

Designing Switched LAN Internetworks

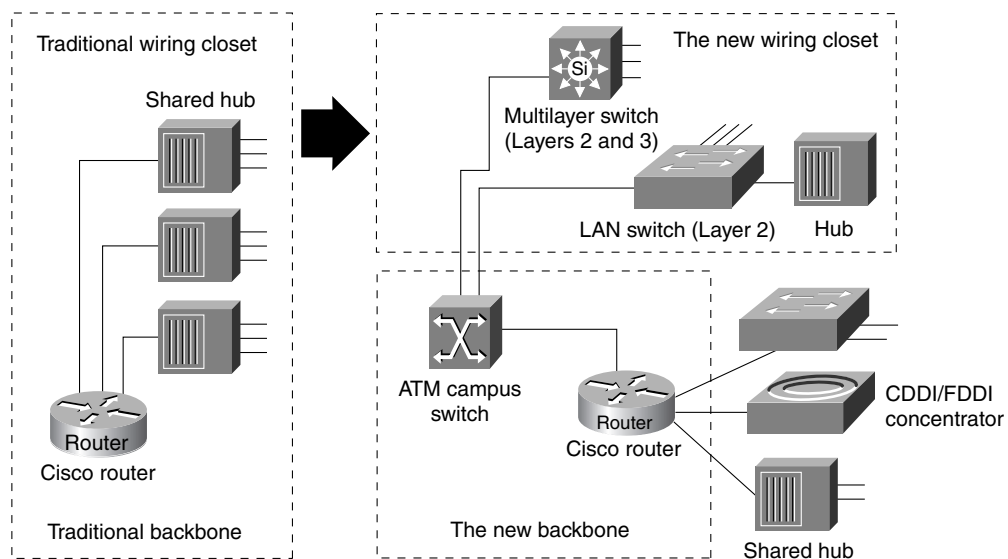
This chapter describes the following three technologies that network designers can use to design switched LAN internetworks:

- LAN switching
- Virtual LANs (VLANs)
- ATM switching

Evolution from Shared to Switched Networks

In the past, network designers had only a limited number of hardware options when purchasing a technology for their campus networks. Hubs were for wiring closets and routers were for the data center or main telecommunications operations. The increasing power of desktop processors and the requirements of client-server and multimedia applications, however, have driven the need for greater bandwidth in traditional shared-media environments. These requirements are prompting network designers to replace hubs in their wiring closets with switches, as shown in Figure 12-1.

Figure 12-1 Evolution from shared to switched internetworks.



This strategy allows network managers to protect their existing wiring investments and boost network performance with dedicated bandwidth to the desktop for each user. Coinciding with the wiring closet evolution is a similar trend in the network backbone. Here, the role of Asynchronous

Transfer Mode (ATM) is increasing as a result of standardizing protocols, such as LAN emulation (LANE), that enable ATM devices to coexist with existing LAN technologies. Network designers are collapsing their router backbones with ATM switches, which offer the greater backbone bandwidth required by high-throughput data services.

Technologies for Building Switched LAN Internetworks

With the advent of such technologies as Layer 3 switching, LAN switching, and VLANs, building campus LANs is becoming more complex than in the past. Today, the following three technologies are required to build successful campus networks:

- LAN switching technologies
 - Ethernet switching—Provides Layer 2 switching and offers broadcast domain segmentation using VLANs. This is the base fabric of the network.
 - Token Ring switching—Offers the same functionality as Ethernet switching but uses Token Ring technology. You can use a Token Ring switch as either a transparent bridge or source-route bridge.
 - Copper Data Distributed Interface (CDDI)—Provides a single-attachment station (SAS) or dual-attachment station (DAS) to two Category 5 unshielded twisted-pair (UTP), 100 Mbps RJ-45 connectors.
 - Fiber Distributed Data Interface (FDDI)—Provides an SAS or DAS connection to the FDDI backbone network using two multimode, media interface connector (MIC) fiber-optic connections.
- ATM switching technologies

ATM switching offers high-speed switching technology for voice, video, and data. Its operation is similar to LAN switching technologies for data operations. ATM, however, offers superior voice, video, and data integration today.

- Routing technologies

Routing is a key technology for connecting LANs in a campus network. It can be either Layer 3 switching or more traditional routing with Layer 3 switching features and enhanced Layer 3 software features.

Note Switched LAN internetworks are also referred to as *campus LANs*.

Role of LAN Switching Technology in Campus Networks

Most network designers are beginning to integrate switching devices into their existing shared-media networks to achieve the following goals:

- Increase the bandwidth that is available to each user, thereby alleviating congestion in their shared-media networks.
- Employ the manageability of VLANs by organizing network users into logical workgroups that are independent of the physical topology of wiring closet hubs. This, in turn, can reduce the cost of moves, adds, and changes while increasing the flexibility of the network.
- Deploy emerging multimedia applications across different switching platforms and technologies, making them available to a variety of users.

- Provide a smooth evolution path to high-performance switching solutions, such as Fast Ethernet and ATM.

Segmenting shared-media LANs divides the users into two or more separate LAN segments, reducing the number of users contending for bandwidth. LAN switching technology, which builds upon this trend, employs *microsegmentation*, which further segments the LAN to fewer users and ultimately to a single user with a dedicated LAN segment. Each switch port provides a dedicated, 10MB Ethernet segment, or dedicated 4/16MB Token Ring segment.

Segments are interconnected by internetworking devices that enable communication between LANs while blocking other types of traffic. Switches have the intelligence to monitor traffic and compile address tables, which then allows them to forward packets directly to specific ports in the LAN. Switches also usually provide nonblocking service, which allows multiple conversations (traffic between two ports) to occur simultaneously.

Switching technology is quickly becoming the preferred solution for improving LAN traffic for the following reasons:

- Unlike hubs and repeaters, switches allow multiple data streams to pass simultaneously.
- Switches have the capability through microsegmentation to support the increased speed and bandwidth requirements of emerging technologies.
- Switches deliver dedicated bandwidth to users through high-density group switched and switched 10BaseT or 100BaseT Ethernet, flexible 10/100 BaseT Ethernet, fiber-based Fast Ethernet, Fast EtherChannel, Token Ring, CDDI/FDDI, and ATM LAN Emulation (LANE).

Switched Internetwork Solutions

Network designers are discovering, however, that many products offered as switched internetwork solutions are inadequate. Some offer a limited number of hardware platforms with little or no system integration with the current infrastructure. Others require complete abandonment of all investments in the current network infrastructure. To be successful, a switched internetwork solution must accomplish the following:

- Leverage strategic investments in the existing communications infrastructure while increasing available bandwidth.
- Reduce the costs of managing network operations.
- Offer options to support multimedia applications and other high-demand traffic across a variety of platforms.
- Provide scalability, traffic control, and security that is at least as good or better than that of today's router-based internetworks.
- Provide support for embedded remote monitoring (RMON) agent.

The key to achieving these benefits is to understand the role of the internetworking software infrastructure within the switched internetworks. Within today's networks, routers allow for the interconnection of disparate LAN and WAN technologies, while also implementing security filters and logical firewalls. It is these capabilities that have allowed current internetworks to scale globally while remaining stable and robust.

As networks evolve toward switched internetworks, similar logical internetworking capabilities are required for stability and scalability. Although LAN and ATM switches provide great performance improvements, they also raise new internetworking challenges. Switched internetworks must integrate with existing LAN and WAN networks. Such services as VLANs, which will be deployed with switched internetworks, also have particular internetworking requirements.

A true switched internetwork, therefore, is more than a collection of boxes. Rather, it consists of a system of devices integrated and supported by an intelligent internetworking software infrastructure. Presently, this network intelligence is centralized within routers. However, with the advent of switched internetworks, the intelligence will often be dispersed throughout the network, reflecting the decentralized nature of switching systems. The need for an internetworking infrastructure, however, will remain.

Components of the Switched Internetworking Model

A switched internetwork is composed of the following three basic components:

- Physical switching platforms
- A common software infrastructure
- Network management tools and applications

Cisco provides network designers with a complete, end-to-end solution for implementing and managing scalable, robust, switched internetworks.

Scalable Switching Platforms

The first component of the switched internetworking model is the physical switching platform. This can be an ATM switch, a LAN switch, or a router.

ATM Switches

Although switched internetworks can be built with a variety of technologies, many network designers will deploy ATM in order to utilize its unique characteristics. ATM provides scalable bandwidth that spans both LANs and WANs. It also promises Quality of Service (QoS) guarantees—bandwidth on demand—that can map into and support higher-level protocol infrastructures for emerging multimedia applications and provide a common, multiservice network infrastructure.

ATM switches are one of the key components of ATM technology. All ATM switches, however, are not alike. Even though all ATM switches perform cell relay, ATM switches differ markedly in the following capabilities:

- Variety of interfaces and services that are supported
- Redundancy
- Depth of ATM internetworking software
- Sophistication of traffic management mechanism
- Blocking and non-blocking switching fabrics
- SVC and PVC support

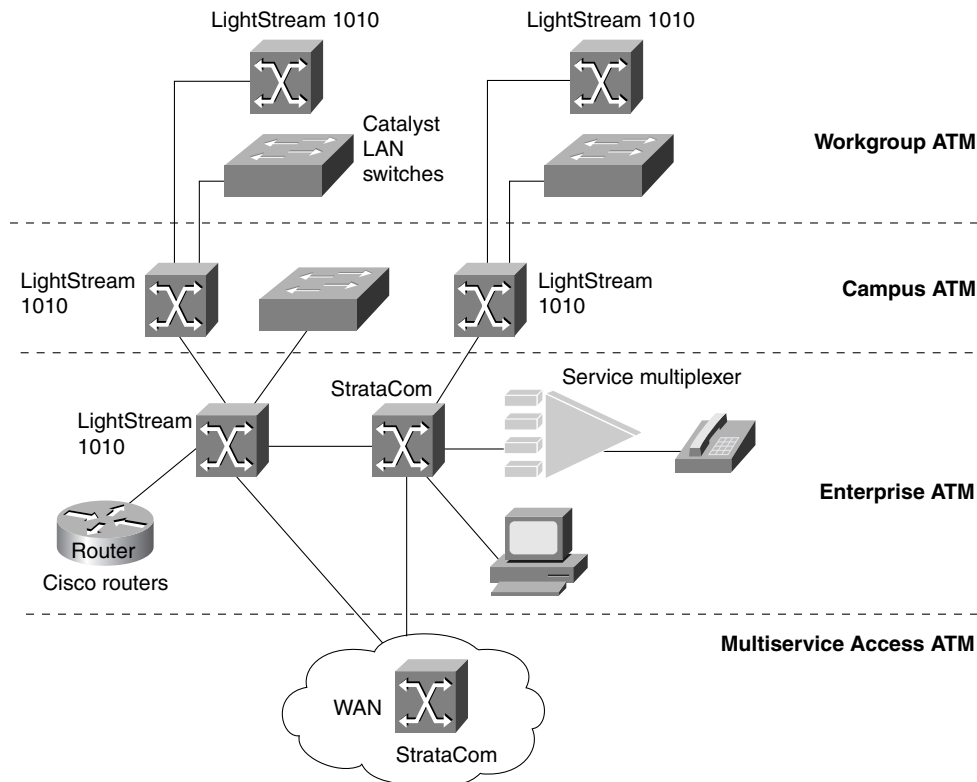
Just as there are routers and LAN switches available at various price/performance points with different levels of functionality, ATM switches can be segmented into the following four distinct types that reflect the needs of particular applications and markets:

- Workgroup ATM switches
- Campus ATM switches
- Enterprise ATM switches

- Multiservice access switches

As Figure 12–2 shows, Cisco offers a complete range of ATM switches.

Figure 12-2 Different types of ATM switches.



Workgroup and Campus ATM Switches

Workgroup ATM switches are optimized for deploying ATM to the desktop over low-cost ATM desktop interfaces, with ATM signaling interoperability for ATM adapters and QoS support for multimedia applications.

Campus ATM switches are generally used for small-scale ATM backbones (for example, to link ATM routers or LAN switches). This use of ATM switches can alleviate current backbone congestion while enabling the deployment of such new services as VLANs. Campus switches need to support a wide variety of both local backbone and WAN types but be price/performance optimized for the local backbone function. In this class of switches, ATM routing capabilities that allow multiple switches to be tied together is very important. Congestion control mechanisms for optimizing backbone performance is also important.

Enterprise ATM Switches

Enterprise ATM switches are sophisticated multiservice devices that are designed to form the core backbones of large, enterprise networks. They are intended to complement the role played by today's high-end multiprotocol routers. Enterprise ATM switches, much as campus ATM switches, are used to interconnect workgroup ATM switches and other ATM-connected devices, such as LAN switches. Enterprise-class switches, however, can act not only as ATM backbones but can serve as the single point of integration for all of the disparate services and technology found in enterprise backbones


today. By integrating all of these services onto a common platform and a common ATM transport infrastructure, network designers can gain greater manageability while eliminating the need for multiple overlay networks.

LAN Switches

A LAN switch is a device that typically consists of many ports that connect LAN segments (Ethernet and Token Ring) and a high-speed port (such as 100-Mbps Ethernet, Fiber Distributed Data Interface [FDDI], or 155-Mbps ATM). The high-speed port, in turn, connects the LAN switch to other devices in the network.

A LAN switch has dedicated bandwidth per port, and each port represents a different segment. For best performance, network designers often assign just one host to a port, giving that host dedicated bandwidth of 10 Mbps, as shown in Figure 12–3, or 16 Mbps for Token Ring networks.

Figure 12-3 Sample LAN switch configuration.



When a LAN switch first starts up and as the devices that are connected to it request services from other devices, the switch builds a table that associates the MAC address of each local device with the port number through which that device is reachable. That way, when Host A on Port 1 needs to transmit to Host B on Port 2, the LAN switch forwards frames from Port 1 to Port 2, thus sparing other hosts on Port 3 from responding to frames destined for Host B. If Host C needs to send data to Host D at the same time that Host A sends data to Host B, it can do so because the LAN switch can forward frames from Port 3 to Port 4 at the same time it forwards frames from Port 1 to Port 2.

Whenever a device connected to the LAN switch sends a packet to an address that is not in the LAN switch's table (for example, to a device that is beyond the LAN switch), or whenever the device sends a broadcast or multicast packet, the LAN switch sends the packet out all ports (except for the port from which the packet originated)—a technique known as *flooding*.

Because they work like traditional “transparent” bridges, LAN switches dissolve previously well-defined workgroup or department boundaries. A network built and designed only with LAN switches appears as a *flat* network topology consisting of a single broadcast domain. Consequently, these networks are liable to suffer the problems inherent in flat (or *bridged*) networks—that is, they do not scale well. Note, however, that LAN switches that support VLANs are more scalable than traditional bridges.

Multiservice Access Switches

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

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

Routing Platforms

In addition to LAN switches and ATM switches, typically network designers use routers as one of the components in a switched internetwork infrastructure. While LAN switches are being added to wiring closets to increase bandwidth and to reduce congestion in existing shared-media hubs, high-speed backbone technologies, such as ATM switching and ATM routers are being deployed in the backbone. Within a switched internetwork, routing platforms also allow for the interconnection of disparate LAN and WAN technologies while also implementing broadcast filters and logical firewalls. In general, if you need advanced internetworking services, such as broadcast firewalling and communication between dissimilar LANs, routers are necessary.

Common Software Infrastructure

The second level of a switched internetworking model is a common software infrastructure. The function of this software infrastructure is to unify the variety of physical switching platforms: LAN switches, ATM switches, and multiprotocol routers. Specifically, the software infrastructure should perform the following tasks:

- Monitor the logical topology of the network
- Logically route traffic
- Manage and control sensitive traffic
- Provide firewalls, gateways, filtering, and protocol translation


Cisco offers network designers Cisco Internetwork Operating System (Cisco IOS) switching software. This subset of the Cisco IOS software is optimized for switching and provides the unifying element to Cisco's line of switching platforms in a switched internetwork. The Cisco IOS software is found on standalone routers, router modules for shared-media hubs, PC and workstations file servers, multiservice WAN access switches, LAN switches, ATM switches, and ATM-capable PBXs. It provides optional levels of routing and switching across a switched internetwork in addition to new capabilities, such as VLANs, ATM internetworking software services, multilayer switching, extensions to support new networked multimedia applications, and traffic management and analysis tools.

VLANs

A VLAN consists of several end systems, either hosts or network equipment (such as switches and routers), all of which are members of a single logical broadcast domain. A VLAN no longer has physical proximity constraints for the broadcast domain. This VLAN is supported on various pieces of network equipment (for example, LAN switches) that support VLAN trunking protocols between them. Each VLAN supports a separate Spanning Tree (IEEE 802.1d).

First-generation VLANs are based on various OSI Layer 2 bridging and multiplexing mechanisms, such as IEEE 802.10, LAN Emulation (LANE), and Inter-Switch Link (ISL), that allow the formation of multiple, disjointed, overlaid broadcast groups on a single network infrastructure. Figure 12-4 shows an example of a switched LAN network that uses VLANs. Layer 2 of the OSI reference model provides reliable transit of data across a physical link. The data link layer is concerned with physical addressing, network topology, line discipline, error notification, ordered delivery frames, and flow control. The IEEE has divided this layer into two sublayers: the MAC sublayer and the LLC sublayer, sometimes simply called link layer.

Figure 12-4 Typical VLAN topology.



In Figure 12-4, 10-Mbps Ethernet connects the hosts on each floor to switches A, B, C, and D. 100-Mbps Fast Ethernet connects these to Switch E. VLAN 10 consists of those hosts on Ports 6 and 8 of Switch A and Port 2 on Switch B. VLAN 20 consists of those hosts that are on Port 1 of Switch A and Ports 1 and 3 of Switch B.

VLANs can be used to group a set of related users, regardless of their physical connectivity. They can be located across a campus environment or even across geographically dispersed locations. The users might be assigned to a VLAN because they belong to the same department or functional team, or because data flow patterns among them is such that it makes sense to group them together. Note, however, that without a router, hosts in one VLAN cannot communicate with hosts in another VLAN.

Network Management Tools and Applications

The third and last component of a switched internetworking model consists of network management tools and applications. As switching is integrated throughout the network, network management becomes crucial at both the workgroup and backbone levels. Managing a switch-based network requires a radically different approach than managing traditional hub and router-based LANs.

As part of designing a switched internetwork, network designers must ensure that their design takes into account network management applications needed to monitor, configure, plan, and analyze switched internetwork devices and services. Cisco offers such tools for emerging switched internetworks.

Cisco Switched Internetworking Products

Cisco offers the following products that meet the needs of a switched internetwork, all discussed in the following sections:

- Cisco ATM Switches
- Cisco LAN Switches
- Cisco Routing Platforms

Cisco ATM Switches

Cisco's LightStream 1010 family of ATM switches is specifically designed for workgroup and campus backbone deployment. It incorporates support for the latest ATM Forum specifications and builds upon the Cisco IOS software. The LightStream 1010 is a five-slot, modular switch that features the option of dual, load-sharing, hot-swappable power supplies, 5-Gbps of shared memory, nonblocking switch fabric, and 65,536 cells of shared ATM cell buffers. It supports a wide range of modular, hot-swappable, desktop, backbone, and wide-area ATM interfaces. These characteristics allow network managers to deploy it in a variety of scenarios ranging from high-density, 155-Mbps copper UTP-5 workgroups to high-performance OC-12 backbones.

In order to support the bursty, best-effort traffic generated by LAN switches and routes, the LightStream 1010 provides advanced traffic management mechanisms. The LightStream 1010's intelligent early packet discard mechanism allows it to discard entire packets rather than individual cells when necessary, which greatly increases performance for current protocols, such as TCP/IP and IPX. It also supports the latest ATM Forum Available Bit Rate (ABR) congestion control specifications, which allows the LightStream 1010 to slow traffic sources before congestion becomes excessive. Because of its support for the ATM Forum private network-network interface (PNNI) protocols, networks of LightStream 1010s can scale to hundreds of nodes.

In addition, the LightStream 1010 offers a high degree of manageability. Advanced port snooping and connection-steering capabilities allow the connections on any port to be directed to a monitor port for analysis by an external ATM analyzer. This capability is critical for the monitoring and troubleshooting of ATM switching systems, which unlike shared-media LANs, cannot be monitored easily with external devices. Simple Network Management Protocol (SNMP) monitoring and configuration invoked through the CiscoView graphical user interface (GUI) device configuration applications and the AtmDirector CiscoWorks ATM system management application, allow for comprehensive network management.

By building on the Cisco IOS software, the LightStream 1010 switch also shares the advanced serviceability capabilities found today on Cisco's multiprotocol routers. As with all Cisco routers, the LightStream 1010 switch supports such protocols as BOOTP, DHCP, Telnet, and Trivial File Transfer Protocol (TFTP) for remote access and autoconfiguration. It also offers the access protections of the Cisco IOS software, from multiple password levels to TACACS for remote access validation, to preclude unauthorized changes to the switch configuration. These capabilities are clearly essential to safeguard the operation of the mission-critical campus backbones in which the LightStream 1010 will typically be deployed.

The Cisco/StrataCom BPX/AXIS is a powerful broadband 9.6-Gbps ATM switch designed to meet the demanding, high-traffic needs of a large private enterprise or public service provider. The Cisco/StrataCom IGX is a 1.2-Gbps ATM-based enterprise WAN switch that can be used to provide enterprise WAN features in your internetwork. For more information on these enterprise ATM switches, see Chapter 8, “Designing ATM Internetworks.”

Cisco LAN Switches

Cisco’s Catalyst family is a comprehensive line of high-performance switches designed to help network managers easily migrate from traditional shared LANs to fully switched internetworks. The Catalyst family delivers the varying levels of flexibility and cost-effectiveness required for today’s desktop, workgroup, and backbone applications while enabling enterprise-wide switched internetworks. Using these LAN switches instead of traditional shared hubs increase performance and provides new capabilities, such as VLANs.

Figure 12–5 shows an example of switches that can be used in a campus backbone. In this example, the Cisco switches are used to interconnect the four buildings that comprise the campus network.

Figure 12-5 LAN switches in a campus backbone.




Table 12–1 summarizes the LAN switches that Cisco offers.

Table 12-1 List of Cisco LAN Switches

| Cisco LAN Switch | Description |
|--|--|
| Catalyst 5500 switching system | <p>The Catalyst 5500 switch chassis has 13 slots. Slot 1 is for the Supervisor Engine II model which provides switching, local and remote management, and dual Fast Ethernet interfaces. Slot 2 contains an additional redundant Supervisor Engine II in case the first module fails.</p> <p>A failure of the active Supervisor Engine II is detected by the standby module, which takes control of the Supervisor Engine II switching functions. If a redundant Supervisor Engine II is not required, slot 2 is available for any interface module.</p> <p>The Catalyst 5500 has a 3.6-Gbps media-independent switch fabric and a 5-Gbps cell-switch fabric. The backplane provides the connection between power supplies, Supervisor Engine II, interface modules, and backbone module. The 3.6-Gbps media-independent fabric supports Ethernet, Fast Ethernet, FDDI/CDDI, ATM LAN Emulation, and RSM modules. The 5-Gbps cell-based fabric supports an ATM switch processor (ASP) module and ATM port adapter modules (PAMS).</p> |
| Route switch module | <p>The Cisco Catalyst 5000 series route switch module builds upon the Route Switch Processor (RSP) featured in Cisco's 7500 routing platform. The route switch module provides high-performance multilayer switching and routing services between switched Virtual LANs (VLANs), emulated LANs (ELANs) within an Asynchronous Transfer Mode (ATM) fabric, or across mixed media via an optional Versatile Interface Processor (VIP) and port adapters.</p> |
| Catalyst 5000 switching system | <p>A modular switching platform that meets high performance needs, bandwidth-intensive networking switching applications. It offers five slots that can be populated with any combination of 10BaseT, 10BaseFL modules, switched 10-Mbps Fast Ethernet, FDDI, or ATM modules.</p> <p>It delivers high performance both for client and server connections as well as for backbone connections. Its switching backplane operates at 1.2 Gbps and provides nonblocking performance for all switched 10-Mbps Ethernet interfaces.</p> <p>It supports enterprise-wide VLAN communications across Ethernet, Fast Ethernet, CDDI/FDDI, and ATM connections via the following protocols: ISL for Fast Ethernet interfaces, 802.10 for FDDI interfaces, and LANE v1.0 for ATM.</p> |
| Catalyst 3000 stackable switching system | <p>A 16-port 10BaseT switch that has two open expansion bays that can be populated with 100BaseTX/FX, 10BaseFL, 10BaseT, 100VG-AnyLAN, or ATM. With the Matrix module, up to eight Catalyst 3000 switches can be stacked together as one logical switching system. The Catalyst 3000 system can also be populated with the Catalyst 3011 WAN router module.</p> <p>A fully loaded Catalyst 3000 system can support up to 192 10BaseT ports, or 128 10BaseT ports with 16 high-speed ports. Supports up to 64 VLANs within the stack. Also supports ISL for Fast Ethernet, and ATM LANE.</p> |
| Catalyst 2900 Fast Ethernet switch | <p>A 14-port, fixed-configuration, Fast Ethernet switch that provides media-rate Fast Ethernet switching in backbone, server cluster, and high-performance workgroup applications. Its software architecture combines superior traffic throughput, complete VLAN solutions, traffic management, and fault tolerance.</p> |
| Catalyst 1800 Token Ring switch | <p>A Token Ring switch that has 16 dedicated or shared ports in the base unit plus two feature-card slots that is designed for the workgroup switching environment. Using the four-port Token Ring unshielded twisted-pair/shielded twisted-pair (UTP/STP) feature cards, it supports eight additional Token Ring ports.</p> |

| Cisco LAN Switch | Description |
|---|---|
| Catalyst 1900 and Catalyst 2820 Ethernet switches | <p>Ideally suited to replace shared 10BaseT hubs in the wiring closet with feature-rich switched Ethernet capability to the desktop. The Catalyst 1900 Ethernet switch features 25 switched Ethernet ports providing attachment to individual workstations and 10BaseT hubs. It also has two 100BaseT ports for high speed connectivity to servers and backbones.</p> <p>The Catalyst 2820 Ethernet switch has 25 switched Ethernet ports and two high-speed expansion slots. Field-installable modules provide configuration, wiring, and backbone flexibility with a choice of 100BaseT, FDDI, and future ATM modules available, which support Category 5 UTP or fiber-optic cabling.</p> |
| Catalyst 1200 workgroup switch | <p>A multilayer switch for workgroup applications that can benefit from OSI Layer 3 as well as Layer 2 capabilities. It offers eight 10BaseT or 10BaseFL ports and one expansion slot. The expansion slot can be populated with either one A/B CDDI interface or one A/B FDDI interface. IP routing only supports 802.10 over FDDI VLANs. In addition to meeting a wide range of performance needs for Ethernet and FDDI, it offers such unique features as embedded Remote Monitoring (RMON) functionality, which helps network managers monitor and control the growth and changes of client-server workgroups.</p> |

Cisco Routing Platforms

Both the Cisco 7000 and Cisco 4000 family of multiprotocol routers are particularly well suited for switched internetworking. In particular, the first native-mode ATM router interface, the ATM Interface Processor (AIP) for the Cisco 7000 family of routers, is a key enabler for integrating existing LAN and WAN networks with evolving, ATM-based switched internetworks.

The sophisticated ATM signaling and traffic management capabilities of the AIP also allows it to play a crucial role in the deployment of new services such as VLANs. The AIP, a key enabler for the production deployment of switched internetworks, allows VLANs to internetwork either with each other or with external networks. The Cisco 4000 family of multiprotocol routers also support such capabilities, thereby providing network designers with a wide choice of price/performance points for ATM-capable routers.

Because the Cisco 7000 and Cisco 4000 families support FDDI, Fast Ethernet, and ATM, they provide network designers with a full set of options for high-speed connectivity. Both router families also support routing between VLANs on all media for ease of migration.

Switched LAN Network Designs

A successful switched internetworking solution must combine the benefits of both routers and switches in every part of the network, as well as offer a flexible evolution path from shared-media networking to switched internetworks.

In general, incorporating switches in campus network designs results in the following benefits:

- High bandwidth
- Quality of Service (QoS)
- Low cost
- Easy configuration

If you need advanced internetworking services, however, routers are necessary. Routers offer the following services:

- Broadcast firewalling
- Hierarchical addressing

- Communication between dissimilar LANs
- Fast convergence
- Policy routing
- QoS routing
- Security
- Redundancy and load balancing
- Traffic flow management
- Multimedia group membership

Some of these router services will be offered by switches in the future. For example, support for multimedia often requires a protocol, such as Internet Group Management Protocol (IGMP), that allows workstations to join a group that receives multimedia multicast packets. In the future, Cisco will allow switches to participate in this process by using the Cisco Group Management Protocol (CGMP). One router will still be necessary, but you will not need a router in each department because CGMP switches can communicate with the router to determine whether any of their attached users are part of a multicast group.

Switching and bridging sometimes can result in nonoptimal routing of packets. This is because every packet must go through the root bridge of the spanning tree. When routers are used, the routing of packets can be controlled and designed for optimal paths. Cisco now provides support for improved routing and redundancy in switched environments by supporting one instance of the spanning tree per VLAN.

When designing switched LAN networks, you should consider the following:

- Comparison of LAN Switches and Routers
- Benefits of LAN Switches (Layer 2 Services)
- Benefits of Routers (Layer 3 Services)
- Benefits of VLANs
- VLAN Implementation
- General Network Design Principles
- Switched LAN Network Design Principles

Comparison of LAN Switches and Routers

The fundamental difference between a LAN switch and a router is that the LAN switch operates at Layer 2 of the OSI model and the router operates at Layer 3. This difference affects the way that LAN switches and routers respond to network traffic. This section compares LAN switches and routers with regard to the following network design issues:


- Loops
- Convergence
- Broadcasts
- Subnetworking
- Security
- Media Dependence

Note Because routers implement Layer 2 functionality and switches are beginning to implement Layer 3 functionality, the functions of a LAN switch and a router are merging.

Loops

Switched LAN topologies are susceptible to loops, as shown in Figure 12–6.

Figure 12-6 Switched LAN topology with loops.



In Figure 12–6, it is possible for packets from Client X to be switched by Switch A and then for Switch B to put the same packet back on to LAN 1. In this situation, packets loop and undergo multiple replications. To prevent looping and replication, topologies that may contain loops need to run the Spanning-Tree Protocol. The Spanning-Tree Protocol uses the spanning-tree algorithm to construct topologies that do not contain any loops. Because the spanning-tree algorithm places certain connections in blocking mode, only a subset of the network topology is used for forwarding data. In contrast, routers provide freedom from loops and make use of optimal paths.

Convergence

In transparent switching, neighboring switches make topology decisions locally based on the exchange of Bridge Protocol Data Units (BPDUs). This method of making topology decisions means that convergence on an alternative path can take an order of magnitude longer than in a routed environment.

In a routed environment, sophisticated routing protocols, such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (Enhanced IGRP), maintain concurrent topological databases of the network and allow the network to converge quickly.

Broadcasts

LAN switches do not filter broadcasts, multicasts, or unknown address frames. The lack of filtering can be a serious problem in modern distributed networks in which broadcast messages are used to resolve addresses and dynamically discover network resources such as file servers. Broadcasts originating from each segment are received by every computer in the switched internetwork. Most devices discard broadcasts because they are irrelevant, which means that large amounts of bandwidth are wasted by the transmission of broadcasts.

In some cases, the circulation of broadcasts can saturate the network so that there is no bandwidth left for application data. In this case, new network connections cannot be established, and existing connections may be dropped (a situation known as a *broadcast storm*). The probability of broadcast storms increases as the switched internetwork grows. Routers do not forward broadcasts, and, therefore, are not subject to broadcast storms. For more information about the impact of broadcasts, see Appendix E, “Broadcasts in Switched LAN Internetworks.”

Subnetworking

Transparently switched internetworks are composed of physically separate segments, but are logically considered to be one large network (for example, one IP subnet). This behavior is inherent to the way that LAN switches work—they operate at OSI Layer 2 and have to provide connectivity to hosts as if each host were on the same cable. Layer 2 addressing assumes a flat address space with universally unique addresses.

Routers operate at OSI Layer 3, so can formulate and adhere to a hierarchical addressing structure. Routed networks can associate a logical addressing structure to a physical infrastructure so that each network segment has, for example, a TCP/IP subnet or IPX network. Traffic flow on routed networks is inherently different from traffic flow on switched networks. Routed networks have more flexible traffic flow because they can use the hierarchy to determine optimal paths depending on dynamic factors such as network congestion.

Security

Information is available to routers and switches that can be used to create more secure networks. LAN switches may use custom filters to provide access control based on destination address, source address, protocol type, packet length, and offset bits within the frame. Routers can filter on logical network addresses and provide control based on options available in Layer 3 protocols. For example, routers can permit or deny traffic based on specific TCP/IP socket information for a range of network addresses.

Media Dependence

Two factors need to be considered with regard to mixed-media internetworks. First, the maximum transfer unit (MTU) differs for various network media. Table 12–2 lists the maximum frame size for various network media.

Table 12-2 MTUs for Various Network Media

| Media | Minimum Valid Frame | Maximum Valid Size |
|------------------|---------------------|--------------------------------|
| Ethernet | 46 bytes | 1500 bytes |
| Token Ring | 32 bytes | 16 KB theoretical, 4 KB normal |
| Fast Ethernet | 46 bytes | 1500 bytes |
| FDDI | 32 bytes | 4468 bytes |
| ATM LANE | 64 bytes | 1518 bytes |
| ATM Classical IP | 64 bytes | 9180 bytes |
| Serial HDLC | 14 bytes | No limit, 4.5 KB normal |

When LANs of dissimilar media are switched, hosts must use the MTU that is the lowest common denominator of all the switched LANs that make up the internetwork. This requirement limits throughput and can seriously compromise performance over a relatively fast link, such as FDDI or ATM. Most Layer 3 protocols can fragment and reassemble packets that are too large for a particular subnetwork, so routed networks can accommodate different MTUs, which maximizes throughput.

Second, because they operate at Layer 2, switches must use a translation function to switch between dissimilar media. The translation function can result in serious problems such as non-canonical versus canonical Token Ring-to-Ethernet MAC format conversion. One issue with moving data from

a Token Ring to an Ethernet network is Layer 2 addressing. Token Ring devices read the Layer 2 MAC address as most significant bit starting from left to right. Ethernet devices read the Layer 2 MAC address as most significant bit starting from right to left.

By working at Layer 3, routers are essentially independent of the properties of any physical media and can use a simple address resolution algorithm (such as *Novell-node-address = MAC-address*) or a protocol, such as the Address Resolution Protocol (ARP), to resolve differences between Layer 2 and Layer 3 addresses.

Benefits of LAN Switches (Layer 2 Services)

An individual Layer 2 switch might offer some or all of the following benefits:

- *Bandwidth*—LAN switches provide excellent performance for individual users by allocating dedicated bandwidth to each switch port (for example, each network segment). This technique is known as *microsegmenting*.
- *VLANs*—LAN switches can group individual ports into logical switched workgroups called VLANs, thereby restricting the broadcast domain to designated VLAN member ports. VLANs are also known as switched domains and autonomous switching domains. Communication between VLANs requires a router.
- *Automated packet recognition and translation*—Cisco's unique Automatic Packet Recognition and Translation (APaRT) technology recognizes and converts a variety of Ethernet protocol formats into industry-standard CDDI/FDDI formats. With no changes needed in either client or server end stations the Catalyst solution can provide an easy migration to 100-Mbps server access while preserving the user's investment in existing shared 10Base-T LANs.

Benefits of Routers (Layer 3 Services)

Because routers use Layer 3 addresses, which typically have structure, routers can use techniques (such as address summarization) to build networks that maintain performance and responsiveness as they grow in size. By imposing structure (usually hierarchical) on a network, routers can effectively use redundant paths and determine optimal routes even in a dynamically changing network. This section describes the router functions that are vital in switched LAN designs:

- Broadcast and Multicast Control
- Broadcast Segmentation
- Media Transition

Broadcast and Multicast Control

Routers control broadcasts and multicasts in the following ways:

- *By caching the addresses of remote hosts.* When a host sends a broadcast packet to obtain the address of a remote host that the router already knows about, the router responds on behalf of the remote host and drops the broadcast packet (sparing hosts from having to respond to it).
- *By caching advertised network services.* When a router learns of a new network service, it caches the necessary information and does not forward broadcasts related to it. When a client of that network service sends a broadcast to locate that service, the router responds on behalf of that service and drops the broadcast packet (sparing hosts from having to respond to it). For example, Novell clients use broadcasts to find local services. In a network without a router, every server responds to every client broadcast by multicasting its list of services. Routers manage Novell

broadcasts by collecting services not local to the switch and sending out periodic updates that describe the services offered on the entire network. Each router sends out one frame for every seven services on the network.

- *By providing special protocols, such as the IGMP and Protocol Independent Multicast (PIM).* These new protocols allow a multicasting application to “negotiate” with routers, switches, and clients to determine the devices that belong to a multicast group. This negotiation helps limit the scope and impact of the multicast stream on the network as a whole. For information about IP multicasting, see Chapter 13, “Designing Internetworks for Multimedia.”

Successful network designs contain a mix of appropriately scaled switching and routing. Given the effects of broadcast radiation on CPU performance, well-managed switched LAN designs must include routers for broadcast and multicast management.

Broadcast Segmentation

In addition to preventing broadcasts from radiating throughout the network, routers are also responsible for generating services to each LAN segment. The following are examples of services that the router provides to the network for a variety of protocols:

- *IP*—Proxy ARP and Internet Control Message Protocol (ICMP)
- *IPX*—SAP table updates
- *AppleTalk*—ZIP table updates
- *Network management*—SNMP queries

In a flat virtual network, a single router would be bombarded by myriad requests needing replies, severely taxing its processor. Therefore, the network designer needs to consider the number of routers that can provide reliable services to a given subset of VLANs. Some type of hierarchical design needs to be considered.

Media Transition

In the past, routers have been used to connect networks of different media types, taking care of the OSI Layer 3 address translations and fragmentation requirements. Routers continue to perform this function in switched LAN designs. Most switching is done within like media (such as Ethernet, Token Ring, and FDDI switches) with some capability of connecting to another media type. However, if a requirement for a switched campus network design is to provide high-speed connectivity between unlike media, routers play a significant part in the design.

Benefits of VLANs

In a flat, bridged network all broadcast packets generated by any node in the network are sent to and received by all other network nodes. The ambient level of broadcasts generated by the higher layer protocols in the network—known as *broadcast radiation*—will typically restrict the total number of nodes that the network can support. In extreme cases, the effects of broadcast radiation can be so severe that an end station spends all of its CPU power on processing broadcasts.

VLANs have been designed to address the following problems inherent in a flat, bridged network:

- Scalability issues of a flat network topology
- Simplification of network management by facilitating network reconfigurations

VLANs solve some of the scalability problems of large flat networks by breaking a single bridged domain into several smaller bridged domains, each of which is a virtual LAN. Note that each virtual LAN is itself constrained by the scalability issues described in Appendix E, “Broadcasts in Switched LAN Internetworks.” It is insufficient to solve the broadcast problems inherent to a flat switched network by superimposing VLANs and reducing broadcast domains. VLANs without routers do not scale to large campus environments. Routing is instrumental in the building of scalable VLANs and is the only way to impose hierarchy on the switched VLAN internetwork.

VLANs offer the following features:

- *Broadcast control*—Just as switches isolate collision domains for attached hosts and only forward appropriate traffic out a particular port, VLANs refine this concept further and provide complete isolation between VLANs. A VLAN is a bridging domain, and all broadcast and multicast traffic is contained within it.
- *Security*—VLANs provide security in two ways:
 - High-security users can be grouped into a VLAN, possibly on the same physical segment, and no users outside of that VLAN can communicate with them.
 - Because VLANs are logical groups that behave like physically separate entities, inter-VLAN communication is achieved through a router. When inter-VLAN communication occurs through a router, all the security and filtering functionality that routers traditionally provide can be used because routers are able to look at OSI Layer 3 information. In the case of nonroutable protocols, there can be no inter-VLAN communication. All communication must occur within the same VLAN.
- *Performance*—The logical grouping of users allows, for example, an engineer making intensive use of a networked CAD/CAM station or testing a multicast application to be assigned to a VLAN that contains just that engineer and the servers he or she needs. The engineer’s work does not affect the rest of the engineering group, which results in improved performance for the engineer (by being on a dedicated LAN) and improved performance for the rest of the engineering group (whose communications are not slowed down by the engineer’s use of the network).
- *Network management*—The logical grouping of users, divorced from their physical or geographic locations, allows easier network management. It is no longer necessary to pull cables to move a user from one network to another. Adds, moves, and changes are achieved by configuring a port into the appropriate VLAN. Expensive, time-consuming recabling to extend connectivity in a switched LAN environment is no longer necessary because network management can be used to logically assign a user from one VLAN to another.

VLAN Implementation

This section describes the different methods of creating the logical groupings (or broadcast domains) that make up various types of VLANs. There are three ways of defining a VLAN:

- *By port*—Each port on the switch can support only one VLAN. With port-based VLANs, no Layer 3 address recognition takes place, so Internet Protocol (IP), Novell, and AppleTalk networks must share the same VLAN definition. All traffic within the VLAN is switched, and traffic between VLANs is routed (by an external router or by a router within the switch). This type of VLAN is also known as a *segment-based VLAN*.
- *By protocol*—VLANs based on network addresses (that is, OSI Layer 3 addresses) can differentiate between different protocols, allowing the definition of VLANs to be made on a per-protocol basis. With network address-based VLANs, it will be possible to have a different virtual topology for each protocol, with each topology having its own set of rules, firewalls, and

so forth. Routing between VLANs comes automatically, without the need for an external router or card. Network address-based VLANs will mean that a single port on a switch can support more than one VLAN. This type of VLAN is also known as a *virtual subnet VLAN*.

- *By a user-defined value*—This type of VLAN is typically the most flexible, allowing VLANs to be defined based on the value of any field in a packet. For example, VLANs could be defined on a protocol basis or could be dependent on a particular IPX or NetBIOS service. The simplest form of this type of VLAN is to group users according to their MAC addresses.

Cisco's initial method of implementing VLANs on routers and Catalyst switches is by port. To efficiently operate and manage protocols, such as IP, IPX, and AppleTalk, all nodes in a VLAN should be in the same subnet or network.

Cisco uses three technologies to implement VLANs:

- IEEE 802.10
- Inter-Switch Link (ISL)
- LAN Emulation

The three technologies are similar in that they are based on OSI Layer 2 bridge multiplexing mechanisms.

Note With respect to this chapter and the discussions in it, VLANs are differentiated by assigning each VLAN a “color” (or VLAN ID). For example, Engineering might be the “blue” VLAN, and Manufacturing might be the “green” VLAN.

IEEE 802.10


IEEE 802.10 defines a method for secure bridging of data across a shared metropolitan area network (MAN) backbone. Cisco has initially implemented the relevant portions of the standard to allow the “coloring” of bridged traffic across high-speed backbones (FDDI, Ethernet, Fast Ethernet, Token Ring, and serial links). There are two strategies using IEEE 802.10 to implement VLANs, depending on how traffic is handled through the backbone:

- Switched Backbone
- Routed Backbone

Switched Backbone

In the switched backbone topology shown in Figure 12-7, you want to ensure that intra-VLAN traffic goes only between Segment A and Segment D (both in VLAN 10) and Segment B and Segment C (both in VLAN 20).

Figure 12-7 IEEE 802.10 switched backbone implementation.




In Figure 12-7, all Ethernet ports on Switches X, Y, and Z are in a VLAN and are to be VLAN interfaces. All FDDI interfaces in Switches X, Y, and Z are called *VLAN trunk interfaces*. To ensure that traffic from Segment A destined for Segment D on Switch Z is forwarded onto Ethernet 3 and not onto Ethernet 2, it is colored when it leaves Switch X. Switch Z recognizes the color and knows that it must forward these frames onto Ethernet 3 and not onto Ethernet 2.

The coloring of traffic across the FDDI backbone is achieved by inserting a 16-byte header between the source MAC address and the Link Service Access Point (LSAP) of frames leaving a switch. This header contains a 4-byte VLAN ID or “color.” The receiving switch removes the header and forwards the frame to interfaces that match that VLAN color.

Routed Backbone

In the routed backbone topology shown in Figure 12–8, the goal is the same as for the switched topology—that is, to ensure that intra-VLAN traffic goes only between Segment A and Segment D (both in VLAN 10) and Segment B and Segment C (both in VLAN 20).

Figure 12-8 IEEE 802.10 routed backbone implementation.




It is important that a single VLAN use only one subnet. In Figure 12-8, VLAN 10 (subnet 10) is “split” and therefore must be “glued” together by maintaining a bridged path for it through the network. For Switch X and nodes in VLAN 20 (subnet 20), traffic is switched locally if appropriate. If traffic is destined for a node in VLAN 30 (subnet 30) from a node in VLAN 20, Router Y routes it through the backbone to Router Z. If traffic from Segment D on VLAN 10 is destined for a node in VLAN 20, Router Y routes it back out the FDDI interface.

Fast EtherChannel

Fast Ether Channel is a trunking technology based on grouping together multiple full duplex 802.3 Fast Ethernets to provide fault-tolerant high-speed links between switches, routers, and servers. Fast EtherChannels can be composed of two to four industry-standard Fast Ethernet links to provide load sharing of traffic with up to 800 Mbps of usable bandwidth. Fast EtherChannels can interconnect LAN switches, routers, servers, and clients. Since its load balancing is integrated with the Catalyst 5000 families LAN switch architectures, there is no performance degradation for adding links to a channel—high throughput and low latencies can be maintained while gaining more total available bandwidth. Fast EtherChannel provides link resiliency within a channel—if links should fail the traffic is immediately directed to the remaining links. Finally, Fast EtherChannel is not

dependent on any type of media—it can be used with Fast Ethernet running on existing Unshielded Twisted Pair (UTP) wiring, or single mode and multimode fiber. Figure 12–9 shows a collapsed backbone topology design using the Fast EtherChannel modules to provide links of 400 Mb between switches in the wiring closets and the data center.

Figure 12-9 Collapsed backbone topology design using the Fast EtherChannel modules.



IEEE 802.10 Design Issues

- Routers fast switch IEEE 802.10, which means that the fast-switching throughput of the platform must be considered.
- VLANs must be consistent with the routed model. (That is, subnets cannot be split.)
 - If subnets must be split, they must be “glued” together by a bridged path.
 - Normal routed behavior needs to be maintained for end nodes to correctly achieve routing between VLANs.
 - Networks need to be designed carefully when integrating VLANs; the simplest choice is to avoid splitting VLANs across a routed backbone.

The difference between these two strategies is subtle. Table 12–3 compares the advantages and disadvantages of the two strategies.

Table 12-3 Advantages and Disadvantages of Switched and Routed Backbones

| Switched Backbone | | Routed Backbone | |
|---|--|---|--|
| Advantages | Disadvantages | Advantages | Disadvantages |
| Propagates color information across entire network. | Backbone is running bridging. | No bridging in backbone. | Color information is not propagated across backbone and must be configured manually. |
| Allows greater scalability by extending bridge domains. | Broadcast traffic increases drastically on the backbone. | Easy to integrate into existing internetwork. | If subnets are split, a bridged path has to be set up between switches. |
| | | Can run native protocols in the backbone. | |


A VLAN interface can have only one VLAN ID, and VLAN trunk interfaces support multiple VLANs across them.

Inter-Switch Link

ISL is a Cisco-proprietary protocol for interconnecting multiple switches and maintaining VLAN information as traffic goes between switches. This technology is similar to IEEE 802.10 in that it is a method of multiplexing bridge groups over a high-speed backbone. It is defined only on Fast Ethernet. The discussion of routing and switching in the backbone in the section “IEEE 802.10,” earlier in this chapter, also applies to ISL.

With ISL, an Ethernet frame is encapsulated with a header that maintains VLAN IDs between switches. A 30-byte header is prepended to the Ethernet frame, and it contains a two-byte VLAN ID. In Figure 12–10, Switch Y switches VLAN 20 traffic between segments A and B if appropriate. Otherwise, it encapsulates traffic with an ISL header that identifies it as traffic for VLAN 20 and sends it through the interim switch to Router X. Router X routes the packet to the appropriate interface, which could be through a routed network beyond Router X (as in this case) out the Fast Ethernet interface to Switch Z. Switch Z receives the packet, examines the ISL header noting that this packet is destined for VLAN 20, and switches it to all ports in VLAN 20 (if the packet is a broadcast or multicast) or the appropriate port (if the packet is a unicast).

Figure 12-10 Inter-switch link design.



Router X routes the packet to the appropriate interface, which could be through a routed network beyond Router X (as in this case) out the Fast Ethernet interface to Switch Z. Switch Z receives the packet, examines the ISL header noting that this packet is destined for VLAN 20, and switches it to all ports in VLAN 20 (if the packet is a broadcast or multicast) or the appropriate port (if the packet is a unicast).

Note Routers fast switch ISL, which means that the fast-switching throughput of the platform must be considered.

LAN Emulation

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


LANE uses MAC encapsulation (OSI Layer 2) because this approach supports the largest number of existing OSI Layer 3 protocols. The end result is that all devices attached to an emulated LAN appear to be on one bridged segment. In this way, AppleTalk, IPX, and other protocols should have similar performance characteristics as in a traditional bridged environment. In ATM LANE environments, the ATM switch handles traffic that belongs to the same emulated LAN (ELAN), and routers handle inter-ELAN traffic. For more information about LANE, see Chapter 8, “Designing ATM Internetworks.”

Virtual Multihomed Servers

In traditional networks, there are usually several well-known servers, such as e-mail and corporate servers, that almost everyone in an enterprise needs to access. If these servers are located in only one VLAN, the benefits of VLANs will be lost because all of the different workgroups will be forced to route to access this common information source.

This problem can be solved with LANE and virtual multihomed servers, as shown in Figure 12–11. Network interface cards (NICs) allow workstations and servers to join up to eight different VLANs. This means that the server will appear in eight different ELANs and that to other members of each ELAN, the server appears to be like any other member. This capability greatly increases the performance of the network as a whole because common information is available directly through the optimal Data Direct VCC and does not need to be routed. This also means that the server must process all broadcast traffic in each VLAN that it belongs to, which can decrease performance.

Figure 12-11 Multihomed servers in an ATM network.



To multihome servers in non-ATM environments, there are two possible choices:

- Use servers with multiple NICs (different connections into each VLAN).
- Use servers with NICs that support the VLAN trunking technology (IEEE 802.10 or ISL) used in the backbone.

Virtual Trunk Protocol

The Catalyst 5000 switch implements Cisco's Virtual Trunk Protocol (VTP). VTP is the industry's first protocol implementation specifically designed for large VLAN deployments. VTP enhances VLAN deployment by providing the following:

- Integration of ISL, 802.10, and ATM LAN-based VLANs
- Auto-intelligence within the switches for configuring VLANs
- Configuration consistency across the network
- An automapping scheme for going across mixed-media backbones

- Accurate tracking and monitoring of VLANs
- Dynamic reporting of added VLANs across the network
- Plug-and-play setup and configuration when adding new VLANs

General Network Design Principles

Good network design is based on many concepts that are summarized by the following key principles:

- *Examine single points of failure carefully*—There should be redundancy in the network so that a single failure does not isolate any portion of the network. There are two aspects of redundancy that need to be considered: backup and load balancing. In the event of a failure in the network, there should be an alternative or backup path. Load balancing occurs when two or more paths to a destination exist and can be utilized depending on the network load. The level of redundancy required in a particular network varies from network to network.
- *Characterize application and protocol traffic*—For example, the flow of application data will profile client-server interaction and is crucial for efficient resource allocation, such as the number of clients using a particular server or the number of client workstations on a segment.
- *Analyze bandwidth availability*—For example, there should not be an order of magnitude difference between the different layers of the hierarchical model. It is important to remember that the hierarchical model refers to conceptual layers that provide functionality. The actual demarcation between layers does not have to be a physical link—it can be the backplane of a particular device.
- *Build networks using a hierarchical or modular model*—The hierarchy allows autonomous segments to be internetworked together.

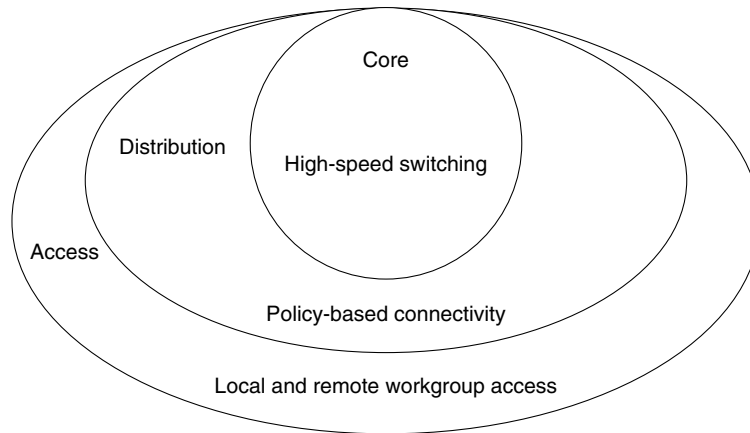
Figure 12–12 shows a high-level view of the various aspects of a hierarchical network design. A hierarchical network design presents three layers—core, distribution, and access—with each layer providing different functionality.

The core layer is a high-speed switching backbone and should be designed to switch packets as fast as possible. This layer of the network should not perform any packet manipulation access lists and filtering that would slow down the switching of packets.

The distribution layer of the network is the demarcation point between the access and core layers and helps to define and differentiate the core. The purpose of this layer is to provide boundary definition and is the place at which packet manipulation can take place. In the campus environment, the distribution layer can include several functions, such as the following:

- Address or area aggregation
- Departmental or workgroup access
- Broadcast/multicast domain definition
- VLAN routing
- Any media transitions that need to occur
- Security

Figure 12-12 Hierarchical network design model.



In the non-campus environment, the distribution layer can be a redistribution point between routing domains or the demarcation between static and dynamic routing protocols. It can also be the point at which remote sites access the corporate network. The distribution layer can be summarized as the layer that provides policy-based connectivity.

The access layer is the point at which local end users are allowed into the network. This layer may also use access lists or filters to further optimize the needs of a particular set of users. In the campus environment, access-layer functions can include the following:

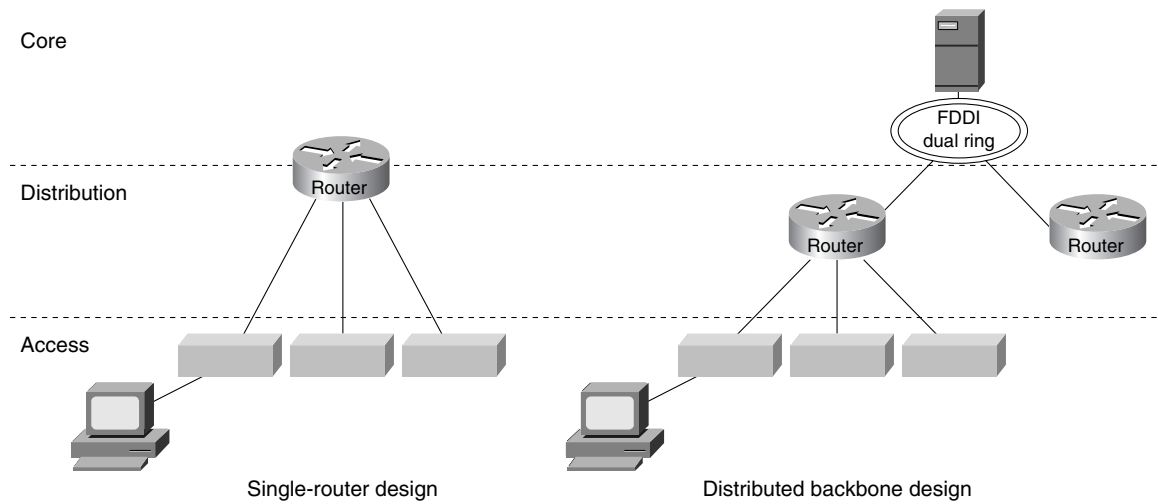
- Shared bandwidth
- Switched bandwidth
- MAC layer filtering
- Microsegmentation

In the non-campus environment, the access layer can give remote sites access to the corporate network via some wide-area technology, such as Frame Relay, ISDN, or leased lines.

It is sometimes mistakenly thought that the three layers (core, distribution, and access) must exist in clear and distinct physical entities, but this does not have to be the case. The layers are defined to aid successful network design and to represent functionality that must exist in a network. The instantiation of each layer can be in distinct routers or switches, can be represented by a physical media, can be combined in a single device, or can be omitted altogether. The way the layers are implemented depends on the needs of the network being designed. Note, however, that for a network to function optimally, hierarchy must be maintained.

With respect to the hierarchical model, traditional campus LANs have followed one of two designs—single router and distributed backbone—as shown in Figure 12–13.

Figure 12-13 Traditional campus design.



In the single-router design, the core and distribution layers are present in a single entity—the router. Core functionality is represented by the backplane of the router and distribution is represented by the router. Access for end users is through individual- or chassis-based hubs. This design suffers from scalability constraints because the router can be only be in one physical location, so all segments end at the same location—the router. The single router is responsible for all distribution functionality, which can cause CPU overload.

The distributed backbone design uses a high-speed backbone media, typically FDDI, to spread routing functionality among several routers. This also allows the backbone to traverse floors, a building, or a campus.

Switched LAN Network Design Principles

When designing switched LAN campus networks, the following factors must be considered:

- *Broadcast radiation*—Broadcast radiation can become fatal—that is, 100 percent of host CPU cycles can be consumed by processing broadcast and multicast packets. Because of delays inherent in carrier sense multiple access collision detect (CSMA/CD) technologies, such as Ethernet, any more than a small amount of broadcast traffic will adversely affect the operation of devices attached to a switch. Although VLANs reduce the effect of broadcast radiation on all LANs, there is still a scaling issue as to how many hosts should reside on a given VLAN. A router allows for larger network designs because a VLAN can be subsegmented depending on traffic patterns. However, in a nonoptimal network design, a single router can be burdened with large amounts of traffic.
- *Well-behaved VLANs*—A well-behaved VLAN is a VLAN in which 80 percent or more of the traffic is local to that VLAN. In an example in which the Marketing, MIS, and Engineering departments each have an individual VLAN segment, the 80 percent rule is violated when a user in the Marketing VLAN reads mail from the MIS VLAN, mounts servers from the Engineering VLAN, and sends e-mail to members of the Engineering VLAN.
- *Available bandwidth to access routing functionality*—Inter-VLAN traffic must be routed, so the network design must allocate enough bandwidth to move inter-VLAN traffic from the source, through the device that provides routing functionality, and to the destination.
- *Appropriate placement of administrative boundaries*—Switching has the effect of flattening networks, and the deployment of switching outside of your administrative boundary can adversely affect the network within your administrative boundary.


Campus network designs are evolving rapidly with the deployment of switching at all levels of the network—from the desktop to the backbone. Three topologies have emerged as generic network designs:

- Scaled Switching
- Large Switching/Minimal Routing
- Distributed Routing/Switching

Scaled Switching

The scaled switching design shown in Figure 12–14 deploys switching at all levels of the network without the use of routers. In this design, each layer consists of switches, with switches in the access layer providing 10-Mbps Ethernet or 16-Mbps Token Ring to end users.

Figure 12-14 Scaled switching design.

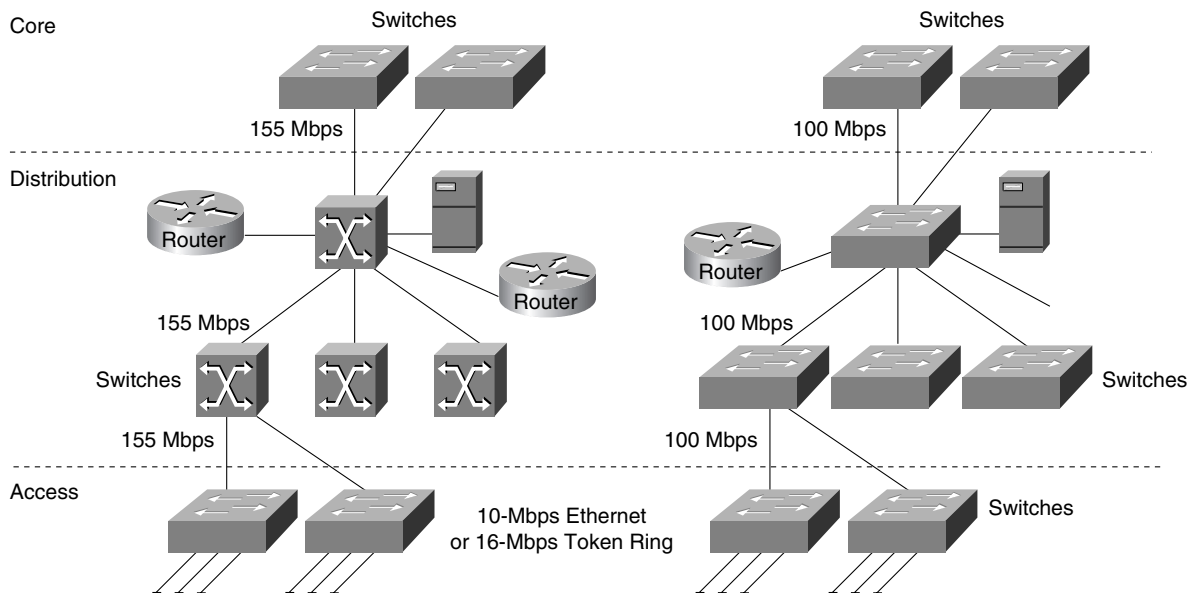


Scaled switching is a low-cost and easy-to-install solution for a small campus network. It does not require knowledge of address structure, is easy to manage, and allows all users to communicate with one another. However, this network comprises a single broadcast domain. If a scaled switched network needs to grow beyond the broadcast domain, it can use VLANs to create multiple broadcast domains. Note that when VLANs are used, end users in one VLAN cannot communicate with end users in another VLAN unless routers are deployed.

Large Switched/Minimal Routing

The large switched/minimal routing design deploys switching at the access layer of the network, either ATM switching or LAN switching at the distribution layer of the network, and ATM/LAN switching at the core. Figure 12–15 shows an example of this network design.

Figure 12-15 Large switched/minimal routing design.



In the case of ATM in the distribution layer, the following key issues are relevant:

- LANE support on routers and switches.
- Support for UNI 3.X signaling (including point-to-multipoint).
- If redundancy is provided by a virtual PVC or SVC mesh, the mesh is a single point of failure.

In the case of LAN switching in the distribution layer, the following key issues are relevant:

- Support for VLAN trunking technology in each device.
- The switches in the distribution layer must run the Spanning-Tree Protocol to prevent loops, which means that some connections will be blocked and load balancing cannot occur.


To scale the large switched/minimal routing design, a logical hierarchy must be imposed. The logical hierarchy consists of VLANs and routers that enable inter-VLAN communication. In this topology, routing is used only in the distribution layer, and the access layer depends on bandwidth through the distribution layer to gain access to high-speed switching functionality in the core layer.

The large switched/minimal routing design scales well when VLANs are designed so that the majority of resources are available in the VLAN. Therefore, if this topology can be designed so that 80 percent of traffic is intra-VLAN and only 20 percent of traffic is inter-VLAN, the bandwidth needed for inter-VLAN routing is not a concern. However, if inter-VLAN traffic is greater than 20 percent, access to routing in the core becomes a scalability issue. For optimal network operation, scalable routing content is needed at the distribution layer of the network.

Distributed Routing/Switching

The distributed routing/switching design deploys switching in the access layer, routing in the distribution layer, and some form of high-speed switching in the core layer, as shown in Figure 12-16.

Figure 12-16 Distributed routing/switching design.



The distributed routing/switching design follows the classic hierarchical network model both physically and logically. Because it provides high bandwidth for access to routing functionality, this design scales very well. This design is optimized for networks that do not have the 80/20 pattern rule. If servers are centralized, most traffic is inter-VLAN; therefore, high routing content is needed.

Summary

Campus LAN designs use switches to replace traditional hubs and use an appropriate mix of routers to minimize broadcast radiation. With the appropriate pieces of software and hardware in place, and adhering to good network design, it is possible to build topologies, such as the examples described in the section “Switched LAN Network Designs” earlier in this chapter.

