be configured to run the LECS. This step is accomplished with the lane config example command, which uses the “example” database for the LECS join decisions. The lane config auto-config-atm-address command starts the LECS process on the Catalyst 5000.

```
interface atm 0
atm pvc 1 0 5 qsaal
atm pvc 1 0 16 ilmi
lane config example
lane config auto-config-atm-address
```

Additional details on configuration of the LECS database include:

- Configure ELAN names at the LEC: In this configuration, all the LECs need to be configured with a ELAN name, which can be embedded in their Configure_Requests. The basic form of the LECS database, this form needs to contain only 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 to the corresponding ATM address of the LES. This scenario is shown in Figure 20. A LEC that does not know which ELAN to join may be assigned to a default ELAN if such an ELAN is configured in the LECS database.
- Configure LEC to ELAN assignment in the LECS database: In this configuration, all the information is centralized in the LECS database. The LECs themselves can just go to the LECS to find out which ELAN they are supposed to join. Although this configuration is more tedious, it provides tighter control over all the ELANs and is useful where security is a big issue. The LECs themselves are identified by their ATM or MAC addresses. In addition, wildcarding of ATM address prefixes is also allowed, making creation of relationships of the form Assign any LEC joining with a prefix of A to ELAN X, and so on, very useful. The configurations are shown in Table 4.
- Hybrid combination: Configuration of a combination of these methods is also possible.

Table 5 TLECS Database Configuration

Config 1
LEC-to-ELAN
Mapping at the LEC

```
lane database test-1
name finance lex-atm-address 47.00091.8100.0000.0800.200c.1001.01
name marketing les-atm-address 47.00091.8100.0000.0800.200c.1001.02
default-name finance
```

Config 2
LEC-to-ELAN
Mapping in the LECS Database

```
lane database test-2
name finance les-atm-address 47.00091.8100.0000.0800.200c.1001.01
name marketing les-atm-address 47.00091.8100.0000.0800.200c.1001.02

client-atm-address 47.0091.8100.000.08... name finance
client-atm-address 47.00091.8100.0000.08... name marketing

mac-address 00c0.0000.0100.0100 name finance
mac-address 00c0.1111.2222 name marketing
```

- LES/BUS configuration: Subinterface 0.1 needs to be configured to provide LES/BUS services for ELAN finance and subinterface 0.2 for ELAN marketing. The command used for starting the LES/BUS for a given ELAN is lane server-bus ethernet <ELAN name>. When this command is issued, the subinterface serves the ELAN, but has no clients actually taking part in the ELAN.

```
interface atm 0.1
lane server-bus ethernet finance
interface atm 0.2
lane server-bus ethernet marketing
```

At this point, the LANE service component is ready for functioning and the only remaining step is to enable the LECs on all the devices. The LEC configuration on each device is treated separately.

LANE Client Configuration

- LEC configuration on the router: The router subinterfaces that are acting as LECs handle the routing between VLANs. A LEC should be configured at each subinterface that needs routing. The `lane client ethernet <elan>` name starts this client software. If you don’t specify an ELAN name, the LECS either finds out the name from its database, or it assigns the default

ELAN. You can specify a specific ELAN to join, but that information must match what is programmed in the LECS in order for the client to successfully join the ELAN.

```
interface atm 1/0.1
ip address <ip address> <netmask>
ipx network <ipx network number>
lane client ethernet finance
interface atm 1/0.2
ip address <ip address> <netmask>
ipx network <ipx network number>
lane client ethernet marketing
```

This configuration is often referred to as router on a stick.

- LEC configuration on the Catalyst 5000: At this point, the Catalyst 5000 ISL VLAN domain meets the ATM ELAN domain, and some way to map VLANs to ELANs must exist. The solution follows:

```
interface atm 0.1
lane client ethernet <VLAN #> finance
interface atm 0.2
lane client ethernet <VLAN #> marketing
```

The `lane client ethernet` command has an additional field that allows the administrator to assign a specific VLAN to an ATM ELAN. This assignment should be made for each VLAN in the Catalyst 5000 that needs to have connectivity to the ATM LANE domain. In this way bridging is accomplished throughout the VLAN and ELANs.

- LEC configuration on the Catalyst 5000 with VTP: The VTP protocol was created to solve the administrative burden of having to create mappings from ISL VLANs to 802.10 VLANs to LANE-based ELANs. With VTP, a single identifier for a VLAN can be created, and the switches automatically map the different VLAN numbers to this VLAN name. ISL and 802.10 have numbering schemes to identify VLANs, and hence VTP will create a mapping from the VLAN numbers to the VLAN name. LANE, on the other hand, already uses ELAN names to identify its VLANs, and hence the VLAN name is retained to be the ELAN name as well. An example of such a mapping is given in Table 5.

Table 6 VTP Example—ELAN, ISL, and 802.10 Mappings

VTP VLAN Name	LANE ELAN Name	ISL VLAN Identifier	802.10 VLAN Identifier (SAID)
Finance	Finance	1	100001
Marketing	Marketing	2	100002

The mappings in Table 5are sent as VTP advertisements by VTP servers (for example, Catalyst 5000s) along trunks (such as ISL or 802.10) to all switches in a switch fabric. All the switches eventually converge and develop a consistent view of the different VLAN mappings across the entire fabric of the switches. Refer to [Merwyn's VTP guide] for a detailed explanation of VTP.

From the perspective of the Catalyst 5000 ATM module, VTP can be useful in reducing the configuration burden since it is possible to automatically create a LEC whenever a new VLAN/ELAN is created. The requirement to use VTP successfully is to have a LES/BUS running for all ELANs. If the LES/BUS for a given ELAN is not configured, then the LECs for that ELAN will never come up. VTP is turned off by default on the ATM module of the Catalyst 5000 and needs to be enabled explicitly by using the global command "vtp enable."

WARNING: Enabling VTP in a LANE network should be examined very carefully since in larger networks it may cause more harm than benefits. Some of the concerns include:

  - SVC budgets—The scalability issue with respect to SVC budgets is discussed in a later section. VTP enables all ELANs on all switches, whether they are needed or not, and hence the SVC budgets for the ATM network may be exceeded.
  - VTP pruning—VTP pruning helps in pruning the broadcast/multicast traffic on trunks to only those VLANs that need to be there. However, this scenario does not work across LANE because of the way LANE was designed.
  - For more details on VTP implementation issues, refer to (4).
- LEC configuration on the Catalyst 3000: As mentioned previously, the Catalyst 3000 follows a menu-based configuration scheme that is slightly different from that of the Catalyst 5000. Enabling a LEC is a two-step process instead of the single command that is used on the Catalyst 5000s and the routers. On the Catalyst 3000s, the VLAN name needs to be changed to whatever the ELAN name is and then the LEC needs to be enabled.

  It should be noted that, in designing networks with the Catalyst 3000 ATM module, the design limitations must be well understood. These limitations are documented in detail in (3).

Cisco LANE Configuration Using User Assigned ATM Addressing

User assigned ATM addressing requires some planning and careful ATM address administration but has some benefits. Once the addresses are assigned to the different LANE components, they are preserved across board swaps, supervisor swaps, RSP swaps and chassis swaps.

While this section covers the configuration of LANE using ‘custom’ ATM addressing, the actual administration of ATM addresses is beyond the scope of this document. Please refer to the ATM Forum signaling specifications to get a clearer understanding of obtaining ATM address prefixes and the authorities managing ATM address spaces.

An example configuration is shown below.

Table 7 LANE Configuration Using User Defined ATM Addressing

User Defined LANE Component Addresses

```
LECS : 47.009188880000000000000000.000000000001.00
LES : 47.009188880000000000000000.000000000002.**
BUS : 47.009188880000000000000000.000000000003.**
LEC : 47.009188880000000000000000.000000000004.**
```

Sample LANE services and client configuration

```
lane database example
name elan1 server-atm-address 47.009188880000000000000000.000000000002.01
name elan2 server-atm-address 47.009188880000000000000000.000000000002.02
!
interface ATM0
atm preferred phy A
atm pvc 1 0 5 qsaal
atm pvc 2 0 16 ilmi
lane config config-atm-address 47.009188880000000000000000.000000000001.00
lane config database example
!
interface ATM0.1 multipoint
lane client-atm-address 47.009188880000000000000000.000000000004.01
lane server-atm-address 47.009188880000000000000000.000000000002.01
lane bus-atm-address 47.0091888800 lane server-bus ethernet elan1
lane client ethernet 1 elan1
!
interface ATM0.2 multipoint
lane client-atm-address 47.009188880000000000000000.000000000004.02
lane server-atm-address 47.009188880000000000000000.000000000002.02
lane bus-atm-address 47.009188880000000000000000.000000000003.02
lane server-bus ethernet elan2
lane client ethernet 2 elan2
```

LANE Redundancy

Introduction

The ATM Forum's work on LANE has given the initial impetus for customers to move workgroup and campus networks to ATM. With LANE, customers can run existing LAN-based applications and broadcast-oriented LAN protocols over ATM. LANE Version 1.0 defines the standards for internetworking legacy LANs such as Ethernet and Token Ring with ATM-attached devices. Such devices include end stations (for example, ATM-attached servers), edge devices that bridge legacy LANs onto an ATM backbone (for example, Catalyst 5000s) and ATM-attached routers to route between ELANs.

However, since LANE Version 1.0 does not define mechanisms for building redundancy and fault tolerance into the LANE services, they become single points of failure. Moreover, the issues of router redundancy and path/link redundancy also need to be resolved. An understanding of these issues and subsequent customer demand to solve them has led Cisco to incorporate multiple mechanisms that can be used to build fault-tolerant ATM networks. This document provides an understanding of these various mechanisms while also highlighting design rules and issues to consider while implementing redundant LANE networks.

This section begins with a discussion on Simple Server Redundancy Protocol (SSRP), which was developed to provide redundant LANE services. Although many vendors have implemented redundant LANE services of some fashion, the services all violate the LANE 1.0 specification and hence are not interoperable with other third-party implementations. However, SSRP does not violate the LANE 1.0 specification and hence is interoperable with third-party LEC implementations, a very important consideration in implementing an interoperable ATM network. The discussion on SSRP is followed by a description of Hot Standby Router Protocol (HSRP) over LANE, which provides a mechanism for building router redundancy. Finally, the Spanning-Tree Protocol and other product-specific features that can be used to build link and path redundancy into edge devices are discussed.

Issues such as network redundancy are covered by protocols such as PNNI and are not discussed in this paper.

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 setup manifests itself 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 an impact on the operation of the network. Specifically:

- LECS failure: A failed LECS impacts all the ELANs since it provides access control for all the ELANs under its control. Although the existing ELANs continue to work normally (assuming only Cisco LECs), no new LEC can join any ELAN under control of that LECS. Also, any LEC that needs to rejoin its ELAN or change its membership to another ELAN is not able to because the LES is unable to 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 limit networks where resiliency and robustness is a requirement and might even be a deciding factor in choosing to implement LANE-based ATM networks.

In addition, other design considerations such as placement of the LANE service components within an ATM network can have implications on the overall robustness of the LANE environment. This subject was discussed in an earlier section of this design guide.

It should be noted that the discussion on resiliency in LANE in this document deals with failures in LANE service mechanisms (LECS/LES/BUS) rather than with ATM network issues such as meltdowns and other network side issues. ATM network redundancy is left for protocols such as PNNI and eventually the network designers themselves to incorporate the needed redundancy in their respective networks.

Resiliency in LANE 1.0 Networks

Increasing the resiliency of a LANE-based network essentially includes delivering increased robustness in the service components of LANE such as the LECS, LES, and BUS. Such robustness is provided by SSRP via 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 handled similarly, where one primary LES/BUS pair is backed up by multiple secondaries. Note that the LES/BUS functions are always colocated 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 an LEC during initialization is to connect with the LECS to obtain the LES ATM address for the ELAN it wishes to join. In order for the LEC to connect to the LECS, multiple mechanisms are defined. The first mechanism that an LEC should use is to query the ATM switch to which it is attached for the LECS address. This address discovery process is accomplished using the ILMI protocol on VPI, VCI—0, 16.

Figure 15 Configuration of Multiple LECSs and LEC Initialization Using LECS Address Discovery

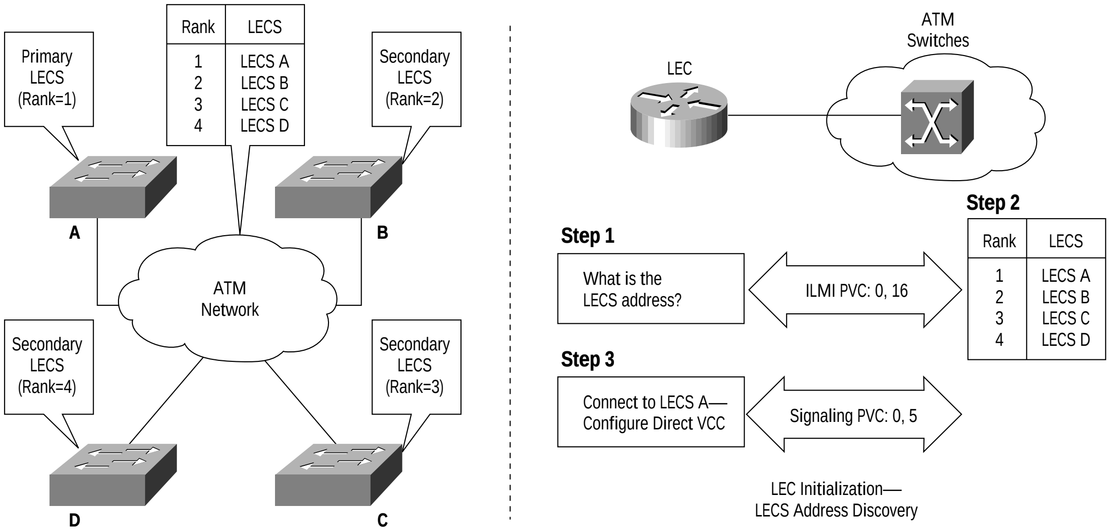

The LECS address is configured into the ATM switches. On the LightStream 1010, the configuration command to add an LECS address follows:

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

With SSRP, multiple LECS addresses are configured into the ATM switches, as shown in Figure 15. A LEC that requests the LECS address from the ATM switch gets the entire table of LECS addresses in response. The LEC should try 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.

While 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. Upon initialization, an LECS obtains the LECS address table from the switch, similar to what the LEC does upon initialization during Steps 1 and 3, as shown in Figure 15. The LECS then tries to connect to all the LECSs that are below itself in rank, as shown in Figure 13(a). The rank is derived from the index entry in the LECS address table.

If an LECS has a connection (VCC) from an LECS whose rank is higher than its own, then it will be in backup mode. The highest-ranking LECS has no other LECS that connect to it from above and hence assumes the role of the primary. The procedure describing how a backup takes over in the case of a failed primary is explained by using the following example.

Figure 13 shows a LANE network with four LECS entities configured. All the ATM switches in the network are configured with the same LECS address table. Upon startup, 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, it tries and connects to LECS B, C, and D. B in turn connects to C and D, and C connects to D. Thus a downward establishment of VCCs occurs, and, since A does not have any VCCs from above, it becomes the primary.

During normal network operation, A responds to all the configure requests while the backups B, C, and D do not respond to any queries. If for some reason the primary (A) fails (because of link failure, box failure, and so on), then B loses its VCC from A, and so do the others. At this point, B does not have any VCCs from above and, therefore, is now the highest-ranking available LECS in the network. B now becomes the primary. C and D still have connections from LECSs that have a higher rank than their own, and hence they continue to be in backup mode, as shown in Figure 16.

Figure 16 (a) Normal Operation of Four Redundant LECSs (b) Operation under a Failure in the Primary LECS

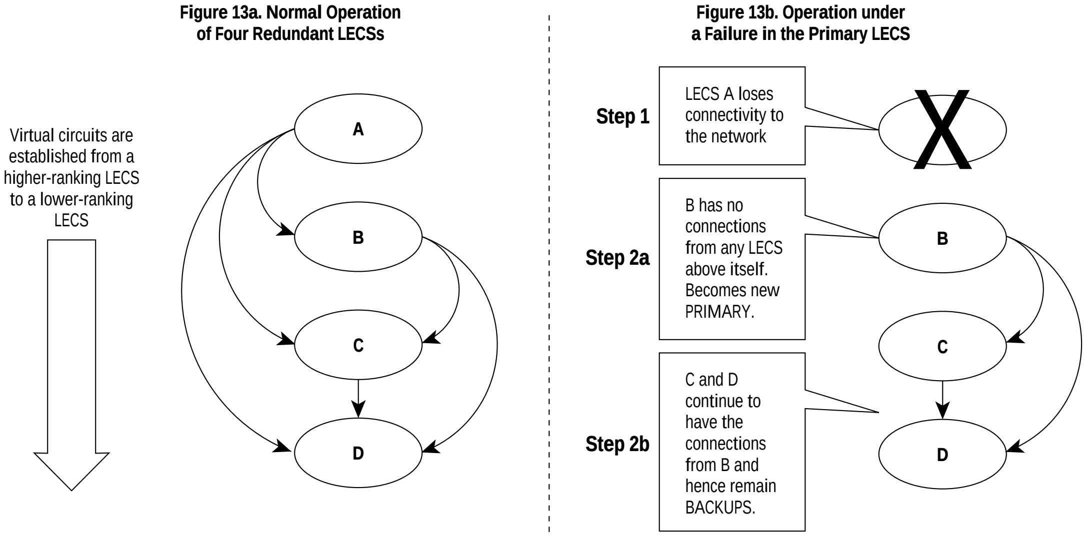

Figure 13a. Normal Operation of Four Redundant LECSs

Virtual circuits are established from a higher-ranking LECS to a lower-ranking LECS

Figure 13b. Operation under a Failure in the Primary LECS

LECS A loses connectivity to the network

B has no connections from any LECS above itself. Becomes new PRIMARY.

C and D continue to have the connections from B and hence remain BACKUPS.

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.

Multiple LES/BUS pairs for a given ELAN are configured first into the LECS database, as shown in Figure 12. Within this database, each LES/BUS pair is assigned a priority. Upon initialization, each LES/BUS opens a VCC, with the primary LECS using the LECS address discovery mechanism described earlier. 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.

Figure 17 (a) Physical View of a Redundant LES/BUS Configuration (b) Logical View with the LECS Database

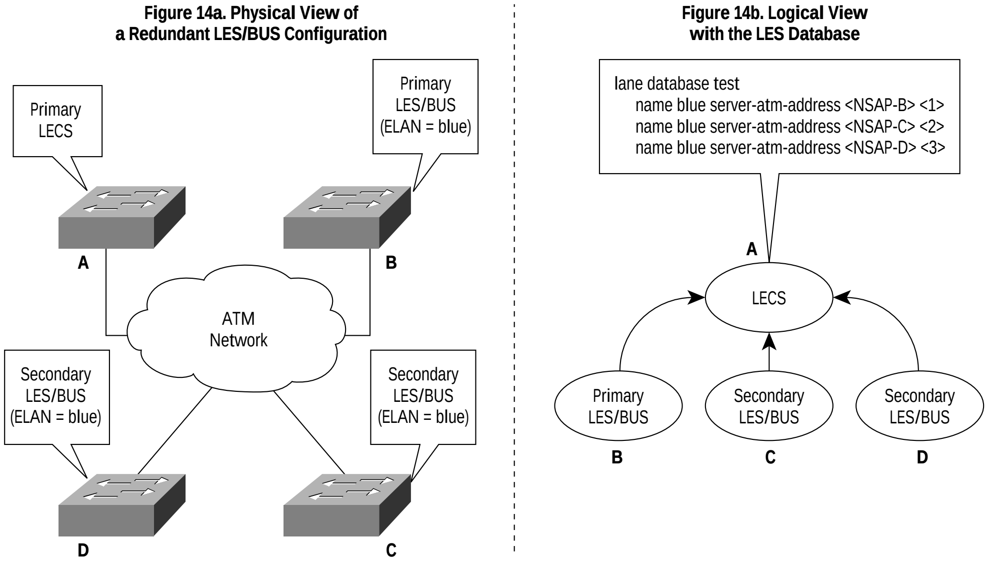

Under normal operating conditions, in the example given all the LES/BUS pairs B, C, and D have VCCs open to the primary LECS. Since B has the highest priority according to the LECS database, all join requests from the LEC are guided to B. If B loses connectivity to the network, its VCC to the LECS is torn down. All the LECs that were part of the ELAN reinitialize since the active LES/BUS for that ELAN has gone down (the control and multicast VCCs for each LEC are torn down). They resend their join requests to the LECS. The LECS guides these join requests to the next LES/BUS pair in order of priority; in this case, it is C. Therefore, the LECs now join the same ELAN, with the backup LES/BUS pair becoming the active one for that ELAN. This process is illustrated in Figure 15.

As can be clearly seen in this process, the only resources consumed while implementing SSRP are a few VCs. No other overhead is incurred in terms of added traffic on the network.

Consult the Cisco IOS configuration guide for configuring SSRP into your network.

SSRP Usage Guidelines

Design Tips

There are no theoretical limits on the number of LECSs that can be configured using SSRP, but 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 since redundancy adds complexity and is very cumbersome when managing and troubleshooting the network.

Configuration Guidelines

In order to support the LECS redundancy scheme, the following configuration rules must be strictly followed. Failure to comply will result in improper operation of SSRP and an improperly functioning network.

1. Each LECS must maintain exactly the same database of ELANs. Therefore, it is imperative for network managers to maintain the same ELAN database across all the LECSs.
2. The LECS addresses in the LECS address table must be configured in the same order on each of the ATM switches in the network.
3. When using SSRP with the well-known address, DO NOT place two LECSs on the same ATM switch. This setup does not work since only one LECS can register the well-known address with the ATM switch (via ILMI). This causes problems during initialization.

Figure 18 (a) LEC VCCs with Primary LES/BUS (b) New LEC VCCs with Backup LES/BUS

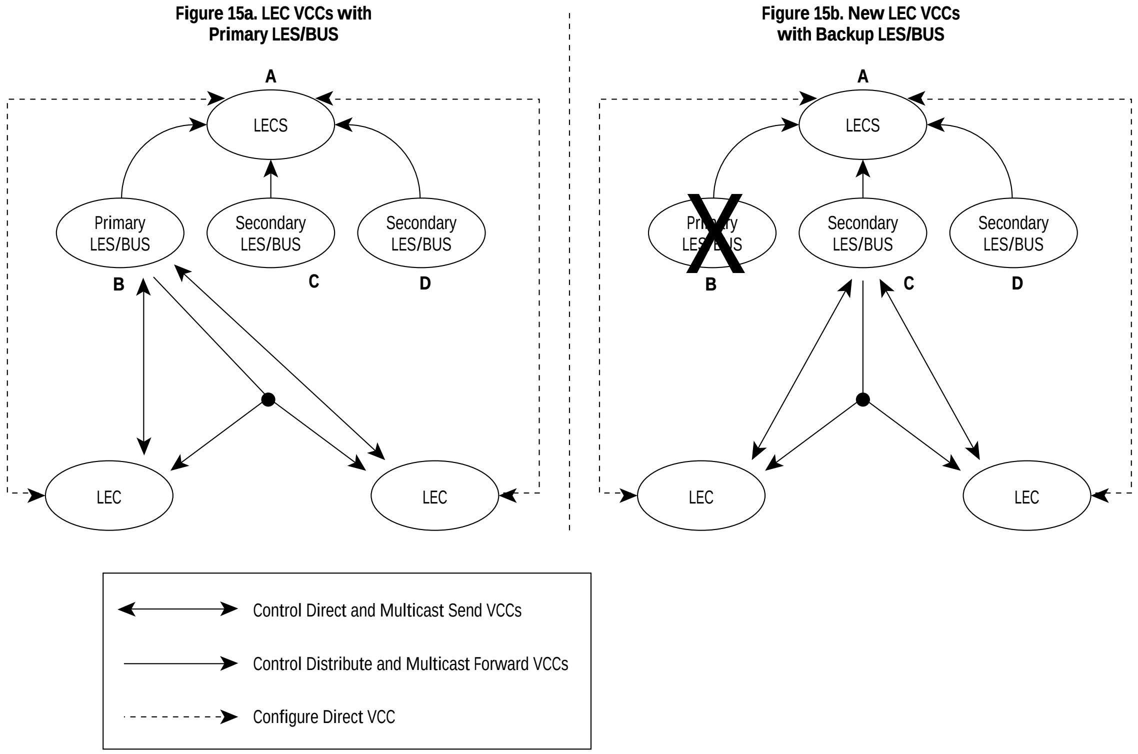

Interoperability Notes

SSRP can be used with independent third-party LECs as long as they use the LECS address discovery mechanism using ILMI and handle multiple LECS addresses returned by the ATM switch appropriately. In other words, 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. Most adapter vendors, including Olicom, Interphase, Adaptec, and Efficient, are implementing SSRP-compatible versions of their LECs that comply with this directive.

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 in addition to the autoconfigured 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 as well as route any call setup requests to the correct place. However, under SSRP, multiple LECSs are in the network, and if each one registers the well-known address to the ATM switches that they are connected to, then call setups are routed to different places in the network.

Hence, under SSRP it is necessary to 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. Thus if the master fails the well-known address moves with the master LECS. The PNNI code on the LightStream 1010 advertises the new route to the well-known address upon change of mastership of the LECS. 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.

Therefore, in order to implement SSRP with the well-known address, use the following steps:

1. First configure the LECS to listen on the autoconfigured address (or 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.
2. Configure each LECS to listen on the well-known address using the command lane config fixed-config-atm-address. Once the master LECS is determined using the LECS redundancy procedure, the master registers the well-known address to the ATM switch to which it is attached.

Caveat: When using SSRP with the well-known address, it is advisable to have each LECS on a separate ATM switch because of the possibility of duplicate address registration during failover on the same switch, which ILMI does not allow.

Behavior of SSRP in Network Partitions

In the event of network partitions where two separate ATM clouds are formed because of interconnecting link or switch failure, each cloud has its own set of LANE services, provided SSRP is configured is such a manner. This setup is illustrated in Figure 16.

Figure 19 Behavior of SSRP during a Network Partition

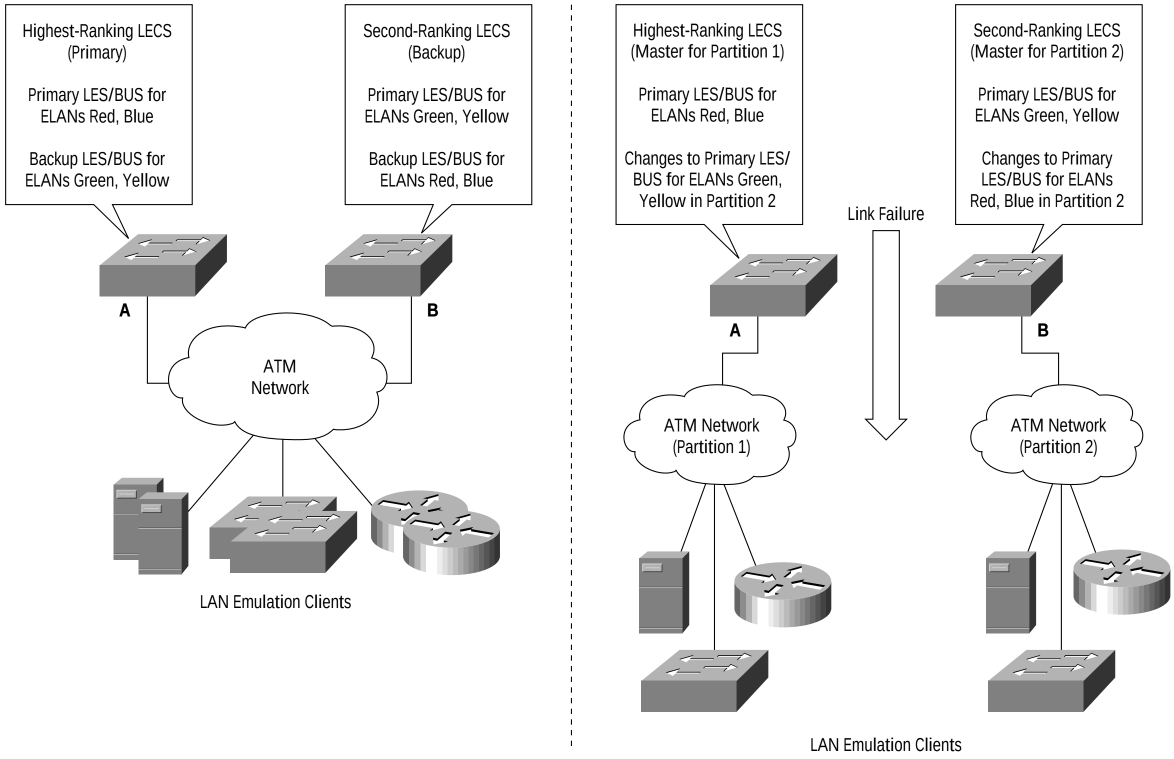

The following concepts should be noted while configuring SSRP with the possibility of network partition:

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

HSRP over LANE

HSRP is a protocol to guard against router failures in the network. The HSRP protocol is exchanged between two routers at its simplest; one of them is elected as the primary router interface (or subinterface) for a given subnet, while the other 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 in order 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.

Since 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.

In Figure 17, an IP end station in a given subnet (which maps to an ELAN) is configured with the default router address, which is the same as the 'shared' HSRP address. Each of the routers has its own IP address that it uses to exchange 'hello' messages with other routers. Each router is a LEC in the given ELAN. A data-direct SVC is opened between router A and router B to exchange the hellos.

If the primary router (router A) goes down, then the redundant router (router B) waits for three hello times before taking over as the new owner of the HSRP IP address and the HSRP MAC address. When the MAC address moves to router B, all LECs must get the new MAC/ATM address mapping and refresh their LE-ARP caches. This process is achieved by having router B send out and LE-NARP (a negative ARP) indicating the change.

Figure 20 HSRP over LANE (Sub)interfaces

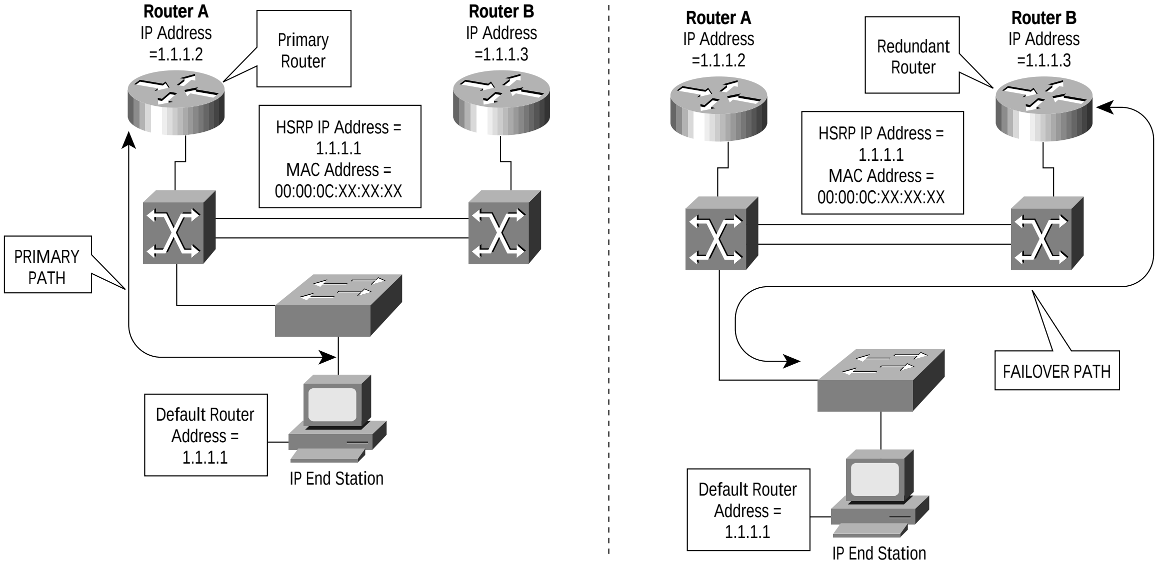

Sample HSRP and SSRP Configuration

In this section, we present a sample configuration of a LANE network which includes SSRP and HSRP.

Catalyst 5000s 1 and 2 are running as primary and secondary LANE services. The protocol they use to co-ordinate the primary and secondary is known as SSRP. The routers 7507-1 and 7507-2 are using HSRP. 7507-1 is the primary for ELAN's 1, 2 and backup for ELAN's 3, 4. 7507-2 is the primary for ELAN's 3,4 and backup for ELAN's 1,2.

Figure 21   Example LANE Network Running SSRP and HSRP

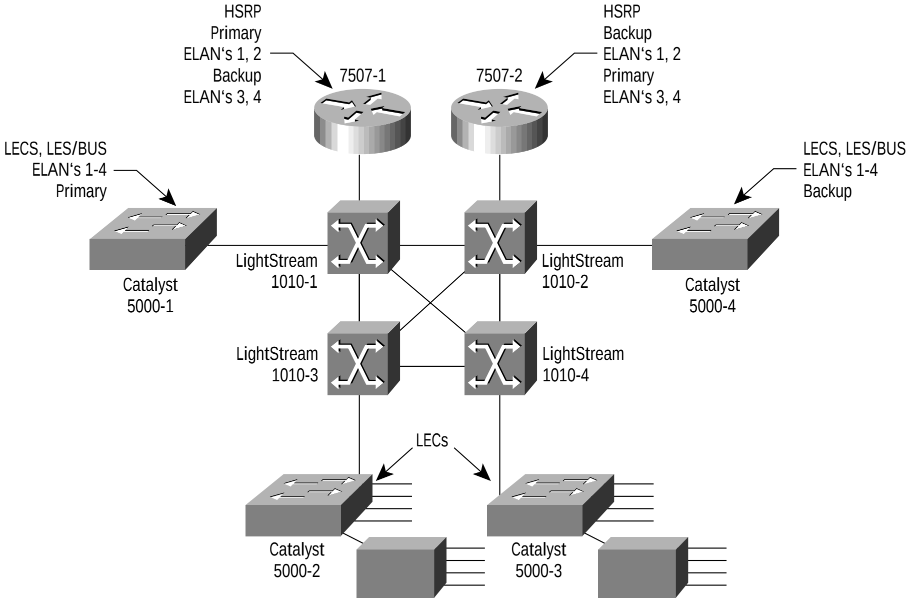

Table 8   SSRP Configuration

Catalyst 5000—1 (Primary)

```
 lane database lecsdb
 name elan1 server-atm-address 47.00918100000000613E5C0901.00400BE69C31.01
 name elan1 server-atm-address 47.009181000000006170599301.00400BE6A041.01
 name elan2 server-atm-address 47.00918100000000613E5C0901.00400BE69C31.02
 name elan2 server-atm-address 47.009181000000006170599301.00400BE6A041.02
 name elan3 server-atm-address 47.00918100000000613E5C0901.00400BE69C31.03
 name elan3 server-atm-address 47.009181000000006170599301.00400BE6A041.03
 name elan4 server-atm-address 47.00918100000000613E5C0901.00400BE69C31.04
 name elan4 server-atm-address 47.009181000000006170599301.00400BE6A041.04
interface ATM0
 mtu 1500
 no ip address
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 lane config auto-config-atm-address
 lane config database lecsdb
!
interface ATM0.1 multipoint
 lane server-bus ethernet elan1
!
interface ATM0.2 multipoint
 lane server-bus ethernet elan2
!
interface ATM0.3 multipoint
 lane server-bus ethernet elan3
!
interface ATM0.4 multipoint
 lane server-bus ethernet elan4
```

Table 8 SSRP Configuration (Continued)

Catalyst 5000—2 (Secondary)

```
 name elan1 server-atm-address 47.00918100000000613E5C0901.00400BE69C31.01
 name elan1 server-atm-address 47.009181000000006170599301.00400BE6A041.01
 name elan2 server-atm-address 47.00918100000000613E5C0901.00400BE69C31.02
 name elan2 server-atm-address 47.009181000000006170599301.00400BE6A041.02
 name elan3 server-atm-address 47.00918100000000613E5C0901.00400BE69C31.03
 name elan3 server-atm-address 47.009181000000006170599301.00400BE6A041.03
 name elan4 server-atm-address 47.00918100000000613E5C0901.00400BE69C31.04
 name elan4 server-atm-address 47.009181000000006170599301.00400BE6A041.04
interface ATM0
 mtu 1500
 no ip address
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
 lane config auto-config-atm-address
 lane config database lecsdb
!
interface ATM0.1 multipoint
 lane server-bus ethernet elan1
!
interface ATM0.2 multipoint
 lane server-bus ethernet elan2
!
interface ATM0.3 multipoint
 lane server-bus ethernet elan3
!
interface ATM0.4 multipoint
 lane server-bus ethernet elan4
```

Table 9 HSRP Configuration

Cisco 7507—1

```
(Primary - ELAN 1,2 Backup - ELAN 3,4)
interface ATM5/0
 mtu 1500
 no ip address
 ip route-cache optimum
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
!
interface ATM5/0.1 multipoint
 ip address 137.80.110.3 255.255.255.0
 no ip redirects
 lane client ethernet elan1
 ipx network 11
 standby 1 priority 101
```

Cisco 7507—2

```
(Primary - ELAN 3,4 Backup - ELAN 1,2)
interface ATM1/0
 mtu 1500
 no ip address
 ip route-cache optimum
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
!
interface ATM1/0.1 multipoint
 ip address 137.80.110.2 255.255.255.0
 no ip redirects
 lane client ethernet elan1
 ipx network 11
 standby 1 preempt
```

Table 9 HSRP Configuration (Continued)

Cisco 7507—1

```
 standby 1 preempt
 standby 1 ip 137.80.110.1
 bridge-group 1
!
interface ATM5/0.2 multipoint
 ip address 137.80.111.3 255.255.255.0
 no ip redirects
 lane client ethernet elan2
 ipx network 12
 standby 2 priority 101
 standby 2 preempt
 standby 2 ip 137.80.111.1
 bridge-group 1
!
interface ATM5/0.3 multipoint
 ip address 137.80.112.3 255.255.255.0
 no ip redirects
 lane client ethernet elan3
 ipx network 13
 standby 3 preempt
 standby 3 ip 137.80.112.1
 bridge-group 1
!
interface ATM5/0.4 multipoint
 ip address 137.80.113.3 255.255.255.0
 no ip redirects
 lane client ethernet elan4
 ipx network 14
 standby 4 preempt
 standby 4 ip 137.80.113.1
 bridge-group 1
```

Cisco 7507—2

```
 standby 1 ip 137.80.110.1
 bridge-group 1
!
interface ATM1/0.2 multipoint
 ip address 137.80.111.2 255.255.255.0
 no ip redirects
 lane client ethernet elan2
 ipx network 12
 standby 2 preempt
 standby 2 ip 137.80.111.1
 bridge-group 1
!
interface ATM1/0.3 multipoint
 ip address 137.80.112.2 255.255.255.0
 no ip redirects
 lane client ethernet elan3
 ipx network 13
 standby 3 priority 101
 standby 3 preempt
 standby 3 ip 137.80.112.1
 bridge-group 1
!
interface ATM1/0.4 multipoint
 ip address 137.80.113.2 255.255.255.0
 no ip redirects
 lane client ethernet elan4
 ipx network 14
 standby 4 priority 101
 standby 4 preempt
 standby 4 ip 137.80.113.1
 bridge-group 1
```

Table 10 LS1010 Configuration and Sample LEC Configuration

LS1010 - 1 LS1010 - 2 LS1010 - 3 LS1010 - 4

```
atm lecs-address-default 47.0091.8100.0000.0061.3e5c.0901.0040.0be6.9c33.00
atm lecs-address-default 47.0091.8100.0000.0061.7059.9301.0040.0be6.a043.00
```

Table 10 LS1010 Configuration and Sample LEC Configuration (Continued)

Sample LEC

```
interface ATM0
 mtu 1500
 no ip address
 supervisor-engine-ip <ip addr>
 atm pvc 1 0 5 qsaal
 atm pvc 2 0 16 ilmi
!
interface ATM0.1 multipoint
 no ip address
 lane client ethernet 1 elan1
!
interface ATM0.2 multipoint
 no ip address
 lane client ethernet 2 elan2
!
interface ATM0.3 multipoint
 no ip address
 lane client ethernet 3 elan3
!
interface ATM0.4 multipoint
 no ip address
 lane client ethernet 4 elan4
```

Dual PHY ATM Card for the Catalyst 5000

Another aspect of addressing the redundancy needs from a physical network perspective is the addition of a dual redundant PHY port on the LANE module of the Catalyst 5000. Dual PHY redundancy is only at a physical level and is useful in cases where the primary link to the ATM switch goes down.

When the card is used as a pure LANE client(s) configuration, then there are no configuration changes for the dual PHY LANE card from the single PHY. However, when the dual PHY LANE module is used for providing LANE services (LECS, LES/BUS) then there are some important configuration guidelines to take into consideration.

What Are the Issues?

As discussed in section 7.3.1, the 13-byte ATM address prefix for the LANE components on the Catalyst 5000 are obtained from the ATM switch using ILMI. In configurations where only a single uplink is utilized, there are no configuration issues since there is only one ATM address prefix. In the case of the dual PHY LANE card, the Catalyst 5000 is dual homed to different LS1010s as shown in Figure 22. Each LS1010 has a different ATM address prefix and therefore the Catalyst 5000 will obtain a different 13-byte prefix depending on which link (PHY-A or PHY-B) is active. An example is illustrated in Table 11.

Figure 22 Dual PHY ATM

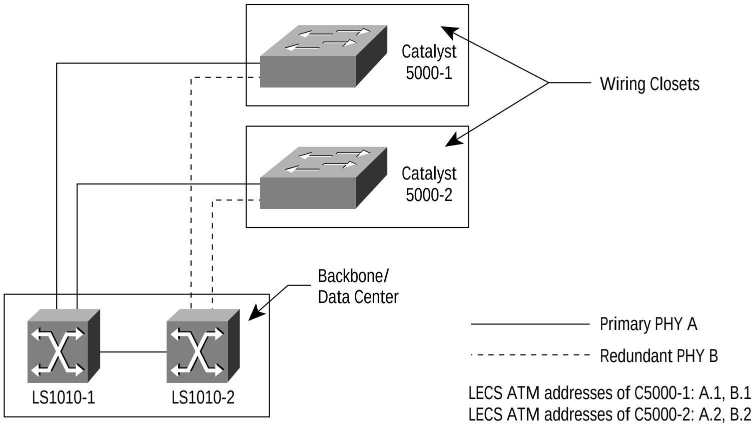

Table 11 Example of Cisco default ATM addressing with the Dual-PHY module

Output of “show lane default” with PHY-A active

```
interface ATM0:
LANE Client: 47.00918100000000605C71F401.00603E9EFC30.**
LANE Server: 47.00918100000000605C71F401.00603E9EFC31.**
LANE Bus: 47.00918100000000605C71F401.00603E9EFC32.**
LANE Config Server: 47.00918100000000605C71F401.00603E9EFC33.00
note: ** is the subinterface number byte in hex
```

Output of “show lane default” with PHY-B active

```
interface ATM0:
LANE Client: 47.00918100000000603E5C4701.00603E9EFC30.**
LANE Server: 47.00918100000000603E5C4701.00603E9EFC31.**
LANE Bus: 47.00918100000000603E5C4701.00603E9EFC32.**
LANE Config Server: 47.00918100000000603E5C4701.00603E9EFC33.00
note: ** is the subinterface number byte in hex
```

These multiple addresses need to be configured properly into the LANE database and the LS1010 LECS address table configurations. In the figure below, the LECS ATM addresses of C5000-1 are A.1 and B.1 where ‘A’ and ‘B’ are 13-byte ATM address prefixes for LS1010-1 and LS1010-2 respectively. For C5000-2 the LECS ATM addresses are A.2 and B.2.

Configuring LANE Clients on the Dual-PHY LANE Card

The ATM address of the LANE client is typically never configured into the LANE database and therefore the LECS will not care about the ATM address of the joining LECs. However, when security is a concern and there is a need to configure the ATM addresses of the LECs allowed to join a certain ELAN then it is important to configure both the ATM addresses associated with the LEC on the Catalyst 5000 into the LANE database.

Configuring LANE Services on the Dual-PHY LANE Card

Using Cisco default addressing : The ATM addresses of the LECS and the LES are used to configure the LECS table on the LS1010s and the LANE database respectively. Since there are multiple addresses for the dual-PHY LANE card, the order in which they appear in the LECS table and the LANE database becomes very important.

In the example shown in Figure 22, the LECS ATM addresses for Cat5000-1 are A.1 and B.1 and those for Cat5000-2 are A.2 and B2. These addresses should go into the LS1010s LECS address table and the different possibilities are explored below:

- LECS order—A.1, A.2, B.1, B.2 : (RECOMMENDED) This ordering indicates that upon failure of the primary LECS (A.1), the secondary LECS that will take over operation as the primary will be A.2. i.e. the primary moves from C5000-1 to C5000-2. Since C5000-2 is already up and running the time it takes to cut over a secondary LECS is minimal. Additionally, this order provides minimal flapping between LECSs as well. The next bullet provides a different LECS ordering and illustrates the point about flapping in a clearer way.
- LECS order—A.1, B.1, A.2, B.2 : This ordering indicates that the secondary LECS is the redundant PHY on the same Catalyst 5000, i.e. B.1. Upon failure of the primary LECS (A.1) the next available LECS that takes over operation will be A.2. This is because there is a small time lag before B.1. comes up and becomes active since it is the redundant PHY port. But, once B.1. does come up and becomes active, it takes over ownership of the primary back from A.2. The LANE clients that have already come up using A.2 will need to re-initialize since the state of the network has changed. As cane be clearly seen, there is one extra step involved in the failover scenario involved with the above ordering scheme. It also adds additional recovery time as well as some flapping of the LECSs which is not desirable. Therefore this ordering scheme is not recommended.

An example configuration is shown below.

Table 12 Example LANE services configuration using Cisco default addressing and the Dual PHY LANE module

LECS Table - LS1010

```
atm lecs-address-default 47.0091.8100.0000.0060.5C71.F401.0060.3E9E.FC33.00 /* A.1 */
atm lecs-address-default 47.0091.8100.0000.0060.5C71.F401.0040.0be6.9c33.00 /* B.1 */
atm lecs-address-default 47.0091.8100.0000.0060.3E5C.4701.0060.3E9E.FC33.00 /* A.2 */
atm lecs-address-default 47.0091.8100.0000.0060.3E5C.4701.0040.0be6.9c33.00 /* B.2 */
```

LANE Database Configuration

```
lane database lecsdb
 name elan1 server-atm-address 47.00918100000000605C71F401.00603E9EFC31.00 /* A.1 */
 name elan1 server-atm-address 47.00918100000000605C71F401.00400be69c31.00 /* B.1 */
 name elan1 server-atm-address 47.00918100000000603E5C4701.00603E9EFC31.00 /* A.2 */
 name elan1 server-atm-address 47.00918100000000603E5C4701.00400be69c31.00 /* B.2 */
```

Summary

The redundant network design schemes discussed herein fill a much-needed gap in the ATM Forum specification for LANE and also cover some hardware redundancy issues surrounding LANE networks. Building fault-tolerant and resilient LANE networks is made possible by the use of:

- Simple Server Replication Protocol (SSRP) for LANE services redundancy, which works with Cisco LECs and any third-party LEC
- Hot Standby Router Protocol (HSRP) over LANE, which provides redundancy for the default router configured at IP endstations
- Dual PHY LANE card on the Catalyst 5000 or multiple ATM uplinks on the Catalyst 3000
- Spanning-Tree Protocol on the Ethernet ATM switches

Appendix A. LANE Network Operation

This section provides a detailed understanding of the different pieces of LANE and is intended for network engineers who get involved with troubleshooting a LANE network.

LANE Components

The LANE components include the following:

- LAN Emulation Client (LEC—End systems such as NIC-connected workstations and servers, Catalyst switches, or Cisco routers that support LANE 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 LES and communicates with other LECs via ATM VCCs.

- LAN Emulation Configuration Server (LECS)—The LECS maintains a database of ELANs and the ATM addresses of the LESs that control the ELAN. 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.

Table 13 Sample LECS Database

ELAN Name	LES ATM Address
Finance	47.00091.8100.0000.0800.200c.1001.01
Marketing	47.00091.8100.0000.0800.200c.1001.02

- LAN Emulation Server (LES)—The LES provides a central control point for all LECs. Each LEC maintains 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, and only control information is forwarded on this VCC. As new LECs join the ELAN, they are added as a leaf to this control distribute tree.

Figure 23 (a) LES Connections A-1(b) BUS Connections

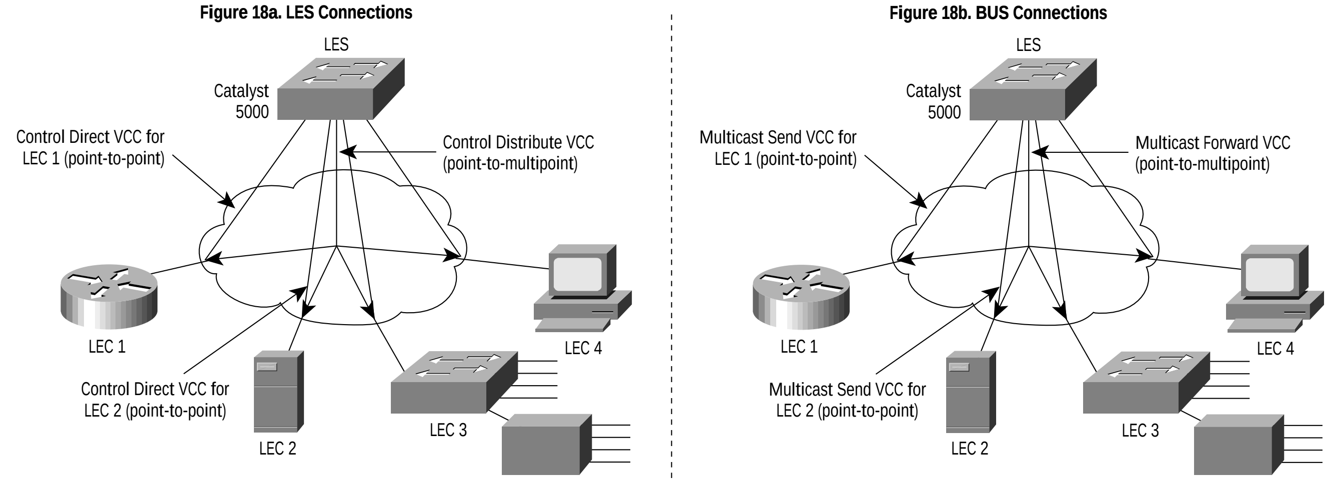

- Broadcast and unknown server (BUS—The BUS acts as a central point to distribute broadcasts and multicasts. ATM is essentially a nonbroadcast multiaccess (NBMA) technology. It does not have "any-to-any" or "broadcast" support, and a special method is needed to implement broadcasting. LANE solved this 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, 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 are in sequence, and when the end of message (EOM) cell arrives, you should just have to reassemble all the cells that have already arrived.

The BUS must take the sequence of cells on each multicast send VCC and reassemble them into frames. When a full frame is received, it is then queued to be sent on the multicast forward VCC. This way all the cells from a particular data frame are guaranteed to be sent in order and not interleaved with cells from any other data frames on the point-to-multipoint VCC. Figure 19 illustrates this process, known as 'serialization.'

Figure 24 Segmentation of Packets into AAL5 Protocol Data Units and Serialization in the BUS

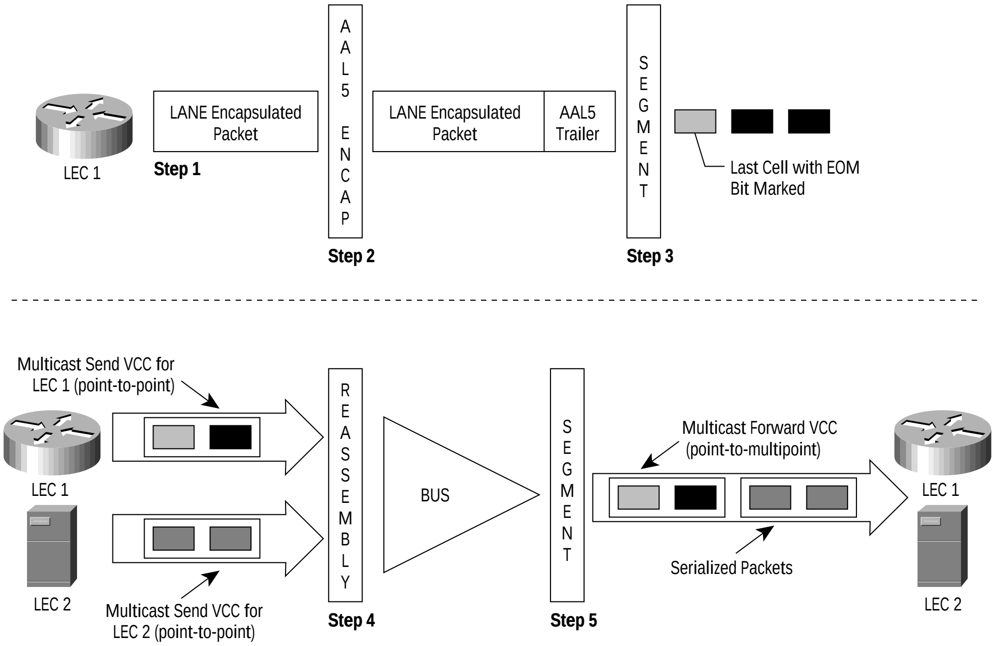

ELAN Operation

An ELAN provides Layer 2 communication between all users on that 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.

Since an ELAN provides Layer 2 communication, it can be equated to a broadcast domain. VLANs can also be thought of as broadcast domains, making it possible to map an ELAN to a VLAN on Layer 2 switches with different VLAN multiplexing technologies such as ISL or 802.10. In addition, IP subnets and IPX networks that are defined on Layer 3-capable devices such as routers also frequently map into broadcast domains (barring secondary addressing), making 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 LECS database as described earlier. ELANs consist of multiple LECs and can be only Ethernet or Token Ring types, but not both at the same time.

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

Registering the ATM Address

The first step for the LEC is to register its ESI with the ATM switch it is attached to and obtain the 13-byte ATM address prefix from the ATM switch. This procedure is known as ‘address registration’ and is a function of the ILMI protocol. This communication typically happens on VPI, VCI : 0,16. Once a LEC is past the address registration phase, it knows its complete ATM address which it uses as the ‘Calling Party address’ in signaling requests. The ATM switch also knows the ATM address of the LEC and hence can route call setups destined to the LEC correctly.

Finding the LECS

In order for a LEC to come up on a given ELAN, first the LEC must find the LECS. Specifically, it looks for the ATM address of the LES that serves the desired ELAN. To find the ATM address of the LECS, the LEC has three choices that it attempts in the following order:

- Query the ATM switch via ILMI. The network administrator is responsible for configuring the ATM switch with the ATM address of the LECS. The LEC can then contact the LECS with UNI signaling.
- Look for a fixed or well-known ATM address that is specified by the ATM Forum as the LECS ATM address, by signaling directly to this address.
- Access PVC 0/17, a "well-known PVC." This option is rarely used and is not included as an option on Cisco LECs.

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 requesting LEC either knows which ELAN it wishes to join and includes this information in the configure request or it just asks the LECS to assign it to an ELAN based on the LECS database. The LEC also includes other information in this request, such as its own MAC address, the type of ELAN (Ethernet or Token Ring), and so forth.

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.

Figure 25 LEC Initialization—LECS Address Discovery and Configure Request Procedure

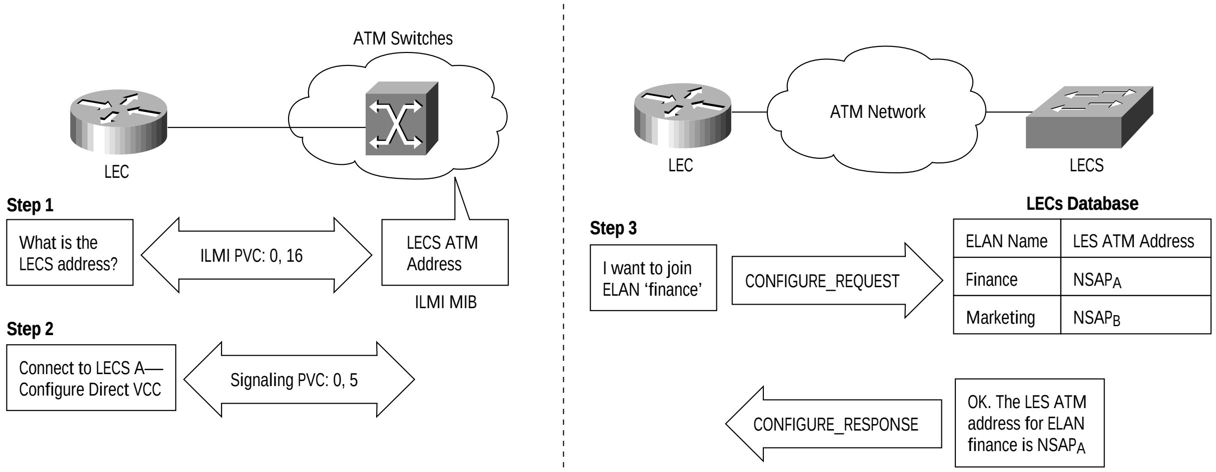

Joining the LES

After the LEC discovers the ATM address of the desired LES, it drops the connection to the LECS[9], 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 register the same MAC or ATM addresses.

Upon receipt of the LE_JOIN_REQUEST, the Cisco 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 identifier (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.

---

[9] Some third-party LECs such as Fore Systems, and so on do not drop their Configure_Direct VCC.

Finding the BUS

Now that 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 a 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 should be 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 point, the LEC is a member of the ELAN and is considered operational. This entire process is illustrated in Figure 21.

LEC-to-LEC Unicast Communication

The real value of LANE is with unicast traffic between LECs. To understand how unicast communication is achieved between two LECs, assume that two IP end stations, A and B, need to talk to each other. For simplicity, let these end stations be directly attached to ATM and be part of the same ‘Ethernet’ ELAN. When A wants to talk to B, it has to go through the following steps:

- Obtain the IP address of B from Domain Name System (DNS) or some other name resolution scheme.
- Issue an IP ARP for the MAC address of B. This packet, a broadcast packet, is sent on the Multicast_Send VCC. The BUS then forwards this broadcast along the Multicast_Forward VCC to all the LECs in the ELAN.

Figure 26 LEC Initialization—Joining the LES and Connecting to the BUS

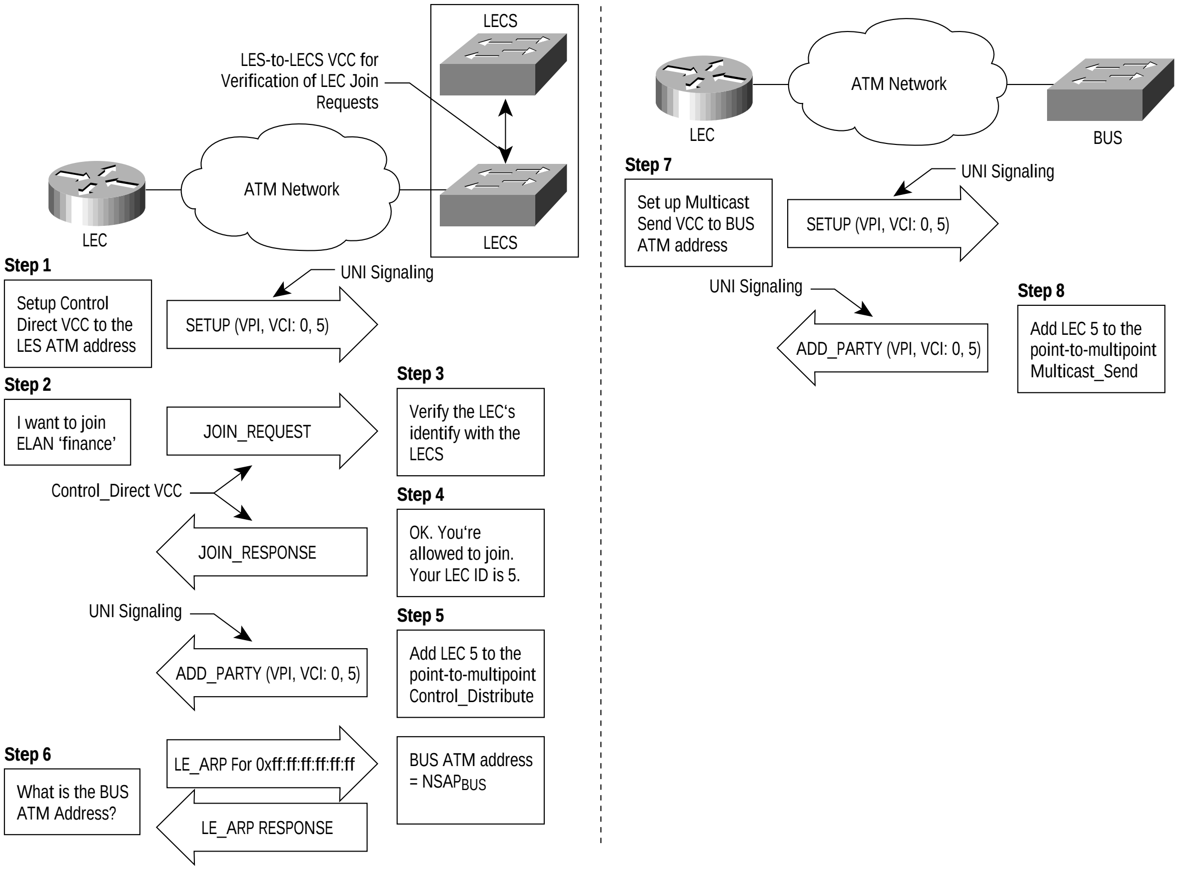

- B receives this IP ARP, recognizes its own IP address in the request, and responds to A as a unicast packet. The LEC on B that needs to transmit this unicast packet recognizes that it has to establish a data-direct VCC to A, but it needs the ATM address of A in order to set up this VCC. To resolve the ATM address of A, B issues a LE_ARP_REQUEST to the LES on the control direct VCC. The LE_ARP_REQUEST contains the MAC address of A. The LES forwards the request on the control distribute VCC so all LEC stations will hear it. In parallel, the unicast data packet (IP ARP response) is sent to the BUS to be forwarded to all endpoints. This "flooding," not the optimal path for unicast traffic, is used only while the LE-ARP procedure is occurring. Since there is a scope for misusing this transmission path, it is rate-controlled to ten packets per second (per the LANE standard [10]). Unicast packets continue using the BUS until the LE_ARP has been resolved.
- A receives the LE_ARP_REQUEST and recognizes its own MAC address in the request. It then issues a LE_ARP_RESPONSE and sends it to the LES, which forwards it on the control distribute VCC so all LECs can learn the new "MAC to ATM" address binding. The LEC on station B also learns the binding.
- B now issues a UNI setup message to the ATM address of A on the signaling PVC 0,5.
- While waiting for the LE_ARP resolution and setting up the data direct VCC, the LEC has been forwarding the 'unknown' unicast packets through the BUS path. When the data direct VCC becomes available, if the LEC switches immediately to the new path, it runs the risk of packets arriving out of order at the destination. This possibility arises since the BUS path is slower than this new 'optimal' path. A packet sent on the BUS path takes longer to reach the destination than a packet that takes the optimal path. The LANE standard has provided for a 'flush protocol' to guard against this situation.

  Under this protocol, after the data-direct VCC becomes available, the LEC generates a flush packet and sends it to the BUS along its Multicast_Send. The destination LEC (A) receives this flush request and responds by sending this packet back as a flush response along the control path (LES). When the source LEC (B) receives its own flush packet on the control distribute VCC, it knows that all previously sent unicasts must have already been forwarded. It should now be safe to begin using the data-direct VCC.

Most of the unicast traffic in a LANE network rides on these data direct VCCs.

The Proxy LEC

A LEC when implemented at an end station needs to represent only one MAC address in the ELAN. It responds to any LE_ARP_REQUESTS for its own MAC address. However, when a LEC is implemented in a bridge/switch, such a device needs to represent (or proxy for) multiple MAC addresses since there are multiple Ethernet end stations attached to it; that is, it needs to respond to multiple LE_ARP_REQUESTS. Hence, these LECs are known as proxy LECs.

LECs on Layer 2 switches such as the Catalyst 5000, 3000, and so on are all proxy LECs. However, LECs on the routers are simple LECs except when bridging is configured for that subinterface on the router.

Bibliography

1. ATM Internetworking—White Paper—Anthony Alles
2. The Next-Generation ATM Switch: From Testbeds to Production Networks—White Paper—Anthony Alles—http://wwwin.cisco.com/Mkt/cmc/cc/cisco/mkt/switch/ls1010/ngatm_wp.htm
3. Designing ATM networks using the Catalyst 3000 ATM module—Implementation Guide—Marcus Phipps—http://wbu-pme/mphipps/3katm-ig.htm.
4. VLAN Technologies Handbook—Design Guide—Merwyn Andrade.

---

[10] Note that the rate control applies only to unknown unicast packets. Broadcast and multicast packets are not rate controlled and are limited only by the peak performance of the BUS.

CISCO SYSTEMS

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
World Wide Web URL:
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic-Batiment L1/L2
16, Avenue du Quebec
BP 706-Villebon
91961 Courtaboeuf Cedex
France
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
Tel: 408 526-7660
Fax: 408 526-4646

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
Tel: 81 3 5219 6000
Fax: 81 3 5219 6010

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the
Cisco Connection Online Web site at http://www.cisco.com.

Argentina · Australia · Austria · Belgium · Brazil · Canada · Chile · China (PRC) · Colombia · Costa Rica · Czech Republic · Denmark
Finland · France · Germany · Hong Kong · Hungary · India · Indonesia · Ireland · Israel · Italy · Japan · Korea · Malaysia · Mexico
The Netherlands · New Zealand · Norway · Philippines · Poland · Portugal · Russia · Singapore · South Africa · Spain · Sweden
Switzerland · Taiwan, ROC · Thailand · United Arab Emirates · United Kingdom · Venezuela

Copyright © 1997 Cisco Systems, Inc. All rights reserved. Printed in USA. AtmDirector, AutoConnect, AutoRoute, AXIS, BPX, Catalyst, CD-PAC, CiscoFusion, Cisco IOS, the Cisco IOS logo, *CiscoLink*, CiscoPro, the CiscoPro logo, CiscoRemote, the CiscoRemote logo, CiscoSecure, Cisco Systems, CiscoView, CiscoVision, CiscoWorks, ClickStart, ControlStream, EdgeConnect, EtherChannel, FairShare, FastCell, FastForward, FastManager, FastMate, FastPADlmp, FastPADmicro, FastPADmp, FragmentFree, FrameClass, Fulcrum INS, IGX, Impact, Internet Junction, JumpStart, LAN^{2}LAN Enterprise, LAN^{2}LAN Remote Office, LightSwitch, MICA, NetBeyond, NetFlow, Newport Systems Solutions, *Packet*, PIX, Point and Click Internetworking, RouteStream, Secure/IP, SMARTnet, StrataSphere, StrataSphere BILLder, StrataSphere Connection Manager, StrataSphere Modeler, StrataSphere Optimizer, Stratm, StrataView Plus, StreamView, SwitchProbe, SwitchVision, SwitchWare, SynchroniCD, *The Cell*, The FastPacket Company, TokenSwitch, TrafficDirector, Virtual EtherSwitch, VirtualStream, VlanDirector, Web Clusters, WNIC, Workgroup Director, Workgroup Stack, and XCI are trademarks; Access by Cisco, Bringing the Power of Internetworking to Everyone, Enter the Net with MultiNet., and The Network Works. No Excuses. are service marks; and Cisco, the Cisco Systems logo, CollisionFree, Combinet, EtherSwitch, FastHub, FastLink, FastNIC, FastPacket, FastPAD, FastSwitch, ForeSight, Grand, Grand Junction, Grand Junction Networks, the Grand Junction Networks logo, HSSI, IGRP, IPX, Kalpana, the Kalpana logo, LightStream, MultiNet, MultiWare, OptiClass, Personal Ethernet, Phase/IP, RPS, StrataCom, TGV, the TGV logo, and UniverCD are registered trademarks of Cisco Systems, Inc. All other trademarks, service marks, registered trademarks, or registered service marks mentioned in this document are the property of their respective owners. 1296R