CHAPTER 1

Overview of a Campus Network

Today's campus networks are comprised of a combination of bridges and routers connected via either Ethernet or Token Ring. This chapter presents a new approach to campus networking called *multilayer switching*. Multilayer switching combines Layer 2 switching functions with Layer 3 routing. This chapter also includes a discussion of the correct placement of Layer 2, Layer 3, and multilayer devices in the campus network. Finally, this chapter shows how you can scale your campus network to meet the ever-increasing demands of your business.

This chapter describes a set of building blocks, which present a logical design for the network irrespective of any product that can be implemented. The success of any campus intranet is based on the placement of network services that, when applied correctly, will guarantee continued scalability.

This chapter covers the following topics:

- Campus network overview
- The emerging campus model
- Switching technologies
- The hierarchical model
- The building block approach
- Campus network availability example

Upon completion of this chapter, given a list of switching functions, you will be able to identify the correct Open System Interconnection (OSI) reference model layer associated with those functions. Given a list of characteristics, you will be able to identify the correct hierarchical model layer. And finally, given a set of user requirements, you will be able to identify the correct Cisco product solution.

Campus Network Overview

This section contains an overview of the traditional campus networks and describes some of the major issues and solutions that have changed the way networks operate. This section also discusses how network traffic patterns are changing. This section covers the following topics:

- The structure and characteristics of current campus networks
- Traditional network problems and the resulting solutions
- Existing and emerging traffic patterns

This section also deals with the demands on today's organizations that have brought about changes in the way campus networks are designed, as well as the components of a campus network and the technologies driving these changes.

Traditional Campus Networks

A *campus* is a building or group of buildings connected into one enterprise network that consists of many LANs. A campus is further defined as a company or a portion of a company contained in a fixed geographic area.

The distinct characteristic of a campus environment is that the company that owns the campus network usually owns the physical wires deployed in the campus. The campus network topology is primarily a LAN technology connecting all the end systems within the building or buildings. Campus networks generally use LAN technologies, such as Ethernet, Token Ring, and Fiber Distributed Data Interface (FDDI). Figure 1-1 shows a sample campus network.

Network designers generally deploy a campus design that is optimized for the fastest functional architecture that runs on the existing physical wire. This self-study book discusses the requirements of emerging applications and how higher-speed technologies, such as Fast Ethernet, Fast EtherChannel, and Gigabit Ethernet, along with multilayer switching (MLS), provide wire-speed data transfer to the desktop.

Regardless of the underlying technology, the main challenge facing network designers and administrators today is to have their campus LAN run effectively and efficiently. In order to achieve this goal, you must understand, implement, and manage the traffic flow throughout your network.

Originally, campus networks consisted of a single LAN to which new users were added. Because of distance limitations, campus networks usually were confined to a building or several buildings in close proximity to each other. The LAN was a physical network that connected the devices. In the case of Ethernet, all the devices shared the available half-duplex 10 Mbps. Because of the carrier sense multiple access collision detect (CSMA/CD) scheme used by the Ethernet, the whole LAN was considered a collision domain.

Figure 1-1 *A Traditional Campus Network*

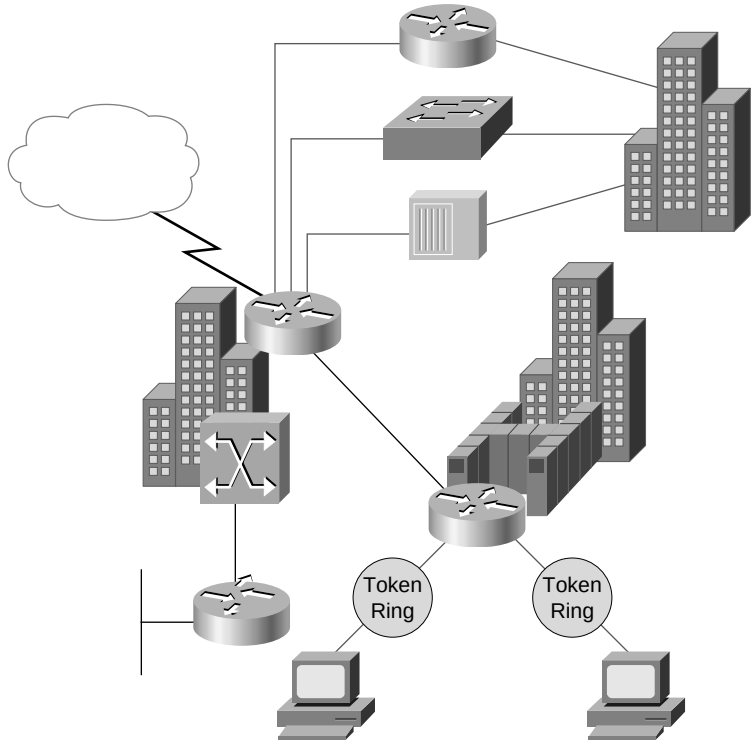

Few design considerations were needed to provide user access to the network backbone. Because of the limitations of Ethernet, physically adjacent users were connected to a single access device to minimize the number of taps into the backbone. Although hubs met this requirement and became standard devices for multiple network access, increased user demand quickly impacted the network's performance.

Traditional Campus Issues

The major problems with traditional networks are availability and performance. These two problems are impacted by the amount of bandwidth in the network.

In a single collision domain, frames are visible to all devices on the LAN and are free to collide. Multiport bridges segment the LAN into discrete collision domains. and forward Layer 2 data frames to only the segment that contains the destination address. Because bridge ports separate the LAN into distinct physical segments, bridges also help resolve Ethernet's distance limitations. Bridges must, however, forward broadcasts, multicasts, and unknown unicasts to all ports. Figure 1-2 shows that bridges terminate collision domains.

Figure 1-2 *A Bridge on a Traditional Campus Network*

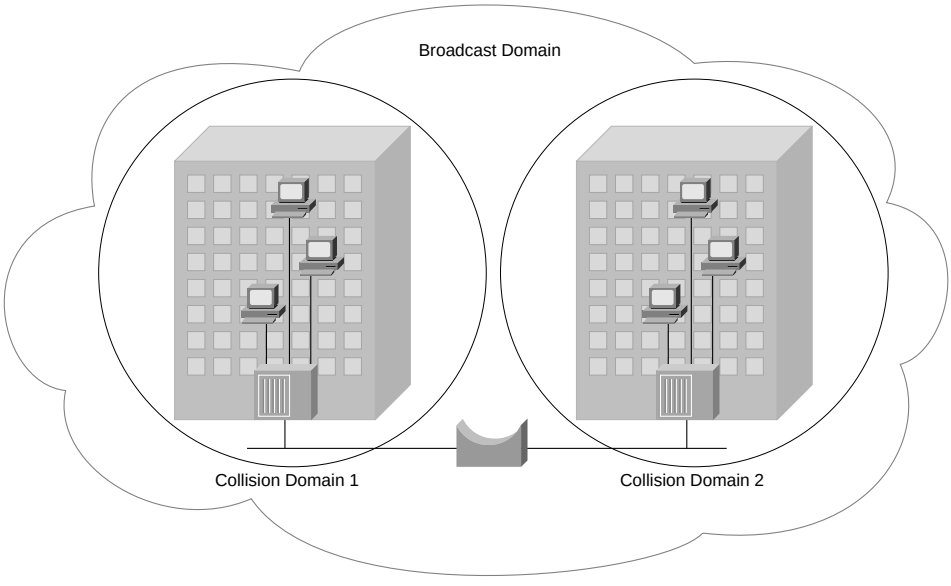

NOTE A collision domain consists of all the devices that can see or be involved in a collision. A Layer 2 device such as a bridge or switch borders a collision domain. A collision domain is different from a broadcast domain. A broadcast domain consists of all the devices that can see the broadcast. A Layer 3 device such as a router borders a broadcast domain. By default, traditional bridge ports create separate collision domains but participate in the same broadcast domain.

Because bridges read only the Media Access Control (MAC) address in the frame, however, frames containing the broadcast MAC address still flood the entire network. Also, a single network device could malfunction and flood the network with indiscriminate jabber, virtually disabling the network. Because routers operate at the network layer, these devices can make intelligent decisions regarding the flow and type of information to and from a network subnet.

Examples of broadcasts that ask questions are IP Address Resolution Protocol (ARP) requests, NetBIOS name requests, and Internetwork Packet Exchange (IPX) Get Nearest Server requests. These types of broadcasts typically flood the entire subnet and have the target device respond directly to the broadcast. Figure 1-3 illustrates how multicasts, broadcasts, and even unknown unicasts become global events in the bridged network because each bridge processes the request for the MAC address for Server A.

Figure 1-3 *The Request for Server A's MAC Address Is Processed on Every Bridge in a Bridged Network*

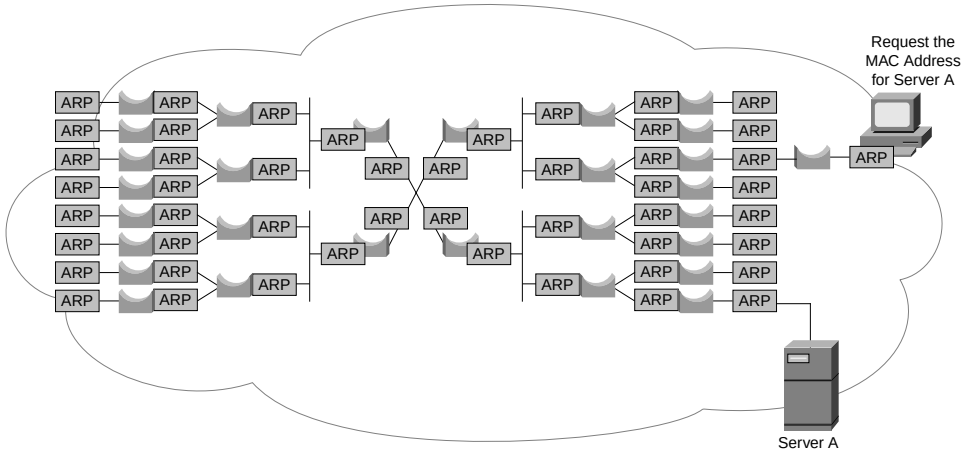

Examples of broadcasts that advertise are IPX Service Advertising Protocol (SAP) packets and routing protocols such as the Routing Information Protocol (RIP) and Interior Gateway Routing Protocol (IGRP).

Finally, multicast traffic can also consume bandwidth. Multicast traffic transmits to a specific group; however, depending on the number of users in a group or the type of application data contained in the multicast packet, this type of transmission can consume most, if not all, of the network resources. Examples of multicast implementations are the Cisco IPTV application using multicast packets to distribute multimedia data, and Novell 5 on IP using multicast packets to locate services.

As networks grow, so does broadcast traffic. Excessive broadcasts reduce the bandwidth available to the end users. In worst-case scenarios, broadcast storms can effectively shut down the network, because the broadcasts monopolize all the available bandwidth.

Also, in a bridge only network, all network-attached workstations and servers are forced to decode all broadcast frames. This action generates additional CPU interrupts and degrades application performance.

A Solution to Broadcast Domain Issues: Localize Traffic

There are two options for addressing the broadcast containment issue for large switched LAN sites.

The first option is to use routers to create many subnets, logically segmenting the traffic. LAN broadcasts do not pass through routers. Although this approach will filter broadcasts, traditional routers process each packet and can create a bottleneck in the network. For example, a Layer 2 bridge or switch can process millions of packets per second, while a

traditional router will process in the hundreds of thousands of packets per second. This difference in packet processing speed can cause a bottleneck in the network if a large amount of traffic has to cross the Layer 3 device.

The second option is to implement virtual LANs (VLANs) within the switched network. For the purpose of this book, VLANs are defined as broadcast domains. A VLAN is a group of end devices, on multiple physical LAN segments or switches, that communicates as if these end devices were located on a single shared-media LAN segment. Devices on the same VLAN need not be physically colocated in the same part of the building or campus; however, Cisco recommends a one-to-one correspondence between VLANs and IP subnets. Figure 1-4 shows the usage of VLANs to separate the network into individual broadcast domains.

Figure 1-4 *VLANs*

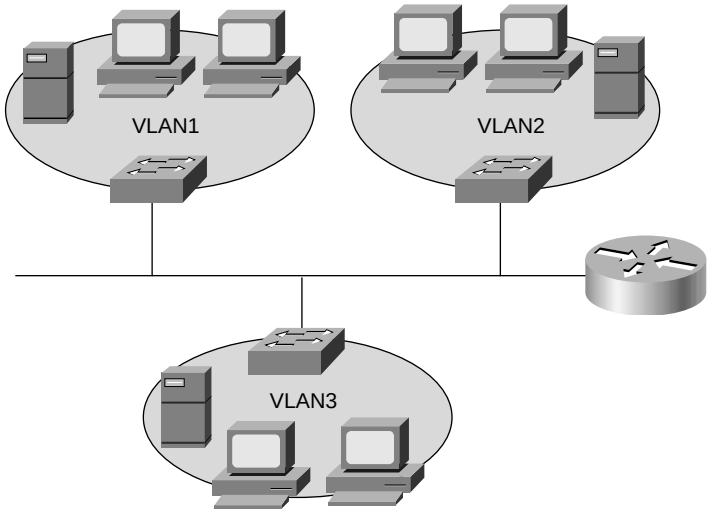

One of the primary benefits of VLANs is that LAN switches can be used to effectively contain broadcast traffic and separate traffic flows.

Because a VLAN is essentially a broadcast domain, the broadcast domain is now defined by a particular set of ports on switches. None of the switches in the set will bridge any frames between two VLANs. *It is important to note that routers are required to move traffic between broadcast domains.*

Current Campus Networks

Most campus networks now consist of two components: LAN switches and routers. By creating smaller Layer 2 broadcast domains and linking them using Layer 3 functionality,

network administrators can filter broadcast traffic, interconnect multiprotocol workgroups, and offer a level of secure traffic.

Traffic in the Network

Devices and associated software applications running on the network all generate data traffic. Your network probably has at least the typical applications, such as word processing, file transfer, and electronic mail. These applications do not require much bandwidth, and their traffic patterns are intuitive.

However, emerging campus LANs have and need much more than these basic applications. Multifaceted applications, such as desktop publishing, videoconferencing, and WebTV multicast programs, are all gaining popularity. The characteristics of these applications are not always as easy to predict.

NOTE It is recommended that you maintain a snapshot of the traffic on your network using a device such as a protocol analyzer or probe. Traffic patterns should be monitored on an ongoing basis, because they will change over time. A good understanding of your network's traffic patterns and applications is essential for planning and managing the network, as well as for future modifications such as quality of service (QoS).

The 80/20 Rule

Ideally, end users with common interests or work patterns are placed in the same logical network as the servers they access most often. With the definition of logical networks, most of the traffic within these workgroups is limited to the local segment. This simple task minimizes the load on the network backbone.

The 80/20 rule states that in a properly designed network environment, 80 percent of the traffic on a given network segment is local. Not more than 20 percent of the network traffic should move across a backbone. Backbone congestion indicates that traffic patterns are not meeting the 80/20 rule. In this case, rather than adding switches or upgrading hubs, network administrators can improve network performance by doing one of the following:

- Moving resources such as applications, software programs, and files from one server to another to contain traffic locally within a workgroup
- Moving users logically, if not physically, so that the workgroups more closely reflect the actual traffic patterns
- Adding servers so that users can access them locally without having to cross the backbone

The New 20/80 Rule

Traffic patterns are moving toward what is now referred to as the 20/80 model. In the 20/80 model, only 20 percent of traffic is local to the workgroup LAN, and 80 percent of the traffic is required to go off the local network.

Two factors contribute to these changing traffic patterns:

- With Web-based computing, such as Internet applications, a PC can be both a subscriber and a publisher of information. As a result, information can come from anywhere in the network, creating massive amounts of traffic that must travel across subnet boundaries. Users hop transparently between servers across the entire enterprise by using hyperlinks, without the need to know where the data is located.
- The second factor leading to the loss of locality is the move toward server consolidation. Enterprises are deploying centralized server farms because of the reduced cost of ownership, security, and ease of management. All traffic from the client subnets to these servers must travel across the campus backbone.

This change in traffic patterns means that 80 percent of the traffic now must cross a Layer 3 device. Because routing is a CPU-intensive process, the point where the Layer 3 processing takes place can lead to bottlenecks in the network. This change in traffic patterns requires the network Layer 3 performance to match the Layer 2 performance.

The new 20/80 rule makes it difficult for network administrators to manage VLANs. Network administrators do not want to spend their time tracking traffic patterns and redesigning the network. Because VLANs are created on the premise that most traffic is interworkgroup, end stations need to be in the same broadcast domain to take advantage of the switched infrastructure.

With the new 20/80 rule, end devices need access to multiple VLANs. However, these end devices still need to operate within their current VLANs.

With current and future traffic patterns moving away from the traditional 80/20 rule, more traffic must flow between subnets (VLANs). Also, access to specific devices needs to be controlled. To perform these functions, routing technology is required within the network.

All these factors together are redesigning the traditional networks into a new campus model.

The Emerging Campus Network

Customer requirements for the campus network are evolving. This section presents the features and technologies that can be used to respond to these requirements.

The key requirements placing pressure on the emerging campus designs are as follows:

- **Fast convergence**—This requirement stipulates that the network must adapt quickly to network topology changes. This requirement becomes critical as the campus network grows in geographic scope.
- **Deterministic paths**—This requirement demands the desirability of a given path to a destination for certain applications or user groups.
- **Deterministic failover**—This requirement specifies that a mechanism be in place to ensure that the network is operational at all times.
- **Scalable size and throughput**—This requirement orders that as the network grows and new applications are added, the infrastructure must handle the increased traffic demands.
- **Centralized applications**—This requirement dictates that centralized applications be available to support most or all users on the network.
- **The new 20/80 rule**—This requirement focuses on the shift in traditional traffic patterns.
- **Multiprotocol support**—This requirement specifies that campus networks must now support multiprotocol environments.
- **Multicasting**—This requirement demands that campus networks support IP multicast traffic in addition to IP unicast traffic.

Emerging Campus Structure

User demands and complex applications force network designers to focus on the traffic patterns in the network. No longer can networks be divided into subnetworks based only on the number of users. The emergence of servers that run global applications also has a direct effect on the load across the network. A higher traffic load across the entire network results in the need for more efficient routing and switching techniques.

In the new campus model, traffic patterns dictate the placement of the services required by the end user. Services can be separated into three separate categories:

- Local services
- Remote services
- Enterprise services

Local Services

A local service is when the entities that provide services reside on the same subnet, and therefore, the same virtual network as the user. Local services remain in specific areas of the network. Traffic to and from local services is confined between the server, switches, and end users. Local traffic does not enter the network backbone or move through a router.

To service the localized traffic, Layer 2 switches are moving to the edge of the network and into the wiring closets. These switches connect end-user devices and servers into common workgroups.

Remote Services

A remote service is an entity that might be geographically close to the end user but is not on the same subnet or VLAN as that user. Traffic to and from remote services might or might not cross the backbone. Because these services are remote to the requesting end user, however, requests for remote services will cross broadcast domain boundaries. Therefore, switches connect to Layer 3 devices to allow for cross-broadcast domain boundary traffic. The router also controls the type of traffic that crosses the network backbone.

Enterprise Services

Enterprise services are services common to all users. Examples of enterprise services are e-mail, Internet access, and videoconferencing. Because all users need to access enterprise services, these servers and services exist within a separate subnet placed close to the network's backbone. Because enterprise services exist outside the end user's broadcast domain, Layer 3 devices are required for access to these services. Enterprise services might or might not be grouped by Layer 2 switches.

Placing the enterprise servers close to the backbone ensures the same distance from each user; however, this also means that all traffic going to an enterprise server crosses the backbone.

Figure 1-5 shows the three services and how traffic patterns dictate the placement of these services.

Figure 1-5 *Sample Emerging Campus Network Structure*

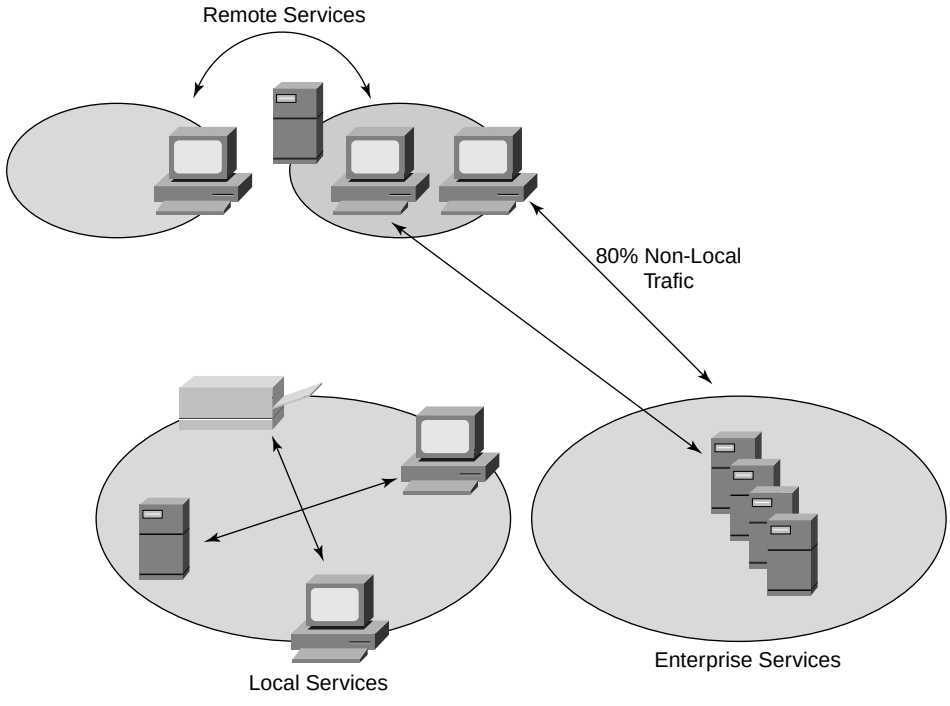

Switching Technologies

Due to the emerging 20/80 rule, network managers want to take advantage of the high throughput benefits of switching technology while still retaining Layer 3 functionality in the network. Therefore, a new model is required to support these requirements. This model employs a concept of providing switching techniques for Layer 2, 3, and 4 functions.

Basic Layer Terminology

Most communication environments use a model that separates the communications functions and applications processing into layers. Each layer serves a specific function. This self-study book focuses on Layers 2, 3, and 4 of this model. Figure 1-6 shows the basic layer terminology.

Figure 1-6 *Basic Layer Terminology*

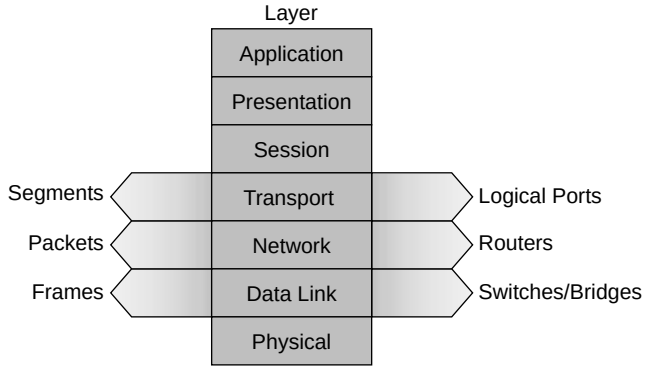

Each layer uses its own layer protocol to communicate with peer layers in another system. Each layer protocol exchanges information, called protocol data units (PDUs), between peer layers. A given layer can use a more specific name for its PDU. Table 1-1 gives examples of specific PDUs for Layers 2, 3, and 4 and the device types that process those PDUs.

Table 1-1 *PDU and Device Types Relating to the OSI Layers*

Model Layer	PDU Type	Device Type
Data Link (Layer 2)	Frames	Switches/bridges
Network (Layer 3)	Packets	Routers
Transport (Layer 4)	TCP segments	TCP ports

Each peer-layer protocol uses the services of the underlying layers. Thus, Transmission Control Protocol (TCP) segments are encapsulated in Layer 3 packets, and Layer 3 packets are encapsulated in Layer 2 frames. The layer-specific device processes only those PDU headers for which the device is responsible.

Layer 2 Switching

Layer 2 switching is hardware-based bridging. In a switch, frame forwarding is handled by specialized hardware called Application-Specific Integrated Circuits (ASICs). Because of ASIC technology, switches also provide scalability to gigabit speeds and low latency at costs significantly lower than Ethernet bridges.

NOTE An ASIC is an Application-Specific Integrated Circuit. This means that an application or process has been implemented in hardware or an integrated circuit chip. ASICs are used in everything from switches to wireless phones. For more information on specific ASICs in Cisco switches, refer to Appendix B, "Switching Architectures and Functional Descriptions."

Layer 2 switches give network managers the ability to increase bandwidth without adding unnecessary complexity to the network. Layer 2 data frames consist of both infrastructure content, such as MAC addresses, and end-user content. At Layer 2, no modification is required to the packet infrastructure content when going between like Layer 1 interfaces, such as from Ethernet to Fast Ethernet. However, changes to infrastructure content might occur when bridging between unlike media types such as FDDI and Ethernet or Token Ring and Ethernet.

Workgroup connectivity and network segmentation are the two primary uses for Layer 2 switches. The high performance of a Layer 2 switch can produce network designs that decrease the number of hosts per physical segment. Decreasing the hosts per segment leads to a flatter network design with more segments in the campus.

However, for all its advantages, Layer 2 switching has all the same characteristics and limitations as bridging, as shown in Figure 1-7. Broadcast domains built with Layer 2 switches still experience the same scaling and performance issues as the large bridged networks of the past. The broadcast and multicast radius increases with the number of hosts, and broadcasts still interrupt all the end stations. The Spanning-Tree Protocol limitations of slow convergence and blocked links still apply.

Given the limitations of Layer 2 switching, there is still a need for Layer 3 functionality within the network.

NOTE Although switches are much faster than the traditional bridge, Layer 2 switching and bridging are logically equivalent and are used synonymously in this book.

Figure 1-7 *Sample Network with Layer 2 Switching*

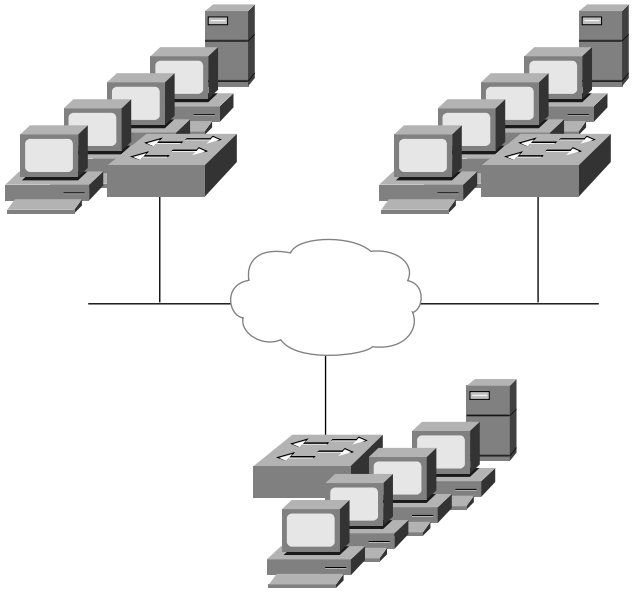

Benefits of Routing

Routers make optimal path decisions based on Layer 3 addressing. Routers improve network segmentation by determining the next network point to which a packet should be forwarded based on an optimal path decision. Routers do not forward Layer 2 broadcast frames, nor do routers forward multicast packets on a network that does not have any multicast clients.

Routing decisions must be performed for every packet received by the routing process. Incoming packets include a destination address (DA) field identifying a unique destination within the network. Routers use the DA to look up the next-hop router address and physical point in their routing tables.

After path determination is complete, packet forwarding is a well-defined set of packet manipulations. Figure 1-8 shows how routers make optimal path decisions based on Layer 3 protocols.

Figure 1-8 *An Example of a Optimized Routed Path*

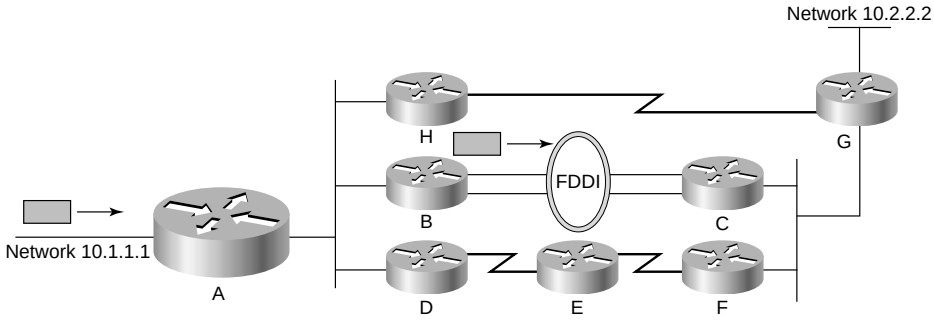

NOTE Each routed packet undergoes a very similar process. This process includes determining the Layer 3 destination by looking at the DA in the Layer 3 packet. The router then looks up the destination in the routing table in order to find the next-hop Layer 3 address. After the next-hop Layer 3 address has been determined, the router looks up the Layer 2 address of the next hop in a table that maps Layer 3 addresses to Layer 2 addresses. This would be, for example, an ARP table for IP. After both of these lookups have occurred, the frame undergoes a process sometimes called an *inline rewrite*. This process overwrites the old Layer 2 information with the new information, including the new source and destination MAC addresses, and it decrements or increments the Time To Live in the Layer 3 packet. After all this is complete, the packet is forwarded out the egress interface.

A network layer address identifies an entity at the network layer of the OSI reference model and is called a *virtual address* or *logical address*. Routers and other internetworking devices require one network-layer address per physical network connection for each network-layer protocol supported.

Because routers map a single Layer 3 logical address to a single network device, routers limit or secure network traffic based on identifiable attributes within each packet. These options can be applied to inbound or outbound packets on any router interface.

Concurrent with the increasing acceptance of Layer 2 switching as an essential component of network infrastructure are two other developments:

- Migration of servers to server farms for increased security and management of data resources
- Deployment of intranets, with organization-wide client/server communications based on Web technology

These factors are moving data flows off local subnets and onto the routed network, where the limitations of router performance can increasingly lead to bottlenecks.

Layer 3 Switching

Layer 3 switching is hardware-based routing. In particular, packet forwarding is handled by specialized hardware ASICs. A Layer 3 switch does everything to a packet that a traditional router does, such as the following:

- Determines the forwarding path based on Layer 3 information
- Validates the integrity of the Layer 3 header via checksum
- Verifies packet expiration and updates accordingly
- Processes and responds to any option information
- Updates forwarding statistics in the Management Information Base (MIB)
- Applies security controls if required

The primary difference between the packet-switching operation of a router and a Layer 3 switch is the physical implementation. In general-purpose routers, microprocessor-based engines typically perform packet switching. A Layer 3 switch performs packet switching with hardware.

NOTE A Layer 3 device, such as a router or switch, performs two basic functions. The first function is to make a path determination based on the information found in the Layer 3 address. This is done through the use of routing protocols to build routing tables. The routing tables are used to determine how a packet should move through the network. The second function that a router must perform is called *packet switching*. Packet switching is the process of rewriting the Layer 2 information, decrementing the Time To Live (TTL) field, and moving the frame from one interface to the next. This process of packet switching should not be confused with basic Layer 2 switching.

Cisco currently has two major implementations of Layer 3 switching for the Catalyst switch product: multilayer switching and Cisco Express Forwarding. Multilayer switching is covered in-depth in this book. Cisco Express Forwarding is covered in Appendix B.

High-performance packet-by-packet Layer 3 switching is achieved in different ways. For example, the Cisco 12000 Gigabit Switch Router (GSR) achieves wire-speed Layer 3 switching with a crossbar switch matrix. The Catalyst family of multilayer switches performs Layer 3 switching with special ASICs.

Because it is designed to handle high-performance LAN traffic, a Layer 3 switch can be placed anywhere within the network, cost-effectively replacing the traditional router.

NOTE Layer 3 switching and routing are logically equivalent and are used synonymously in this book.

Layer 4 Switching

Layer 4 switching refers to Layer 3 hardware-based routing that considers the application. Information in packet headers typically includes Layer 2 and Layer 3 addressing and the Layer 3 protocol type, plus more fields relevant to Layer 3 devices, such as TTL and checksum. The packet also contains information relevant to the higher layers within the communicating hosts, such as the protocol type and port number.

A simple definition of Layer 4 switching is the ability to make forwarding decisions based on not just the MAC address or source/destination IP addresses but on these Layer 4 parameters. In TCP or User Datagram Protocol (UDP) flows, the application is encoded as a port number in the segment header. Layer 4 switching is vendor-neutral and is beneficial even when added to preexisting network environments.

Cisco routers can control traffic based on Layer 4 information. One method of controlling Layer 4 traffic is by using standard or extended access lists. Another method is to provide Layer 4 accounting of flows using NetFlow switching.

Finally, when performing Layer 4 functions, a switch reads the TCP and UDP fields to determine what type of information the packet is carrying. The network manager can program the switch to prioritize traffic by application. This function allows network managers to define a quality of service (QoS) for end users. When being used for QoS purposes, Layer 4 switching means that a videoconferencing application might be granted more bandwidth than an e-mail message.

Layer 4 switching is necessary if your policy dictates granular control of traffic by application, or if you require accounting of traffic by application.

However, it should be noted that switches performing Layer 4 switching need the ability to identify and store large numbers of forwarding table entries. This is especially true if the switch is within the core of an enterprise network. Many Layer 2 and Layer 3 switches have forwarding tables that are sized in proportion to the number of network devices.

With Layer 4 switches, the number of network devices must be multiplied by the number of different application protocols and conversations in use in the network. Thus, the size of the forwarding table can quickly grow as the numbers of end devices and types of applications increase. This large table capacity is essential to creating a high-performance switch that supports wire-speed forwarding of traffic at Layer 4.

Multilayer Switching

Multilayer switching combines Layer 2 switching and Layer 3 routing functionality. Multilayer switching moves campus traffic at wire speed while at the same time satisfying Layer 3 routing requirements. This combination not only solves throughput problems but also removes the conditions under which Layer 3 bottlenecks form. Multilayer switching is based on the "route once, switch many" model.

Multilayer switching in the Catalyst family of switches can operate as a Layer 3 switch or a Layer 4 switch. When operating as a Layer 3 switch, the Catalyst family of switches caches flows based on IP addresses. When operating as a Layer 4 switch, the Catalyst family of switches caches conversations based on source address, destination address, source port, and destination port.

Because Layer 3 or Layer 4 switching is performed in hardware, there is no performance difference between the two modes of operation.

Multilayer switching is discussed in more detail in Chapter 6, "Improving IP Routing Performance with Multilayer Switching."

The Hierarchical Model

Campus network designs have traditionally placed basic network-level intelligence and services at the center of the network and shared bandwidth at the user level. Over the past few years, distributed network services and switching have migrated to the user level, and a distinct model has taken shape.

Figure 1-9 illustrates the hierarchical model and the devices within that model. The layers are defined as follows:

- Access layer
- Distribution layer
- Core layer

This approach allows designers to define building blocks that interconnect users and services. The building block encompasses both distributed network services and network intelligence.

The best-managed campus networks are typically designed following the hierarchical model. This model simplifies the management of the network and allows for controlled growth. The following sections describe the components of the hierarchical model.

Figure 1-9 *The Hierarchical Model*

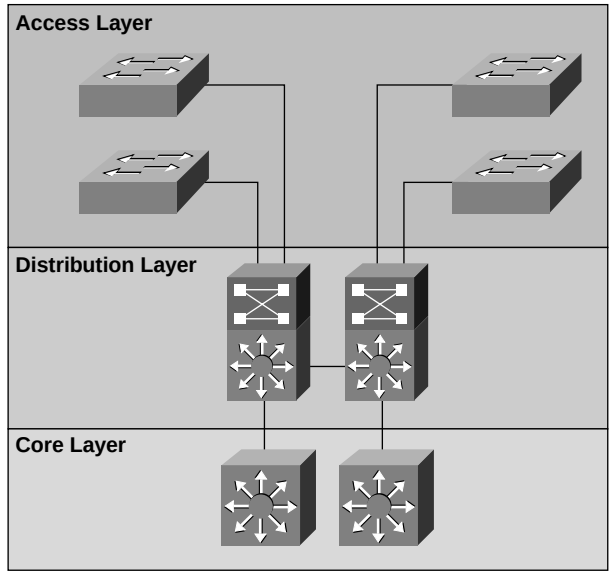

The Access Layer

The access layer of the network is the point at which end users are allowed into the network. This layer can provide further tuning in terms of filtering or access lists; however, the key function of this layer is to provide access for end users into the network. In the campus environment, some of the functions represented by the access layer are as follows:

- Shared bandwidth
- Switched bandwidth
- Layer 2 services, such as VLAN membership and traffic filtering based on broadcast or MAC addresses

The main criterion for access devices is to provide this functionality with low-cost, high-port density devices.

The Distribution Layer

The distribution layer of the network marks the point between the access and core layers of the network. The distribution layer also helps define and differentiate the core. This layer provides a boundary definition and designates where potentially expensive packet manipulations are handled. In the campus environment, the distribution layer can represent a multitude of functions, some of which are as follows:

- VLAN aggregation
- Departmental or workgroup access
- Broadcast or multicast domain definition
- Inter-VLAN routing
- Media translations
- Security

The distribution layer can be summarized as the layer that provides policy-based connectivity.

The Core Layer

The core layer is sometimes referred to as the backbone of a campus network. The primary purpose of the core layer is to switch traffic as fast as possible. This layer of the network should not be involved in expensive packet manipulation or any processing that slows down traffic switching. Functions such as access lists and packet filtering should be avoided in the core. The core layer is responsible for the following functions:

- Providing connectivity between switch blocks
- Providing access to other blocks, such as the WAN block
- Switching frames or packets as quickly as possible

Choosing a Cisco Product

Campus size is an important factor in network design. A large campus has several or many colocated buildings. A medium campus has one or several colocated buildings, and a small campus might have only one building.

The selection of Cisco products at a specific layer depends on the required functionality of each device. The following sections discuss each layer and the appropriate devices. For pictures and further explanation of the products explained in this section, see the Cisco Connection Online (CCO) Web site's product listings at http://www.cisco.com/public/products_prod.shtml.

Access Layer Switches

The Catalyst 1900 or 2820 series switch is an effective access device in a small or medium campus network, connecting individual desktops and 10BaseT hubs to distribution switches with high-speed connections.

The Catalyst 2900 series switch is also effective in providing network access to server clusters or end-user populations of less than 50 users that have high bandwidth

requirements. In these types of applications, such as in CAD/CAM and IC design environments, the 2900 series switch provides up to 1000 Mbps throughput for client/server applications and enterprise servers.

The Catalyst 4000 series provides an advanced high-performance enterprise switching solution optimized for connecting up to 96 users and data center server environments that require up to 36 Gigabit Ethernet ports. The Catalyst 4000 series leverages a multigigabit architecture for 10/100/1000 Mbps Ethernet switching.

The Catalyst 5000 series is an effective access device in large campus networks that need to provide network access for more than 100 end users. This switch supports 10/100/1000 Mbps Ethernet switching.

Distribution Layer Switches

Because the distribution layer switch is the aggregation point for multiple access switches, it must be able to handle the total amount of traffic from these devices. The distribution layer switch also must be able to participate in multilayer switching with a route processor. Therefore, the most effective distribution switch devices are the Catalyst 5000 series and 2926G switches. For Layer 3 switching, the Catalyst 5000 series switches support an internal route processor module, and the 2926G switch works with an external router, such as the 4x00 and 7x00 series routers.

The Catalyst 6000 switches are effective at the distribution level, where users require very high densities of Fast or Gigabit Ethernet—up to 384 10/100, 192 100FX, and 130 Gigabit Ethernet ports.

NOTE When choosing a distribution layer switch, remember to consider the amount of aggregate bandwidth to support the access layer devices and connectivity to the core layer. The Catalyst 5000 is best used in small to medium networks where the primary connectivity is 10/100 Mb Ethernet. The Catalyst 5000 has a 3.6 Gbps switch fabric, which can be oversubscribed if it supports many Gigabit Ethernet ports. The Catalyst 6000 is recommended in medium to large network environments due to its ability to handle Gigabit Ethernet better than the Catalyst 5000. The Catalyst 6000 has a 32 Gbps switch fabric, which allows it to handle a larger number of Gigabit Ethernet connections. Refer to Appendix B for a more detailed discussion of switch fabrics and oversubscription.

Core Layer Switches

For core backbone implementations, the Catalyst 6500 and 8500 series provide wire-speed multicast forwarding and routing, as well as the Protocol-Independent Multicast (PIM) protocol for scalable multicast routing.

Both of these switches provide the high bandwidth and performance that is required for a campus backbone. They are ideal for aggregating multiprotocol traffic from multiple wiring closets or from workgroup switches, such as the Catalyst 5000/6000.

The Building Block Approach

In an attempt to ease the design process of small to large campus networks, Cisco has created a set of building blocks for the network. These building blocks allow a design to scale as small or as large as is necessary.

Network building blocks can be any one of the following fundamental campus elements or contributing variables.

Campus elements:

- Switch block
- Core block

Contributing variables:

- Server block
- WAN block
- Mainframe block

The fundamental campus elements are described in this section. Figure 1-10 shows the campus elements as well as the contributing blocks, or variables, connected by the core block.

NOTE For more information on the various components of the building blocks, refer to the Campus Network Design Case Study white paper, located at the CCO Web site (http://www.cisco.com).

Figure 1-10 *Campus Network Building Blocks*

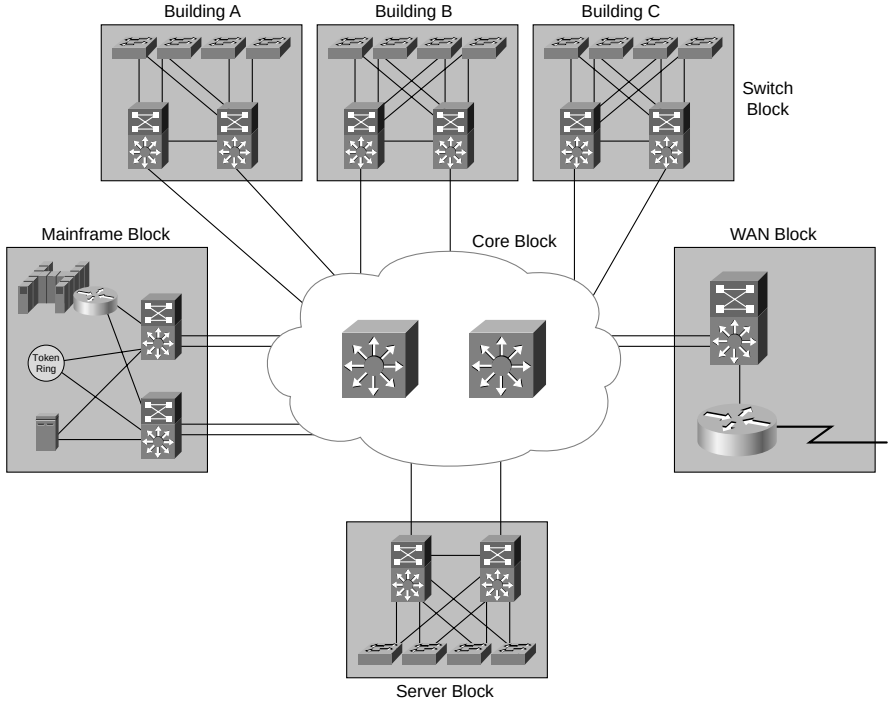

The Switch Block

The switch block contains a balanced implementation of scalable Layer 2 switching and Layer 3 services. Although the current generation of LAN switches is replacing shared media concentrators, LAN switches are not replacements for Layer 3 devices. Therefore, the switch block consists of both switch and router functionality. The switch block shown in Figure 1-11 prevents all broadcast traffic as well as network problems from traversing the core block and from reaching other switch blocks.

Figure 1-11 *The Switch Block*

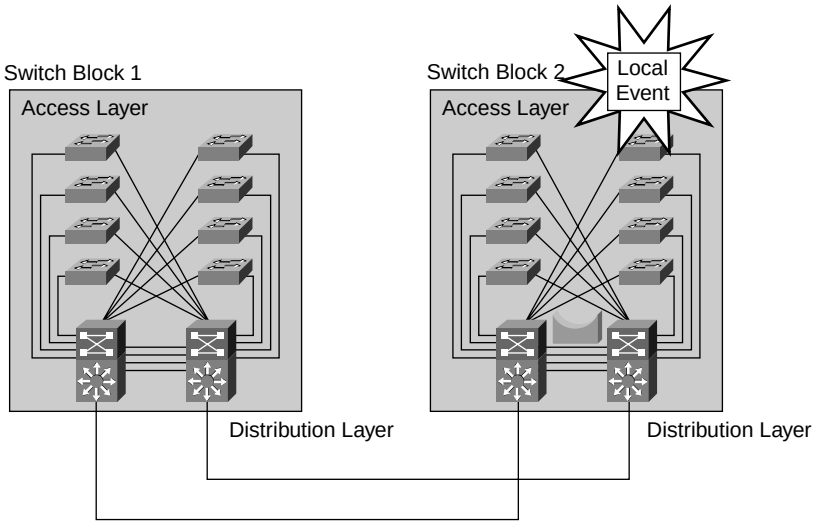

Layer 2 switches in the wiring closets connect users to the network at the access layer and provide dedicated bandwidth to each port. The Catalyst 2900 and 2820/1900 series switches provide cost-effective wiring closet connectivity.

The access devices merge into one or more distribution devices. The distribution device provides Layer 2 connectivity between access switches and acts as a central connection point for all of the switches in the wiring closets.

The distribution layer also provides Layer 3 functionality, which supports routing and networking services. The distribution layer shields the switch block against failures in other parts of the network.

The distribution device can be one of the following:

- A switch and external router combination
- A multilayer switch

These distribution layer devices are discussed in more detail in Chapter 5, “Inter-VLAN Routing.”

If the switch block experiences a broadcast storm, the router prevents the storm from propagating into the core and into the rest of the network. Each block is protected from the other blocks when failures occur. However, the switch block, in which the broadcast storm occurs, still experiences network problems until the device generating the broadcasts is found and removed from the network.

NOTE Cisco has attempted to address the problem of broadcast storms by adding broadcast thresholds to their switches. A broadcast threshold prevents the port from receiving broadcasts until the number of broadcasts drops below a specific number of broadcasts per second or a specific percentage of traffic. This prevents the broadcasts from overwhelming the device, such as a workstation, on the other end of the port.

Switch Block Characteristics

Access layer switches may support one or more subnets. A subnet must reside within one broadcast domain. This means that all stations residing in or ports configured on the same VLAN are assigned network addresses within the same subnet. However, a single VLAN can support multiple subnets.

The broadcast isolation feature of VLANs is the characteristic that allows VLANs to be identified with subnets. For example, the IP Address Resolution Protocol (ARP) propagates only within the VLAN of the originating request. All subnets terminate on Layer 3 devices, such as a router or a Route Switch Module (RSM). To connect to devices in other VLANs, the frame must traverse a router. In this model, VLANs should not extend beyond the distribution switch.

Access layer switches have redundant connections, or *uplinks*, to the distribution switch to maintain resiliency. The Spanning-Tree Protocol allows these redundant links to exist while preventing undesirable loops in the switch block. The Spanning-Tree Protocol terminates at the boundary of the switch block.

All switches in the network may be connected to a common management default subnet. VLAN management domains are discussed in greater detail in Chapter 3, “Defining Common Workgroups with VLANs.”

NOTE Cisco philosophy discourages trunking across the core unless discrete VLAN1 connections are made between the distribution and core switches.

Sizing the Switch Block

Although the size of a switch block is flexible, certain factors limit the size of this block. The number of switches that collapse into the distribution layer depends on several factors:

- Different types and patterns of traffic
- Amount of Layer 3 switching capacity at the distribution layer
- Number of users per access layer switch

- Extent to which subnets need to traverse geographical locations within the network
- Size to which the Spanning-Tree domains should be allowed to grow

There are two main factors in sizing the switch block:

- Traffic types and behavior
- Size and number of workgroups

NOTE This model is based on no more than 2000 users per switch block.

A switch block is too large if

- A traffic bottleneck occurs in the routers at the distribution layer due to intensive CPU processing required by services such as access lists
- Broadcast or multicast traffic slows down the switches and routers

NOTE The decision to break up the block should be based on the traffic going across the network, rather than the specific number of nodes in a building block. It is important to take periodic snapshots of traffic flows in order to be able to determine when to break up the switch block.

The Core Block

A core is required when there are two or more switch blocks. The core block is responsible for transferring cross-campus traffic without any processor-intensive operations, such as routing. All the traffic going to and from the switch blocks, the server blocks, the Internet, and the wide-area network passes through the core.

Traffic going from one switch block to another also must travel through the core. Because of these traffic patterns, the core handles much more traffic than any other block. Therefore, the core must be able to pass the traffic to and from the blocks as quickly as when there are two or more switch blocks.

In Figure 1-12, the core block supports frame, packet, or cell-based technologies, depending on your specific needs. This self-study book discusses and demonstrates an Ethernet core. Because the distribution switch provides Layer 3 functionality, individual subnets will connect all distribution and core devices.

Figure 1-12 *The Core Layer*

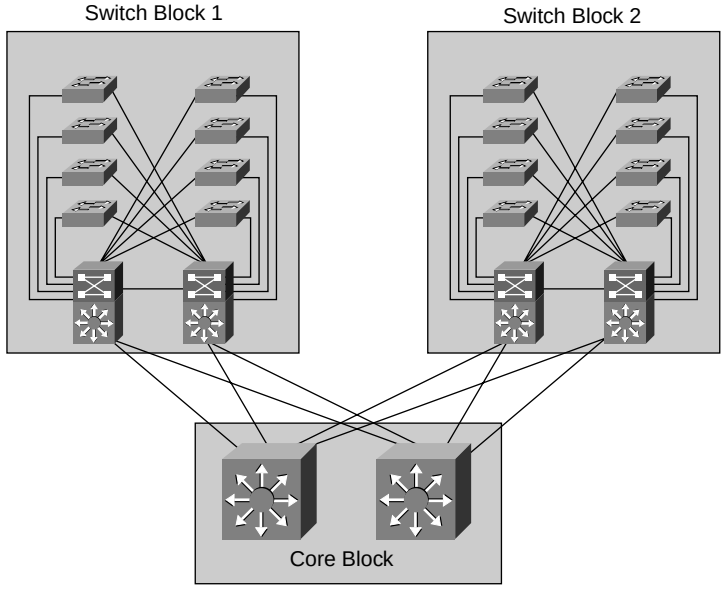

The core can consist of one subnet; however, for resiliency and load balancing, at least two subnets are configured. Because VLANs terminate at the distribution device, core links are not trunk links, and traffic is routed across the core. Therefore, core links do not carry multiple subnets per link.

One or more switches make up the core subnet; however, Cisco recommends that a minimum of two devices be present in the core to provide redundancy. The core block can consist of high-speed Layer 2 devices such as Catalyst 5500 or 6500 series switches or Layer 3 devices such as the 8500 series routers.

At a minimum, the media between the distribution switches and the core layer switches should be capable of supporting the amount of load handled by the distribution switch.

At a minimum, the links between core switches in the same core subnet should be sufficient to switch the aggregate amount of traffic with respect to the input aggregation switch traffic. The design of the core should consider average link utilization while still allowing for future traffic growth.

There are two basic core designs:

- Collapsed core
- Dual core

Collapsed Core

The collapsed core exists when both the distribution and core layer functions are performed in the same device. A collapsed core design is prevalent in small campus networks. Although the functions of each layer are contained in the same device, the functionality remains distinctly separate. Figure 1-13 shows a sample collapsed core.

Figure 1-13 *A Collapsed Core*

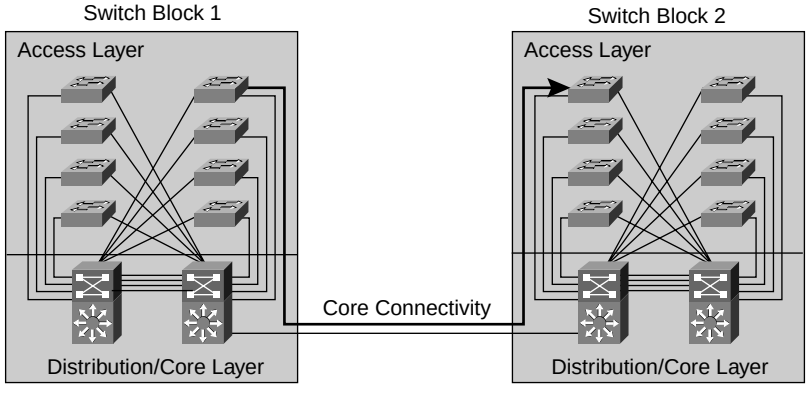

In a collapsed core design, each access layer switch has a redundant link to the distribution layer switch. Each access layer switch may support more than one subnet; however, all subnets terminate on Layer 3 ports on the distribution switch.

Redundant uplinks provide Layer 2 resiliency between the access and distribution switches. Spanning Tree blocks the redundant links to prevent loops.

Redundancy is provided at Layer 3 by the dual distribution switches with the Hot Standby Router Protocol (HSRP), providing transparent default gateway operations for IP. In the event that the primary routing process fails, connectivity to the core is maintained. HSRP is discussed in more detail in Chapter 7, “Configuring HSRP for Fault-Tolerant Routing.”

Dual Core

A dual-core configuration is necessary when two or more switch blocks exist and redundant connections are required. Figure 1-14 shows a dual-core configuration in which the core contains only Layer 2 switches for the backbone. The core devices are not linked to avoid any bridging loops.

Figure 1-14 *A Dual Core*

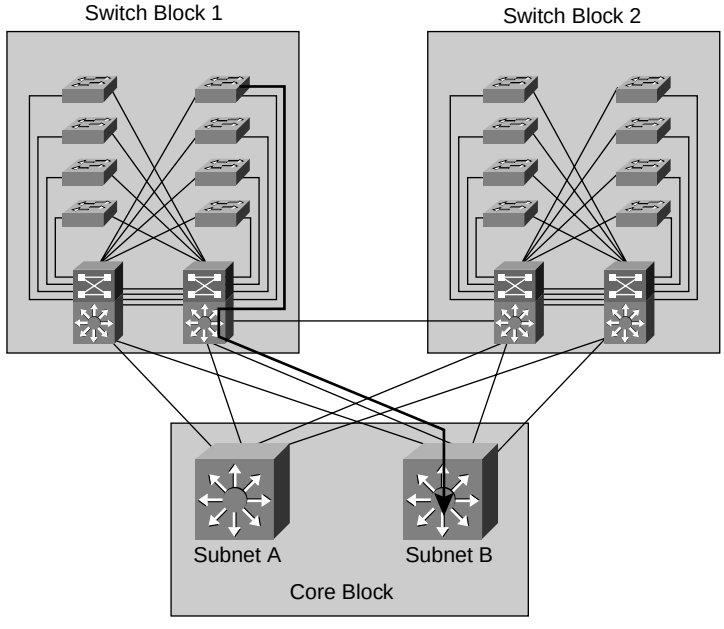

A dual-core topology provides two equal-cost paths and twice the bandwidth. Each core switch carries a symmetrical number of subnets to the Layer 3 function of the distribution device. Each switch block is redundantly linked to both core switches, allowing for two distinct and equal path links. If one core device fails, convergence is not an issue, because the routing tables in the distribution devices already have an established route to the remaining core device. The Layer 3 routing protocol provides the link determination across the core, and HSRP provides quick failover determination. Spanning Tree is not needed in the core, because there are no redundant links between the core switches.

Sizing the Core

Because Layer 3 devices isolate the core, routing protocols are used to maintain the current state of the network. As the routing protocol sends these updates and changes to the routers throughout the network, the network topology might also change. The more routers connected to the network, the longer it takes these updates and changes to propagate throughout the network and change the topology. Also, one or more of the routers might connect to a WAN or the Internet, which adds more sources of routing updates and topology changes.

The routing protocol used on the Layer 3 devices determines the number of distribution devices that can be attached to the core. Table 1-2 gives examples of some widely used

protocols and the maximum number of peer routers for which these protocols can maintain state information.

Table 1-2 *Maximum Number of Supported Blocks by Routing Protocol*

Routing Protocol	Maximum Number of Supported Routing Peers	Number of Subnet Links to the Core	Maximum Number of Supported Blocks
Open Shortest Path First (OSPF)	50	2	25
Enhanced Interior Gateway Routing Protocol (EIGRP)	50	2	25
Routing Information Protocol (RIP)	30	2	15

All blocks, including server, WAN, and mainframe blocks, are included in the maximum number of switch blocks supported in Table 1-2. The figures given in Table 1-2 are very optimistic for the total number of peers supported by each of the routing protocols. In reality, the maximum number of peers for each of the routing protocols should be closer to 15.

Layer 2 Backbone Scaling

Switched Layer 2 Ethernet cores are very cost-effective and provide high-performance connectivity between switch blocks. The classic design model has several switch blocks, each supporting Layer 2 devices in a wiring closet terminating into a Layer 3 device. The Layer 3 devices are connected by a core composed of Layer 2 devices.

The Spanning-Tree Protocol represents a practical limit to the scale of a Layer 2 switched backbone. As you increase the number of core devices, you need to increase the number of links from the distribution switches to maintain redundancy. Because routing protocols dictate the number of equal-cost paths, the number of independent core switches is limited. Interconnecting the core switches creates bridging loops. Introducing the Spanning-Tree Protocol into the core compromises the high-performance connectivity between switch blocks. Ideally, Layer 2 switched backbones consist of two switches with no Spanning-Tree loops in the topology.

Figure 1-15 has eight distribution layer switches that connect to the core. The distribution layer switches have connections to each of the core layer switches. Note that the core layer switches are not connected to each other in order to prevent a loop. With the two core layer switches there are now two equal-cost paths to all VLAN destinations.

Figure 1-15 *Scaling the Core Block with Layer 2*

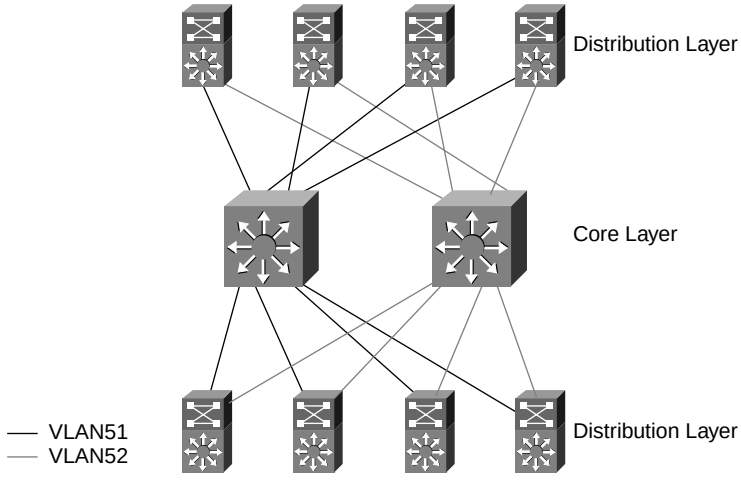

In Figure 1-15, the Layer 2 switched backbone provides redundancy without any Spanning-Tree loops. Because the two core switches are not linked, loops do not occur. Each distribution switch in every switch block has a separate path to each core switch. The dual connection between each core and distribution device provides each Layer 3 device with two equal-cost paths to every other router in the campus.

Layer 3 Backbone Scaling

Most of the designs successfully follow the model that has been presented so far, with Layer 2 at the access layer, Layer 3 at the distribution layer, and Layer 2 at the core layer. This is sometimes referred to as the Layer 2-Layer 3-Layer 2 model. However, there are designs in which having a Layer 3 core layer is advantageous.

You would implement a Layer 3 core for the following reasons:

- Fast convergence
- Automatic load balancing
- Eliminate peering problems

Figure 1-16 shows a Layer 3 core. Each connection from the distribution layer devices is a separate subnet, and the core layer devices are responsible for routing between each of these subnets.

Figure 1-16 *Layer 3 Backbone Scaling*

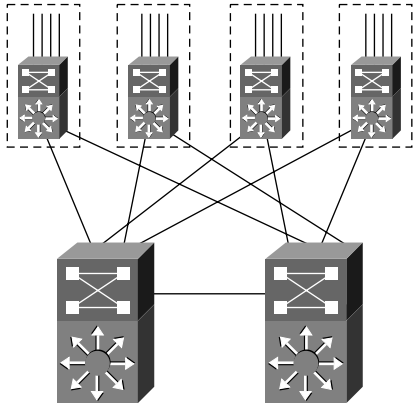

Fast Convergence

As you increase the number of switch blocks and servers, each distribution layer device must be connected to the core. Because there is a limit to the number of switch blocks to a dual Layer 2 core, increasing the number of connections means increasing the number of core devices. To maintain redundancy, the core devices must be connected. After you interconnect Layer 2 devices, bridging loops appear. To eliminate bridging loops in the core, you must enable the Spanning-Tree Protocol. The Spanning-Tree Protocol might have a convergence time of over 50 seconds. If there is a fault in the network's core, Spanning-Tree Protocol convergence can disable a network core for more than one minute.

With the implementation of Layer 3 devices in the core, the Spanning-Tree Protocol becomes unnecessary. In this design, routing protocols are used to maintain the network topology. Convergence for routing protocols takes anywhere from 5 to 10 seconds, depending on the routing protocol.

Automatic Load Balancing

Load balancing tries to achieve a traffic distribution pattern that will best utilize the multiple links that provide redundancy. With multiple, interconnected Layer 2 devices in the core, you must selectively choose the root for utilizing more than one path. You then manually configure the links to support specific VLAN traffic.

With Layer 3 devices in the core, the routing protocols can load balance over multiple equal-cost paths.

Eliminate Peering Problems

Another issue with the Layer 2 core in a very large network involves router peering. Router peering ensures that the routing protocol running within a router maintains state and reachability information about other neighboring routers. In this scenario, each distribution device becomes a peer with every other distribution device in the network. Scalability becomes an issue in the configuration, because each distribution device must maintain state for all other distribution devices.

With the implementation of Layer 3 devices in the core, a hierarchy is created, and the distribution device is not considered a peer with all other distribution devices. This type of core might appear in very large campus networks in which the network supports more than 100 switch blocks.

Implementing Layer 3 devices in the core is expensive. As stated earlier, the main purpose of the core is to move packets as quickly and efficiently as possible. Although Layer 3 devices can support the switching of some protocols, both performance and equipment cost become an issue.

Campus Network Availability Example

The design shown in Figure 1-17 consists of two buildings: North and South. Each building has 10 floors and 1000 users. Each floor is connected to an access switch in the wiring closet. Each access switch is linked to a distribution layer switch.

If a link from a wiring closet switch to the distribution-layer switch is disconnected, 100 users on a floor could lose their connections to the backbone. To prevent this, each access switch has a link to each distribution switch in the building. The Spanning-Tree Protocol blocks the redundant link to prevent loops.

Load balancing across the core is achieved by intelligent Layer 3 routing protocols implemented in the Cisco IOS software. In Figure 1-17, there are four equal-cost paths between the two buildings. The four paths from the North building to the South building are AXC, AYD, BXC, and BYD. These four Layer 2 paths are considered equal by Layer 3 routing protocols. Note that all paths from both buildings to the backbone are single, logical hops.

Figure 1-17 *A Campus Network Availability Example*

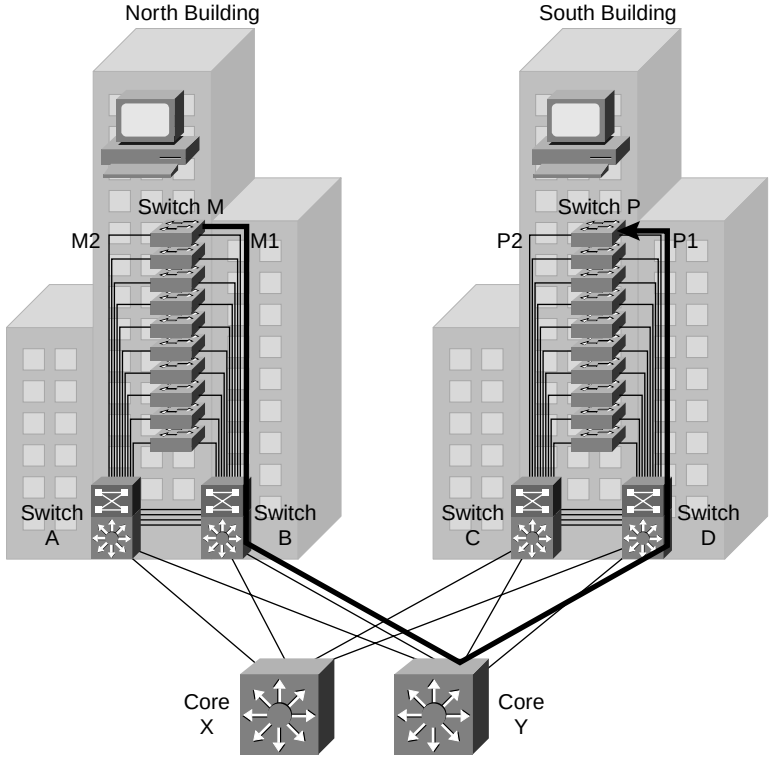

In this scenario, a user attached to access switch M wants to transfer data to a user attached to switch P. The active router is multilayer switch B, and HSRP is enabled. As shown in Figure 1-17, the logical path for the data from M to P is as follows:

- Access switch M switches the data over link M1 to multilayer switch B.
- Multilayer switch B routes the data out subnet BY to core Y.
- Core Y switches the information out link YD to multilayer switch D.
- Multilayer switch D switches the data over link P1 to access switch P.

If the M1 link fails, however, as shown in Figure 1-18, the redundant M2 link becomes the primary link, and the data path from M to P is as follows:

- Access switch M switches the data over link M2 to multilayer switch A.
- Multilayer switch A switches the data to the active HSRP router, multilayer switch B.
- Multilayer switch B routes the data out subnet BY to core Y.
- Core Y switches the information out link YD to multilayer switch D.
- Multilayer Switch D switches the data over link P1 to access switch P.

Figure 1-18 *Access Layer Switch Failure Causes Data to Fail Over to the Redundant Link*

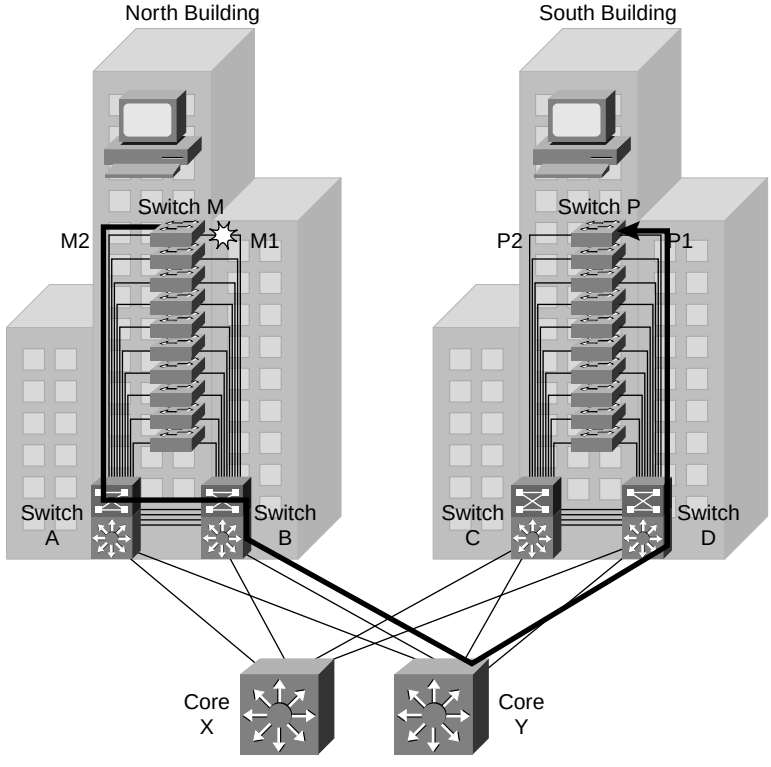

A distribution layer switch represents a point of failure at the building level. One thousand users in the North building could lose their connections to the backbone in the event of a disabled route processor. To provide fault-tolerant access to each user, redundant links connect access layer switches to a pair of Catalyst multilayer switches in the distribution layer, as shown in Figure 1-19.

In Figure 1-19, the route processors in both distribution devices are connected, and HSRP is configured. This configuration allows for fast failover at Layer 3 if one of the distribution switches fails. In this scenario, the data path from M to P is as follows:

- Access switch M switches the data over link M2 to multilayer switch A.
- Because multilayer switch B is disabled, multilayer switch A becomes the active HSRP router and routes the data out subnet AY to core Y.
- Core Y switches the information out link YD to multilayer switch D.
- Multilayer switch D switches the data over link P1 to access switch P.

Figure 1-19 *Redundancy in the Distribution Layer*

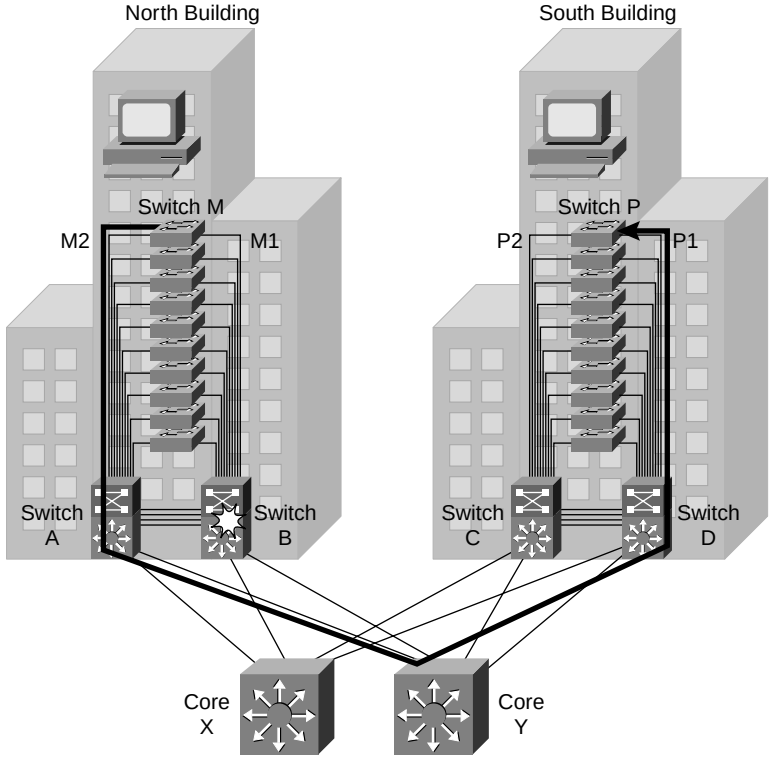

Redundancy in the backbone is achieved by installing two or more Catalyst switches in the core. Each link from the distribution switch to the core is an equal-cost path. This topology provides failover as well as load balancing from each distribution switch across the backbone.

Load balancing across the core is achieved by intelligent Layer 3 routing protocols implemented in the Cisco IOS software. Figure 1-20 shows an example.

If one switch in the core fails, the data path from M to P is as follows:

- Access switch M switches the data over link M2 to multilayer switch A. Link M1 is still disabled.
- Because multilayer switch B is disabled, multilayer switch A becomes the active HSRP router.
- Because core Y is disabled, the data is routed out subnet AX to core X.
- The path at the distribution layer depends on how the ports and VLANs on those devices are configured. In this scenario, the data core X switches the information out link XC to multilayer switch C.

- Multilayer switch C has the MAC address of workstation P in the CAM table, and it switches the data over link P2 to access switch P.

Figure 1-20 *Example of a Core Layer Failure*

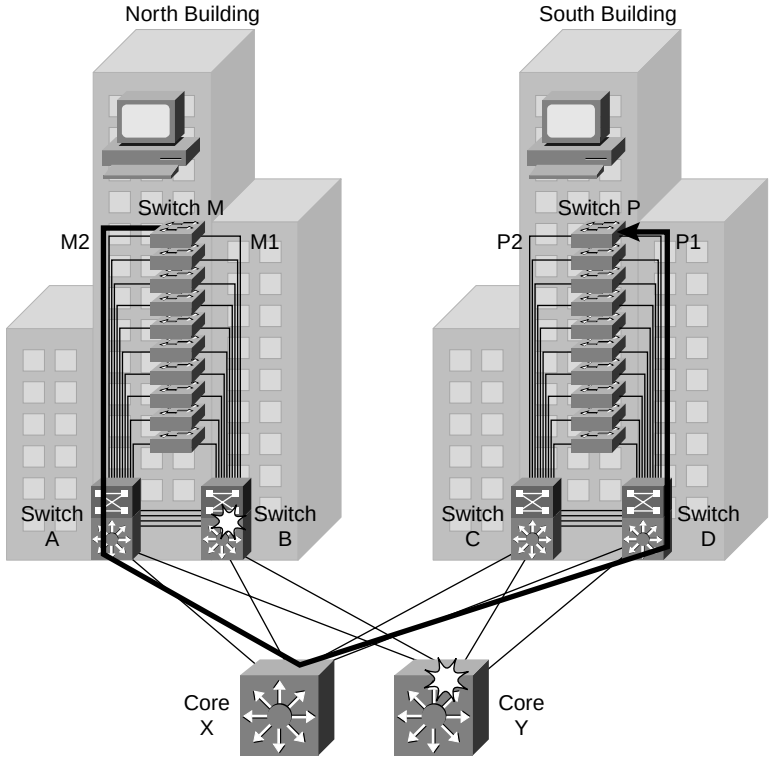

Summary

Many problems in the current campus networks can be solved with proper design and an understanding of the ever-changing traffic patterns. A redundant, well-planned design will provide users of the campus network with availability and fault tolerance and will provide the network administrator with a deterministic flow of traffic.

Here are some of the issues and solutions discussed in this chapter:

- Broadcasts are useful and necessary traffic; however, too much broadcast traffic can cause network performance problems. Managing broadcast traffic is a critical aspect of campus LAN design.
- The location of common workgroups and servers can have a significant impact on traffic patterns.

- Adding more bandwidth is not the long-term solution to meeting the needs of high-priority traffic.
- Multilayer switching combines Layer 2 switching and Layer 3 routing functionality.
- The multilayer design model is inherently scalable. Layer 3 switching performance scales because it is distributed. Backbone performance scales as you add more links or switches.
- A switch block is the unit that contains distributed network services and network intelligence. A switch block consists of Layer 2 switches, Layer 3 routers, and, sometimes, distributed servers.
- A core block is the unit that transfers cross-campus traffic. It can consist of Layer 2 or Layer 3 devices.
- The purpose of network link redundancy is to provide alternative physical pathways through the network in case one pathway fails.
- Routing updates and route changes might propagate to all of the routers on the campus, depending on which routing protocol is in use.
- The number of equal-cost paths supported by the routing protocol also determines the size of the core. For instance, if the routing protocol supports up to six equal-cost paths, these paths should be distributed over a different physical Layer 2 device to maintain a 1:1 redundancy. Hence, six equal-cost paths require six Layer 2 devices.
- Finally, at a minimum, the core layer switches must be capable of scaling to the following:

  $n \times$ amount of load per link at 100 percent capacity = total core switch capacity

  where n is equal to the number of distribution links.

Review Questions

The following questions test your retention of the material presented in this chapter. The answers to the Review Questions can be found in Appendix A, “Answers to Review Questions.”

1. Discuss the various trends that have forced the redesign of campus networks.
2. Describe the different switching technologies and how they enable multilayer switching.
3. Explain the multilayer model and how it affects traffic flows in the network.

Written Exercises: Overview of a Campus Network

These exercises help you identify the different switching technologies implemented at OSI Layers 2, 3, and 4. The answers to each task can be found at the end of this chapter.

Task 1: Describing Layer 2, 3, and 4 and Multilayer Switching Functions

For each network function listed in the following table, choose the switching technology that most accurately supports each function, and place the letter of that switching technology next to the function. Each function can have more than one correct answer, so list all possible answers. The first answer is given as an example.

A. Layer 2 switching
B. Layer 3 switching
C. Layer 4 switching
D. Multilayer switching

B	Hardware-based routing
	Hardware-based bridging
	Forwards packets based on application port numbers
	Based on the principle of “route once, switch many”
	Promotes flatter network designs with fewer subnets
	Controls traffic using access lists based on port identifiers
	Provides traffic flow accounting through NetFlow switching
	Combines Layer 2 switching and Layer 3 routing functionality
	Uses frames to communicate with peer layers in another system
	Handles frame forwarding using specialized hardware called ASICs
	Uses packets to communicate with peer layers in another system
	Allows the switch to be programmed to prioritize traffic by application
	Interrogates the initial packet in a flow, which is then forwarded to a cache
	Allows the prioritization of traffic based on specific applications
	Does not modify frame infrastructure when moving frames between like media

	Reads the TCP or UDP field to determine what type of information the packet is carrying
	Repackages multiport bridging technology with significant performance improvements and increased scalability
	Implements a level of security, interrogating traffic based on the source address, destination addresses, and port number

Task 2: Identifying the Switch Layer Solution for a Given Network Requirement

Read each statement carefully. From the following list, choose the corresponding letter of the layer that best fits the description given in each statement. There is only one possible answer for each description:

A. Access layer
B. Distribution layer
C. Core layer

1 This layer is responsible for routing traffic between VLANs.

2 The primary objective of this layer is to switch traffic as fast as possible.

3 This layer provides communication between switch blocks and to the enterprise servers.

4 The key function of this layer is to provide access for end users into the network.

5 This layer provides segmentation and terminates collision and broadcast domains.

6 This layer blocks or forwards traffic into or out of the switch block, based on Layer 3 parameters.

7 Some of the functions represented by this layer are shared bandwidth, VLAN membership, and traffic filtering based on broadcast addresses.

8 The main criterion for devices at this layer is to provide LAN segmentation with low-cost, high-port density devices.

Task 3: Given a Set of User Requirements, Identify the Correct Cisco Product Solution

Read the following scenarios and choose the most appropriate Cisco solution from the list following each scenario.

1 The ABC Company is a small widget-distributing company that wants to interconnect users on multiple floors in the same building. To date, the company has only 15 employees, but it plans to triple that number in the next two years. Because of the nature of the business, users need to access large files on the workgroup servers. What is the most appropriate device for the access layer?

A. Catalyst 8500 series switch

B Catalyst 5500 series switch with an internal RSM

C. Catalyst 1900 series switch with 10BaseT ports

D. Catalyst 2900 series switch with 100BaseTX ports

2 The Acme Engineering Company has redesigned its campus network to support three switch blocks containing 2000 end users in each block. The network administrator wants to control broadcast domains to each individual switch block while still allowing inter-VLAN routing within and between switch blocks. What is the most appropriate device for the distribution layer?

A Catalyst 8500 series switch

B. Catalyst 4000 series switch

C. Catalyst 5500 series switch with an internal RSM

D. Catalyst 1900 series switch with a two-port 100Base uplink module

3 The Tool & Die Manufacturing Company has experienced 300 percent growth over the last year and has recently installed a new multimedia center for distributing company information throughout the campus. The company has requirements for gigabit-speed data transfer, high availability, and inter-VLAN routing between the end users and the enterprise server farms. What is the most appropriate device for the distribution layer?

A Catalyst 8500 series switch

B. Catalyst 1900 series switch with 12 10BaseT ports

C Catalyst 4000 series switch with a six-port Gigabit Ethernet module

D. Catalyst 6000 series switch with a 16-port Gigabit Ethernet module

4 The Rataxes Toy Company needs to interconnect its four separate campus buildings with a high-speed, high-bandwidth backbone. Each building supports a separate department, and each department supports a different network protocol. The network designer has already recommended a Catalyst 6000 series switch at the distribution layer. What is the most appropriate device for the core layer?

A. Catalyst 8500 series switch

B. Catalyst 1900 series switch with 12 10BaseT ports

C. Catalyst 4000 series switch with a six-port Gigabit Ethernet module

D. Catalyst 5500 series switch with a 24-port group switched 100BaseT Ethernet module

Task 1 Answers: Describing Layer 2, 3, and 4 and Multilayer Switching Functions

The following table contains the correct answers for Task 1.

B	Hardware-based routing
A	Hardware-based bridging
B & C	Forwards packets based on application port numbers
D	Based on the principle of “route once, switch many”
A	Promotes flatter network designs with fewer subnets
B & C	Controls traffic using access lists based on port identifiers
B & C	Provides traffic flow accounting through NetFlow switching
D	Combines Layer 2 switching and Layer 3 routing functionality
A	Uses frames to communicate with peer layers in another system
A	Handles frame forwarding using specialized hardware called ASICs
B	Uses packets to communicate with peer layers in another system
C	Allows the switch to be programmed to prioritize traffic by application
D	Interrogates the initial packet in a flow, which is then forwarded to a cache
C	Allows the prioritization of traffic based on specific applications
A	Does not modify frame infrastructure when moving frames between like media
C	Reads the TCP or UDP field to determine what type of information the packet is carrying

A	Repackages multiport bridging technology with significant performance improvements and increased scalability
C	Implements a level of security, interrogating traffic based on the source address, destination addresses, and port number

Task 2 Answers: Identifying the Switch Layer Solution for a Given Network Requirement

The following are the correct answers for Task 2.

1. This layer is responsible for routing traffic between VLANs.

   B. Distribution layer

2. The primary objective of this layer is to switch traffic as fast as possible.

   C. Core layer

3. This layer provides communication between switch blocks and to the enterprise servers.

   C. Core layer

4. The key function of this layer is to provide access for end users into the network.

   A. Access layer

5. This layer provides segmentation and terminates collision and broadcast domains.

   B. Distribution layer

6. This layer blocks or forwards traffic into or out of the switch block, based on Layer 3 parameters.

   B. Distribution layer

7. Some of the functions represented by this layer are shared bandwidth, VLAN membership, and traffic filtering, based on broadcast addresses.

   A. Access layer

8. The main criterion for devices at this layer is to provide LAN segmentation with low-cost, high-port density devices.

   A. Access layer

Task 3 Answers: Given a Set of User Requirements, Identify the Correct Cisco Product Solution

The following are the correct answers for Task 3.

1 The ABC Company is a small widget-distributing company that wants to interconnect users on multiple floors in the same building. To date, the company has only 15 employees, but it plans to triple that number in the next two years. Because of the nature of the business, users need to access large files on the workgroup servers. What is the most appropriate device for the access layer?

C. Catalyst 1900 series switch with 10BaseT ports

2 The Acme Engineering Company has redesigned its campus network to support three switch blocks containing 2000 end users in each block. The network administrator wants to control broadcast domains to each individual switch block while still allowing inter-VLAN routing within and between switch blocks. What is the most appropriate device for the distribution layer?

C. Catalyst 5500 series switch with an internal RSM

3 The Tool & Die Manufacturing Company has experienced 300 percent growth over the last year and has recently installed a new multimedia center for distributing company information throughout the campus. The company has requirements for gigabit-speed data transfer, high availability, and inter-VLAN routing between the end users and the enterprise server farms. What is the most appropriate device for the distribution layer?

D. Catalyst 6000 series switch with a 16-port Gigabit Ethernet module

4 The Rataxes Toy Company needs to interconnect its four separate campus buildings with a high-speed, high-bandwidth backbone. Each building supports a separate department, and each department supports a different network protocol. The network designer has already recommended a Catalyst 6000 series switch at the distribution layer. What is the most appropriate device for the core layer?

A. Catalyst 8500 series switch