CHAPTER 4

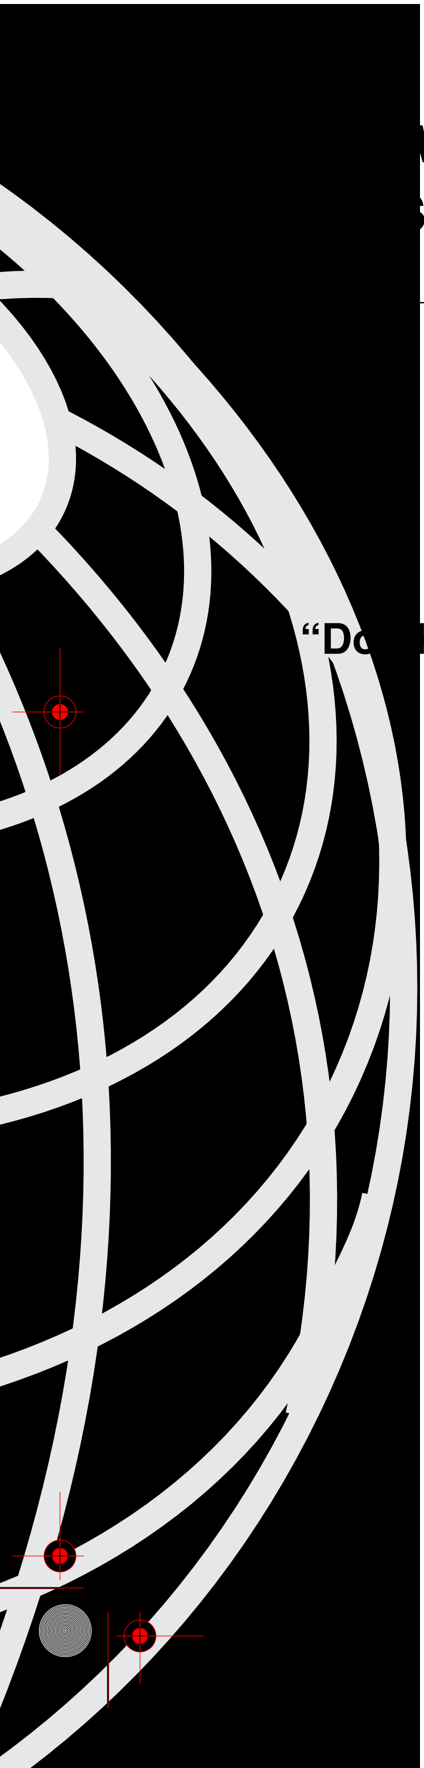

work Topologies and LAN sign

This chapter reviews the topologies used in network design and covers the technologies and design approaches used when designing a local-area network (LAN). The hierarchical, redundant, and secure topology models are covered. Technologies like Ethernet, Fast Ethernet, FDDI, and Token Ring are also covered in this chapter. This chapter also discusses the characteristics of repeaters, bridges, switches, and routers, as well as how to apply these devices in a LAN environment. Finally, this chapter covers the Cisco products used in local-area networks.

“Do Know This Already?” Quiz

The questions in the following quiz are designed to help you gauge how well you know the material covered in this chapter. Compare your answers with those found in Appendix A, “Answers to Quiz Questions.” If you answer most or all of the questions thoroughly and correctly, you might want to skim the chapter and proceed to the “Q&A” and “Case Studies” sections at the end of the chapter. If you find you need to review only certain subject matter, search the chapter for only those sections that cover the objectives you need to review and then test yourself with those question again, as well as the “Q&A” and “Case Studies” questions. If you find the following questions too difficult, read the chapter carefully until you feel you can easily answers these and the “Q&A” and “Case Studies” questions.

1 What OSI layer does a bridge operate?

2 The 10Base2 Ethernet media is commonly referred as?

3 What is the recommended maximum number of nodes that should be used in a multiprotocol LAN segment?

4 Bridges control collision domains, broadcast domains, or both?

5 What is the maximum segment size in a 100BaseT network?

6 What is the maximum segment size in a 10Base2 network?

7 Routers operate on what OSI layer?

8 Fast Ethernet is covered by which IEEE standard?

9 What is 10Base5 commonly referred to as?

10 What device controls a broadcast domain?

You can find the answers to these questions in Appendix A, "Answers to Quiz Questions."

Foundation Topics

LAN Topology Design

The CCDA objectives covered in this section are as follows:

13	Describe the advantages, disadvantages, scalability issues, and applicability of standard internetwork topologies.
14	Draw a topology map that meets the customer's needs and includes a high-level view of internetworking devices and interconnecting media.

This section covers CCDA exam objectives about designing network topologies for the LAN. LANs provide data transfer rates that are typically much faster than wide-area networks (WANs). While most companies own their own LAN infrastructure, wide-area connections between LANs are usually leased on a monthly basis from an outside carrier. With the recent developments in Gigabit Ethernet technologies, LAN designs are now capable of 1000 Mbps speeds. High-speed Gigabit links can connect servers to LAN switches. At these speeds, the capacity is there to meet the performance requirements of current high-bandwidth applications.

Various speeds of Ethernet have evolved into the de facto standard for LANs. Ethernet uses a contention-based access method, meaning each device competes simultaneously for access to the network. All devices attached to the same Ethernet segment form a collision domain. Each device transmitting on that segment may attempt to transmit at the same time as another device on the same segment, resulting in a collision. As the number of devices in the same collision domain increases, so do the collisions, resulting in poorer performance.

Although not discussed in newer switched (bridged) networks, legacy Ethernet networks with repeaters and hubs should limit the size of the collision domain. To scale multiprotocol networks and networks with high-bandwidth applications, limit the size of collision domains using bridges, switches, and routers. This is covered in the section "LAN Hardware" later in the chapter.

Three different network topology models are discussed in the following sections:

- Hierarchical models
- Redundant models
- Secure models

Hierarchical Models

Hierarchical models enable you to design internetworks in layers. To understand the importance of layering, consider the Open System Interconnection (OSI) reference model, which is a layered model for implementing computer communications. Using layers, the OSI model simplifies the tasks required for two computers to communicate. Hierarchical models for internetwork design also use layers to simplify the tasks required for internetworking. Each layer can be focused on specific functions, allowing you to choose the right systems and features for each layer. Hierarchical models apply to both LAN and WAN design.

Benefits of Hierarchical Models

The many benefits of using hierarchical models for your network design include the following:

- Cost savings
- Ease of understanding
- Easy network growth
- Improved fault isolation

After adopting hierarchical design models, many organizations report cost savings because they are no longer trying to do it all in one routing/switching platform. The modular nature of the model enables appropriate use of bandwidth within each layer of the hierarchy, reducing wasted capacity.

Keeping each design element simple and small facilitates ease of understanding, which helps control training and staff costs. Management responsibility and network management systems can be distributed to the different layers of modular network architectures, which also helps control management costs.

Hierarchical design facilitates changes. In a network design, modularity allows creating design elements that can be replicated as the network grows, facilitating easy network growth. As each element in the network design requires change, the cost and complexity of making the upgrade is contained to a small subset of the overall network. In large, flat, or meshed network architectures, changes tend to impact a large number of systems.

Improved fault isolation is facilitated by structuring the network into small, easy-to-understand elements. Network managers can easily understand the transition points in the network, which helps identify failure points.

Today’s fast-converging protocols were designed for hierarchical topologies. To control the impact of routing overhead processing and bandwidth consumption, modular hierarchical topologies must be used with protocols designed with these controls in mind, such as EIGRP.

Route summarization is facilitated by hierarchical network design. Route summarization reduces the routing protocol overhead on links in the network and reduces routing protocol processing within the routers.

Hierarchical Network Design

As Figure 4-1 illustrates, a hierarchical network design has three layers:

- The core layer provides optimal transport between sites.
- The distribution layer provides policy-based connectivity.
- The access layer provides workgroup/user access to the network.

Figure 4-1 *A Hierarchical Network Design Has Three Layers: Core, Distribution, and Access*

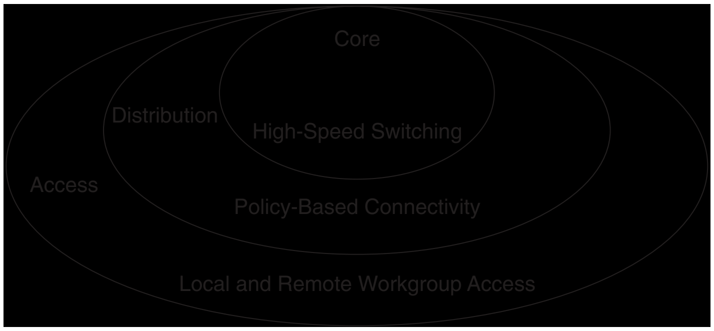

Each layer provides necessary functionality to the network. The layers do not need to be implemented as distinct physical entities. Each layer can be implemented in routers or switches, represented by a physical media, or combined in a single box. A particular layer can be omitted altogether, but for optimum performance, a hierarchy should be maintained.

Core Layer

The core layer is the high-speed switching backbone of the network, which is crucial to enable corporate communications. The core layer should have the following characteristics:

- Offer high reliability
- Provide redundancy
- Provide fault tolerance
- Adapt to changes quickly
- Offer low latency and good manageability
- Avoid slow packet manipulation caused by filters or other processes
- Have a limited and consistent diameter

NOTE When routers are used in a network, the number of router hops from edge to edge is called the *diameter*. As noted, it is considered good practice to design for a consistent diameter within a hierarchical network. This means that from any end station to another end station across the backbone, there should be the same number of hops. The distance from any end station to a server on the backbone should also be consistent.

Limiting the diameter of the internetwork provides predictable performance and ease of troubleshooting. Distribution layer routers and client LANs can be added to the hierarchical model without increasing the diameter because neither will affect how existing end stations communicate.

Distribution Layer

The distribution layer of the network is the demarcation point between the access and core layers of the network. The distribution layer can have many roles, including implementing the following functions:

- Policy (for example, to ensure that traffic sent from a particular network should be forwarded out one interface, while all other traffic should be forwarded out another interface)
- Security
- Address or area aggregation or summarization
- Departmental or workgroup access
- Broadcast/multicast domain definition
- Routing between virtual LANs (VLANs)
- Media translations (for example, between Ethernet and Token Ring)
- Redistribution between routing domains (for example, between two different routing protocols)
- Demarcation between static and dynamic routing protocols

Several Cisco IOS software features can be used to implement policy at the distribution layer, including the following:

- Filtering by source or destination address
- Filtering on input or output ports
- Hiding internal network numbers by route filtering
- Static routing

- Quality of service mechanisms (for example, to ensure that all devices along a path can accommodate the requested parameters)

Access Layer

The access layer provides user access to local segments on the network. The access layer is characterized by switched and shared bandwidth LANs in a campus environment. Microsegmentation, using LAN switches, provides high bandwidth to workgroups by dividing collision domains on Ethernet segments and reducing the number of stations capturing the token on Token Ring LANs.

For small office/home office (SOHO) environments, the access layer provides access for remote sites into the corporate network by using WAN technologies such as ISDN, Frame Relay, and leased lines. Features such as dial-on-demand routing (DDR) and static routing can be implemented to control costs.

Hierarchical Model Examples

For small- to medium-sized companies, the hierarchical model is often implemented as a hub-and-spoke topology, as shown in Figure 4-2. Corporate headquarters forms the hub and links to the remote offices form the spokes.

Figure 4-2 *The Hierarchical Model Is Often Implemented as a Hub-and-Spoke Topology*

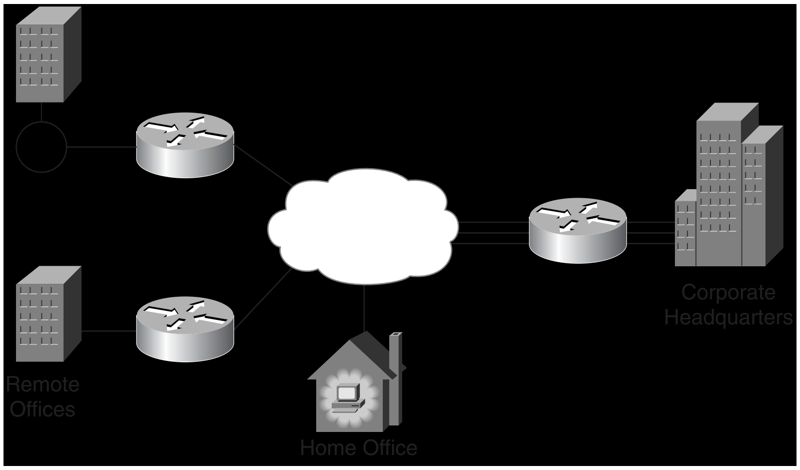

You can implement the hierarchical model by using either routers or switches. Figure 4-3 is an example of a switched hierarchical design, while Figure 4-4 shows examples of routed hierarchical designs.

Figure 4-3 *An Example of a Switched Hierarchical Design*

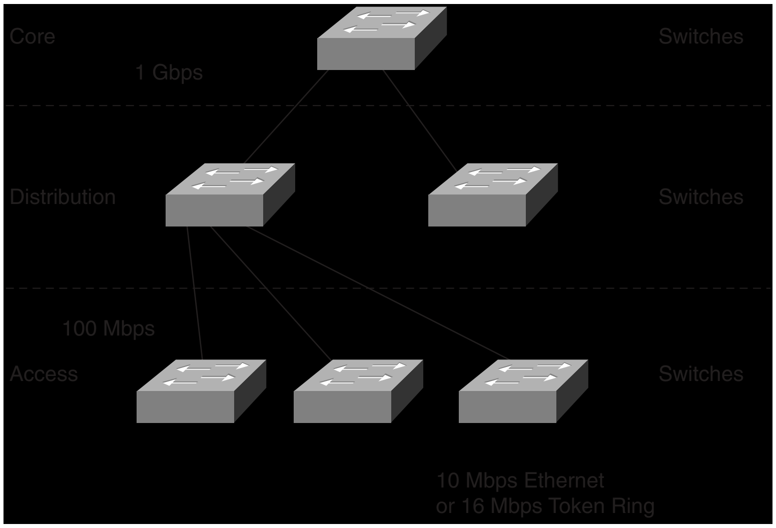

Figure 4-4 *Examples of Routed Hierarchical Designs*

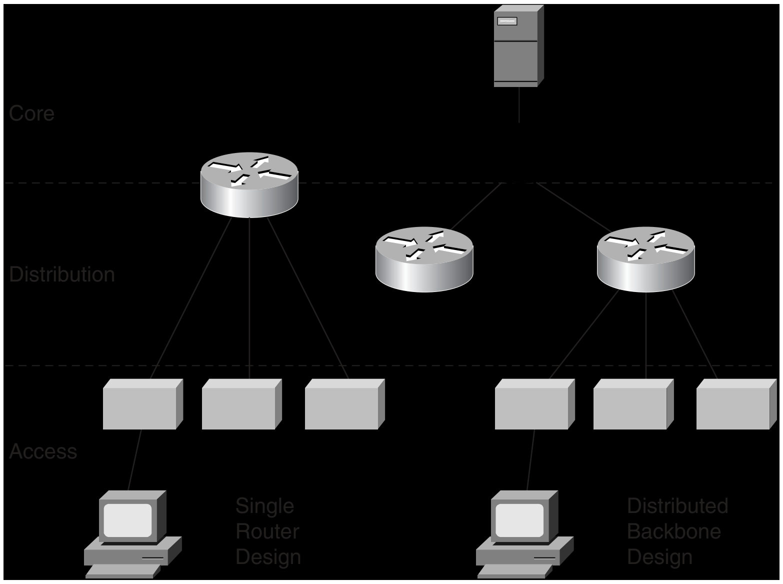

Redundant Models

When designing a network topology for a customer who has critical systems, services, or network paths, you should determine the likelihood that these components will fail and design redundancy where necessary.

Consider incorporating one of the following types of redundancy into your design:

- Workstation-to-router redundancy
- Server redundancy
- Route redundancy
- Media redundancy

Each of these types of redundancy is elaborated in the sections that follow.

Workstation-to-Router Redundancy

When a workstation has traffic to send to a station that is not local, the workstation has many possible ways to discover the address of a router on its network segment, including the following:

- Address Resolution Protocol (ARP)
- Explicit configuration
- Router Discovery Protocol (RDP)
- Routing Information Protocol (RIP)
- Internetwork Packet Exchange (IPX)
- AppleTalk
- Hot Standby Router Protocol (HSRP)

The sections that follow cover each of these methods.

ARP

Some IP workstations send an ARP frame to find a remote station. A router running proxy ARP can respond with its data link layer address. Cisco routers run proxy ARP by default.

Explicit Configuration

Most IP workstations must be configured with the IP address of a default router. This is sometimes called the *default gateway*.

In an IP environment, the most common method for a workstation to find a server is via explicit configuration (default router). If the workstation's default router becomes unavailable, the workstation must be reconfigured with the address of a different router. Some IP stacks enable you to configure multiple default routers, but many other IP stacks do not support redundant default routers.

RDP

RFC 1256 specifies an extension to the Internet Control Message Protocol (ICMP) that allows an IP workstation and router to run RDP to facilitate the workstation learning the address of a router.

RIP

An IP workstation can run RIP to learn about routers. RIP should be used in passive mode rather than active mode. (Active mode means that the station sends RIP frames every 30 seconds.) The Open Shortest Path First (OSPF) protocol also supports a workstation running RIP.

IPX

An IPX workstation broadcasts a find network number message to find a route to a server. A router then responds. If the client loses its connection to the server, it automatically sends the message again.

AppleTalk

An AppleTalk workstation remembers the address of the router that sent the last Routing Table Maintenance Protocol (RTMP) packet. As long as there are one or more routers on an AppleTalk workstation's network, it has a route to remote devices.

HSRP

Cisco's HSRP provides a way for IP workstations to keep communicating on the internetwork even if their default router becomes unavailable. HSRP works by creating a phantom router that has its own IP and MAC addresses. The workstations use this phantom router as their default router.

HSRP routers on a LAN communicate among themselves to designate two routers as active and standby. The active router sends periodic hello messages. The other HSRP routers listen for the hello messages. If the active router fails and the other HSRP routers stop receiving hello messages, the standby router takes over and becomes the active router. Because the new active router assumes both the IP and MAC addresses of the phantom, end nodes see no change at all.

They continue to send packets to the phantom router's MAC address, and the new active router delivers those packets.

HSRP also works for proxy ARP. When an active HSRP router receives an ARP request for a node that is not on the local LAN, the router replies with the phantom router's MAC address instead of its own. If the router that originally sent the ARP reply later loses its connection, the new active router can still deliver the traffic.

Figure 4-5 shows a sample implementation of HSRP.

Figure 4-5 *An Example of HSRP: The Phantom Router Represents the Real Routers*

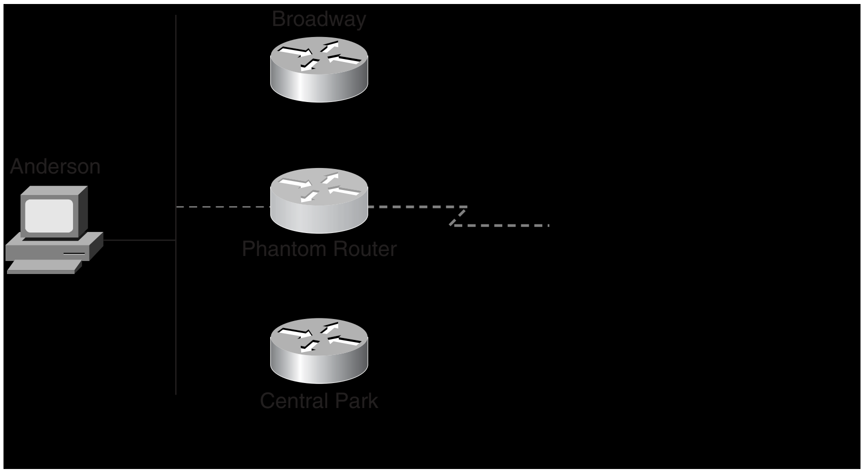

In Figure 4-5, the following sequence occurs:

1. The Anderson workstation is configured to use the Phantom router as its default router.
2. Upon booting, the routers elect Broadway as the HSRP active router. The active router does the work for the HSRP phantom. Central Park is the HSRP standby router.
3. When Anderson sends an ARP frame to find its default router, Broadway responds with the Phantom router's MAC address.
4. If Broadway goes off line, Central Park takes over as the active router, continuing the delivery of Anderson's packets. The change is transparent to Anderson. If a third HSRP router was on the LAN, that router would begin to act as the new standby router.

Server Redundancy

In some environments, fully redundant (mirrored) file servers should be recommended. For example, in a brokerage firm where traders must access data in order to buy and sell stocks, the data can be replicated on two or more redundant servers. The servers should be on different networks and power supplies.

If complete server redundancy is not feasible due to cost considerations, mirroring or duplexing of the file server hard drives is a good idea. *Mirroring* means synchronizing two disks, while *duplexing* is the same as mirroring with the additional feature that the two mirrored hard drives are controlled by different disk controllers.

Route Redundancy

Designing redundant routes has two purposes: load balancing and minimizing downtime.

Load Balancing

AppleTalk and IPX routers can remember only one route to a remote network by default, so they do not support load balancing. You can change this for IPX by using the **ipx maximum-paths** command and for AppleTalk by using the **appletalk maximum-paths** command on a Cisco router.

Most IP routing protocols can load balance across up to six parallel links that have equal cost. Use the **maximum-paths** command to change the number of links that the router will load balance over for IP; the default is four, the maximum is six. To support load balancing, keep the bandwidth consistent within a layer of the hierarchical model so that all paths have the same cost. (Cisco's IGRP and EIGRP are exceptions because they can load balance traffic across multiple routes that have different metrics by using a feature called *variance*.)

A hop-based routing protocol does load balancing over unequal bandwidth paths as long as the hop count is equal. After the slower link becomes saturated, the higher-capacity link cannot be filled; this is called *pinhole congestion*. Pinhole congestion can be avoided by designing equal bandwidth links within one layer of the hierarchy or by using a routing protocol that takes bandwidth into account.

IP load balancing depends on which switching mode is used on a router. Process switching load balances on a packet-by-packet basis. Fast, autonomous, silicon, optimum, distributed, and NetFlow switching load balance on a destination-by-destination basis because the processor caches the encapsulation to a specific destination for these types of switching modes.

Minimizing Downtime

In addition to facilitating load balancing, redundant routes minimize network downtime.

As already discussed, you should keep bandwidth consistent within a given layer of a hierarchy to facilitate load balancing. Another reason to keep bandwidth consistent within a layer of a hierarchy is that routing protocols converge much faster if multiple equal-cost paths to a destination network exist.

By using redundant, meshed network designs, you can minimize the effect of link failures. Depending on the convergence time of the routing protocols being used, a single link failure will not have a catastrophic effect.

A network can be designed as a full mesh or a partial mesh. In a full mesh network, every router has a link to every other router, as shown in Figure 4-6. A full mesh network provides complete redundancy and also provides good performance because there is just a single-hop delay between any two sites. The number of links in a full mesh is n(n–1)/2, where *n* is the number of routers. Each router is connected to every other router. (Divide the result by 2 to avoid counting Router X to Router Y and Router Y to Router X as two different links.)

Figure 4-6 *Full Mesh Network: Every Router Has a Link to Every Other Router in the Network*

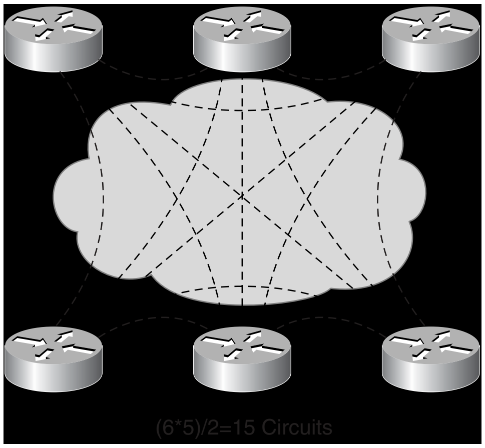

A full mesh network can be expensive to implement in wide-area networks due to the required number of links. In addition, practical limits to scaling exist for groups of routers that broadcast routing updates or service advertisements. As the number of router peers increases, the amount of bandwidth and CPU resources devoted to processing broadcasts increases.

A suggested guideline is to keep broadcast traffic at less than 20 percent of the bandwidth of each link; this will limit the number of peer routers that can exchange routing tables or service advertisements. When planning redundancy, follow guidelines for simple, hierarchical design. Figure 4-7 illustrates a classic hierarchical and redundant enterprise design that uses a partial mesh rather than a full mesh architecture. For LAN designs, links between the access and distribution layer can be Fast Ethernet, with links to the core at Gigabit Ethernet speeds.

Figure 4-7 *Partial Mesh Design with Redundancy*

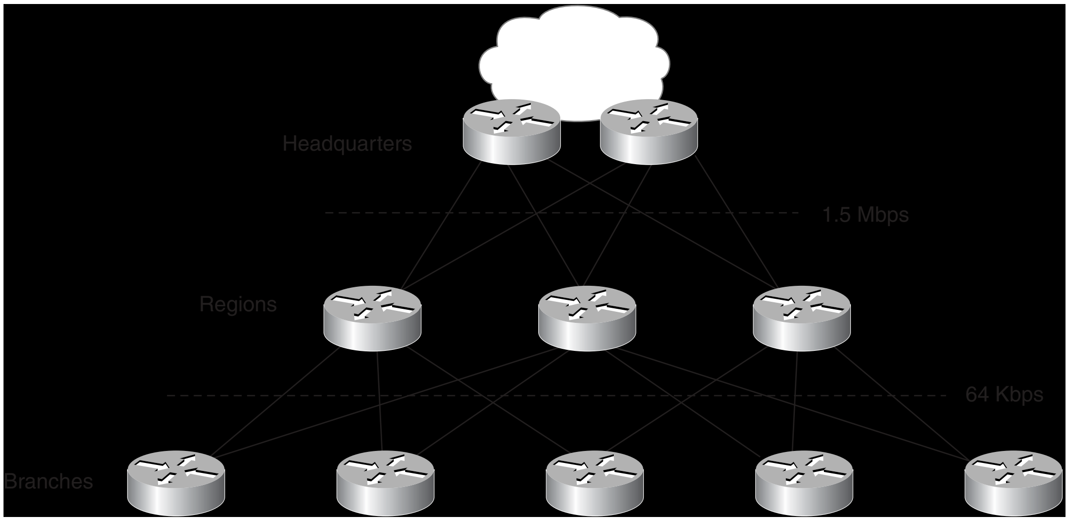

Media Redundancy

In mission-critical applications, it is often necessary to provide redundant media.

In switched networks, switches can have redundant links to each other. This redundancy is good because it minimizes downtime, but it may result in broadcasts continuously circling the network, which is called a *broadcast storm*. Because Cisco switches implement the IEEE 802.1d Spanning-Tree Algorithm, this looping can be avoided in the Spanning-Tree Protocol. The Spanning-Tree Algorithm guarantees that only one path is active between two network stations. The algorithm permits redundant paths that are automatically activated when the active path experiences problems.

Because WAN links are often critical pieces of the internetwork, redundant media is often deployed in WAN environments. As shown in Figure 4-8, backup links can be provisioned so they become active when a primary link goes down or becomes congested.

Figure 4-8 *Backup Links Can Be Used to Provide Redundancy*

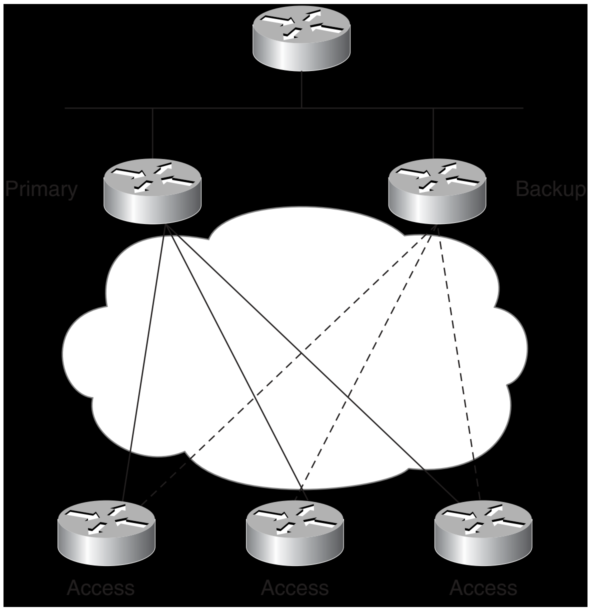

Often, backup links use a different technology. For example, a leased line can be in parallel with a backup dialup line or ISDN circuit. By using *floating static routes*, you can specify that the backup route has a higher administrative distance (used by Cisco routers to select which routing information to use) so that it is not normally used unless the primary route goes down.

NOTE When provisioning backup links, learn as much as possible about the actual physical circuit routing. Different carriers sometimes use the same facilities, meaning that your backup path is susceptible to the same failures as your primary path. You should do some investigative work to ensure that your backup really is acting as a backup.

Backup links can be combined with load balancing and channel aggregation. *Channel aggregation* means that a router can bring up multiple channels (for example, Integrated Services Digital Network [ISDN] B channels) as bandwidth requirements increase.

Cisco supports the Multilink Point-to-Point Protocol (MPPP), which is an Internet Engineering Task Force (IETF) standard for ISDN B channel (or asynchronous serial interface) aggregation. MPPP does not specify how a router should accomplish the decision-making process to bring up extra channels. Instead, it seeks to ensure that packets arrive in sequence at the receiving router. Then, the data is encapsulated within PPP and the datagram is given a sequence number.

At the receiving router, PPP uses this sequence number to re-create the original data stream. Multiple channels appear as one logical link to upper-layer protocols.

Secure Models

This section introduces secure topology models. The information in this book is not sufficient to learn all the nuances of internetwork security. To learn more about internetwork security, you might want to read the book *Firewalls and Internet Security*, by Bill Cheswick and Steve Bellovin, published by Addison Wesley. Also, by searching for the word “security” on Cisco’s web site (www.cisco.com), you can keep up to date on security issues.

Secure topologies are often designed by using a firewall. A firewall protects one network from another untrusted network. This protection can be accomplished in many ways, but in principle, a firewall is a pair of mechanisms: One blocks traffic and the other permits traffic.

Some firewalls place a greater emphasis on blocking traffic, and others emphasize permitting traffic. Figure 4-9 shows a simple firewall topology using routers.

Figure 4-9 *A Simple Firewall Network, Using Routers*

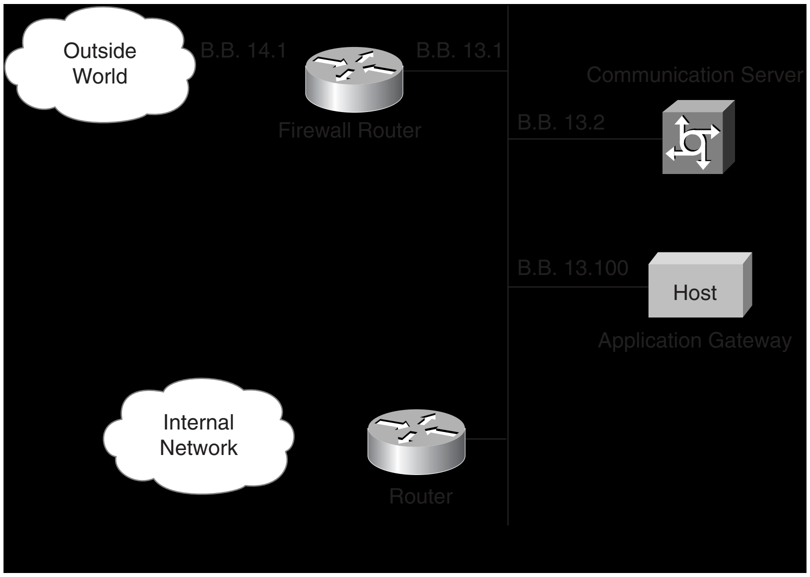

You can design a firewall system using packet-filtering routers and bastion hosts. A *bastion host* is a secure host that supports a limited number of applications for use by outsiders. It holds data that outsiders access (for example, web pages) but is strongly protected from outsiders using it for anything other than its limited purposes.

Three-Part Firewall System

The classic firewall system, called the *three-part firewall system*, has the following three specialized layers, as shown in Figure 4-10:

- An isolation LAN that is a buffer between the corporate internetwork and the outside world. (The isolation LAN is called the demilitarized zone (DMZ) in some literature.)
- A router that acts as an inside packet filter between the corporate internetwork and the isolation LAN.
- Another router that acts as an outside packet filter between the isolation LAN and the outside internetwork.

Figure 4-10 *Structure and Components of a Three-Part Firewall System*

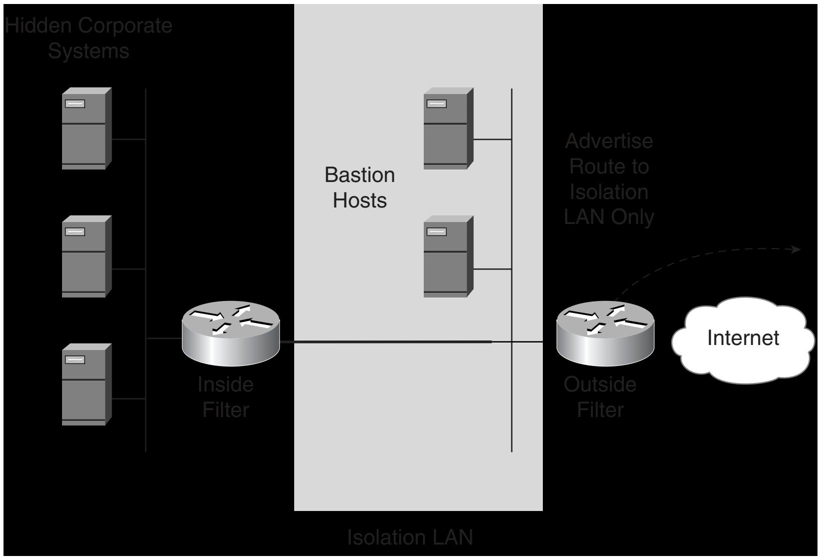

Services available to the outside world are located on bastion hosts in the isolation LAN. Example services in these hosts include:

- Anonymous FTP server
- Web server
- Domain Name System (DNS)
- Telnet
- Specialized security software such as Terminal Access Controller Access Control System (TACACS)

The isolation LAN has a unique network number that is different than the corporate network number. Only the isolation LAN network is visible to the outside world. On the outside filter, you should advertise only the route to the isolation LAN.

If internal users need to get access to Internet services, allow TCP outbound traffic from the internal corporate internetwork. Allow TCP packets back into the internal network only if they are in response to a previously sent request. All other TCP traffic should be blocked because new inbound TCP sessions could be from hackers trying to establish sessions with internal hosts.

NOTE To determine whether TCP traffic is a response to a previously sent request or a request for a new session, the router examines some bits in the code field of the TCP header. If the acknowledgement field (ACK) is valid or reset the connection (RST) bits are set in a TCP segment header, the segment is a response to a previously sent request. The established keyword in Cisco IOS access lists (filters) is used to indicate packets with ACK or RST bits set.

The following list summarizes some *rules* for the three-part firewall system:

- The inside packet filter router should allow inbound TCP packets from established sessions.
- The outside packet filter router should allow inbound TCP packets from established TCP sessions.
- The outside packet filter router should also allow packets to specific TCP or UDP ports going to specific bastion hosts (including TCP SYN packets that are used to establish a session).

Always block traffic from coming in from between the firewall routers and hosts and the internal network. The firewall routers and hosts themselves are likely to be a jumping-off point for hackers, as shown in Figure 4-11.

Figure 4-11 *Firewall Routers and Hosts May Make Your Network Vulnerable to Hacker Attacks*

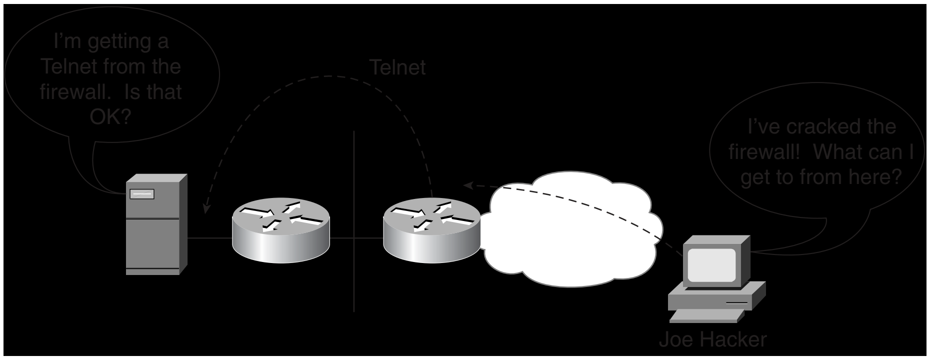
I'm getting a Telnet from the firewall. Is that OK?
Telnet
I've cracked the firewall! What can I get to from here?
Joe Hacker

Keep bastion hosts and firewall routers simple. They should run as few programs as possible. The programs should be simple because simple programs have fewer bugs than complex programs. Bugs introduce possible security holes.

Do not enable any unnecessary services or connections on the outside filter router. A list of suggestions for implementing the outside filter router follows:

- Turn off Telnet access (no virtual terminals defined).
- Use static routing only.
- Do not make it a TFTP server.
- Use password encryption.
- Turn off proxy ARP service.
- Turn off finger service.
- Turn off IP redirects.
- Turn off IP route caching.
- Do not make the router a MacIP server (MacIP provides connectivity for IP over AppleTalk by tunneling IP datagrams inside AppleTalk).

Cisco PIX Firewall

To provide stalwart security, hardware firewall devices can be used in addition to or instead of packet-filtering routers. For example, in the three-part firewall system illustrated earlier in Figure 4-10, a hardware firewall device could be installed on the isolation LAN. A hardware firewall device offers the following benefits:

- Less complex and more robust than packet filters

- No required downtime for installation
- No required upgrading of hosts or routers
- No necessary day-to-day management

Cisco's PIX Firewall is a hardware device that offers the features in the preceding list, as well as full outbound Internet access from unregistered internal hosts. IP addresses can be assigned from the private ranges, as defined in RFC 1918 (available at http://info.internet.isi.edu/in-notes/rfc/files/rfc1918.txt). The PIX Firewall uses a protection scheme called *Network Address Translation (NAT)*, which allows internal users access to the Internet while protecting internal networks from unauthorized access.

Further details on the PIX Firewall are available on Cisco's web site at www.cisco.com/warp/public/cc/cisco/mkt/security/pix/.

The PIX Firewall provides firewall security without the administrative overhead and risks associated with UNIX-based or router-based firewall systems. The PIX Firewall operates on a secure real-time kernel, not on UNIX. The network administrator is provided with complete auditing of all transactions, including attempted break-ins.

The PIX Firewall supports data encryption with the Cisco PIX Private Link, a card that provides secure communication between multiple PIX systems over the Internet using the data encryption standard (DES).

The PIX Firewall provides TCP and UDP connectivity from internal networks to the outside world by using a scheme called *adaptive security.* All inbound traffic is verified for correctness against the following connection state information:

- Source and destination IP addresses
- Source and destination port numbers
- Protocols
- TCP sequence numbers (which are randomized to eliminate the possibility of hackers guessing numbers)

LAN Types

The CCDA objective covered in this section is as follows:

14	Draw a topology map that meets the customer's needs and includes a high-level view of internetworking devices and interconnecting media.

Local-area networks can be classified as a large building LAN, campus LAN, or small/remote LAN. The large building LAN contains the major data center with high-speed access and floor communications closets; the large building LAN is usually the headquarters in larger companies. Campus LANs provide connectivity between buildings on a campus; redundancy is usually a requirement. Small/remote LANs provide connectivity to remote offices with a small number of nodes.

It is important to remember the Cisco hierarchical approach of network design. First, build a high-speed core backbone network. Second, build the distribution layer, where policy can be applied. Finally, build the access layer, where LANs provide access to the network end stations.

Large Building LANs

Large building LANs are segmented by floors or departments. Company mainframes and servers reside in a computing center. Media lines run from the computer center to the wiring closets at the various segments. From the wiring closets, media lines run to the offices and cubicles around the work areas. Figure 4-12 depicts a typical large building design.

Figure 4-12 *Large Building LAN Design*

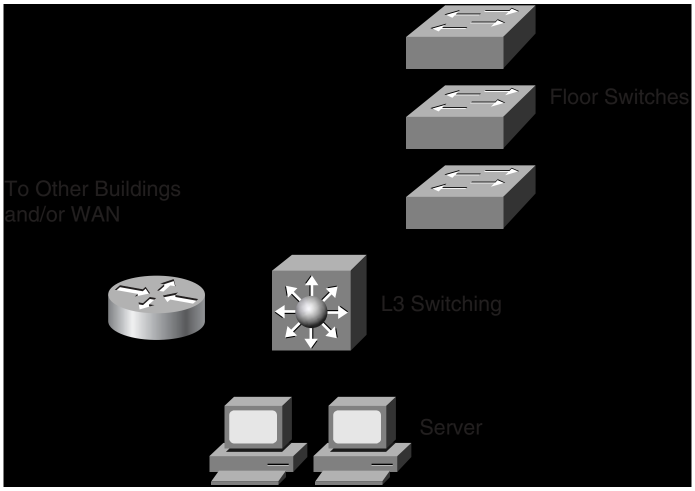

Each floor may have more than 200 users. Following a hierarchical model of access, distribution, and core, Ethernet and Fast Ethernet nodes may connect to hubs and switches in the communications closet. Uplink ports from closet switches connect back to one or two (for redundancy) distribution switches. Distribution switches may provide connectivity to server farms that provide business applications, DHCP, DNS, intranet, and other services.

Campus LANs

A campus LAN connects two or more buildings located near each other using high-bandwidth LAN media. Usually the media (for example, copper or fiber) is owned. High-speed switching devices are recommended to minimize latency. In today's networks, Gigabit Ethernet campus backbones are the standard for new installations. In Figure 4-13, campus buildings are connected by using Layer 3 switches with Gigabit Ethernet media.

Figure 4-13 *Campus LANs*

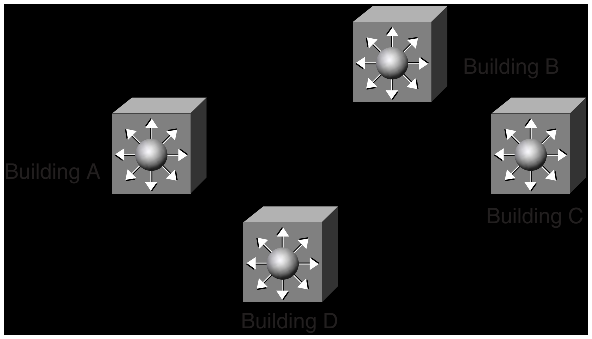

Ensure that a hierarchical design is implemented on the campus LAN and that network layer addressing is assigned to control broadcasts on the networks. Each building should have addressing assigned in such a way as to maximize address summarization. Apply contiguous subnets to buildings at the bit boundary to apply summarization and ease the design. Campus networks can support high-bandwidth applications such as video conferencing. Although most WAN implementations are configured to support only IP, legacy LANs may still be configured to support IPX and AppleTalk.

Small/Remote Site LANs

Small/remote sites usually connect back to the corporate network via a small router (Cisco 2500). The local-area network service is provided by a small hub or LAN switch (Catalyst 1900). The router filters broadcasts to the WAN circuit and forwards packets that require services from the corporate network. A server may be placed at the small/remote site to provide DHCP and other local applications such as NT backup domain controller and DNS; if not, the router will need to be configured to forward DHCP broadcasts and other types of services. Figure 4-14 shows a typical architecture of a small or remote LAN. *Building Cisco Remote Access Networks* from Cisco Press is an excellent resource for more information on remote access.

Figure 4-14 *Small/Remote Office LAN*

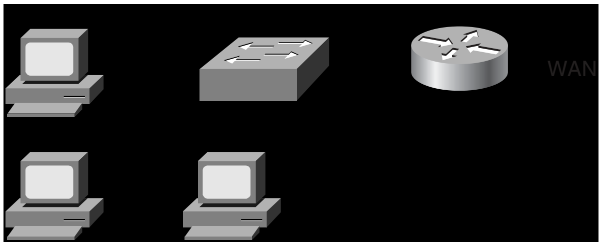

LAN Media

The CCDA objectives covered in this section are as follows:

15 Recognize scalability constraints and issues for standard LAN technologies.

16 Recommend Cisco products and LAN technologies that will meet a customer's requirements for performance, capacity, and scalability in small- to medium-sized networks.

This section identifies some of the constraints that should be considered when provisioning various LAN media types. For additional reference material on this subject, refer to Appendix D, "LAN Media Reference."

Ethernet Design Rules

Table 4-1 provides scalability information that you can use when provisioning IEEE 802.3 networks.

Table 4-1 *Scalability Constraints for IEEE 802.3*

	10Base5	10Base2	10BaseT	100BaseT
Topology	Bus	Bus	Star	Star
Maximum Segment Length (meters)	500	185	100 from hub to station	100 from hub to station

Table 4-1 *Scalability Constraints for IEEE 802.3 (Continued)*

	10Base5	10Base2	10BaseT	100BaseT
Maximum Number of Attachments per Segment	100	30	2 (hub and station or hub-hub)	2 (hub and station or hub-hub)
Maximum Collision Domain	2500 meters of 5 segments and 4 repeaters; only 3 segments can be populated	2500 meters of 5 segments and 4 repeaters; only 3 segments can be populated	2500 meters of 5 segments and 4 repeaters; only 3 segments can be populated	See the details in the section "100 Mbps Fast Ethernet Design Rules" later in this chapter.

The most significant design rule for Ethernet is that the round-trip propagation delay in one collision domain must not exceed 512 bit times, which is a requirement for collision detection to work correctly. This rule means that the maximum round-trip delay for a 10 Mbps Ethernet network is 51.2 microseconds. The maximum round-trip delay for a 100 Mbps Ethernet network is only 5.12 microseconds because the bit time on a 100 Mbps Ethernet network is 0.01 microseconds as opposed to 0.1 microseconds on a 10 Mbps Ethernet network.

To make 100 Mbps Ethernet work, distance limitations are much more severe than those required for 10 Mbps Ethernet. The general rule is that a 100 Mbps Ethernet has a maximum diameter of 205 meters when unshielded twisted-pair (UTP) cabling is used, whereas 10 Mbps Ethernet has a maximum diameter of 500 meters with 10BaseT and 2500 meters with 10Base5.

10 Mbps Fiber Ethernet Design Rules

Table 4-2 provides some guidelines to help you choose the right media for your network designs. 10BaseF is based on the fiber-optic interrepeater link (FOIRL) specification, which includes 10BaseFP, 10BaseFB, 10BaseFL, and a revised FOIRL standard. The new FOIRL allows data terminal equipment (DTE) end-node connections rather than just repeaters, which were allowed with the older FOIRL specification.

Table 4-2 *Scalability Constraints for 10 Mbps Fiber Ethernet*

	10BaseFP	10BaseFB	10BaseFL	Old FOIRL	New FOIRL
Topology	Passive star	Backbone or repeater fiber system	Link	Link	Link or star
Allows DTE (End Node) Connections?	Yes	No	No	No	Yes

continues

Table 4-2 *Scalability Constraints for 10 Mbps Fiber Ethernet (Continued)*

	10BaseFP	10BaseFB	10BaseFL	Old FOIRL	New FOIRL
Maximum Segment Length (Meters)	500	2000	1000 or 2000	1000	1000
Allows Cascaded Repeaters?	No	Yes	No	No	Yes
Maximum Collision Domains in Meters	2500	2500	2500	2500	2500

100 Mbps Fast Ethernet Design Rules

100 Mbps Ethernet, or Fast Ethernet, topologies present some distinct constraints on the network design because of their speed. The combined latency due to cable lengths and repeaters must conform to the specifications in order for the network to work properly. This section discusses these issues and provides example calculations.

Understanding Collision Domains

The overriding design rule for 100 Mbps Ethernet networks is that the round-trip collision delay must not exceed 512 bit times. However, the bit time on a 100 Mbps Ethernet network is 0.01 microseconds, as opposed to 0.1 microseconds on a 10 Mbps Ethernet network. Therefore, the maximum round-trip delay for a 100 Mbps Ethernet network is 5.12 microseconds, as opposed to the more lenient 51.2 microseconds in a 10 Mbps Ethernet network.

100BaseT Repeaters

For a 100 Mbps Ethernet to work, you must impose distance limitations based on the type of repeaters used.

The IEEE 100BaseT specification defines two types of repeaters: Class I and Class II. Class I repeaters have a latency (delay) of 0.7 microseconds or less. Only one repeater hop is allowed. Class II repeaters have a latency (delay) of 0.46 microseconds or less. One or two repeater hops are allowed.

Table 4-3 shows the maximum size of collision domains, depending on the type of repeater.

Table 4-3 *Maximum Size of Collision Domains for 100BaseT*

	Copper	Mixed Copper and Multimode Fiber	Multimode Fiber
DTE-DTE (or Switch-Switch)	100 meters		412 meters (2000 if full duplex)
One Class I Repeater	200 meters	260 meters	272 meters
One Class II Repeater	200 meters	308 meters	320 meters
Two Class II Repeaters	205 meters	216 meters	228 meters

The Cisco FastHub 316 is a Class II repeater, as are all the Cisco FastHub 300 series hubs. These hubs actually exceed the Class II specifications, which means that they have even lower latencies and therefore allow longer cable lengths. For example, with two FastHub 300 repeaters and copper cable, the maximum collision domain is 223 meters.

Example of 100BaseT Topology

Figure 4-15 shows examples of 100BaseT topologies with different media.

Figure 4-15 *Examples of 100BaseT Topologies with Various Media and Repeaters*

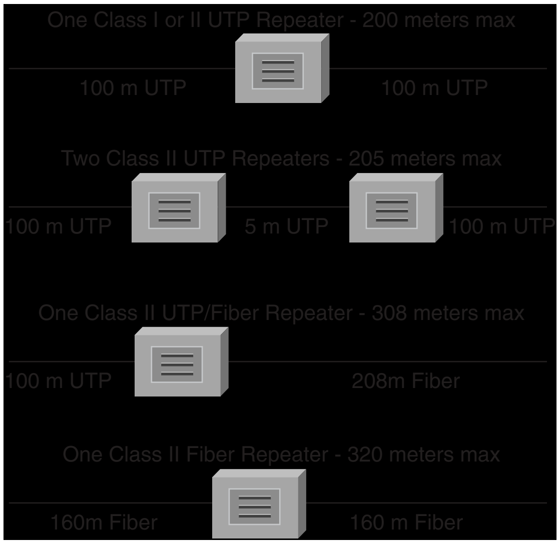
One Class I or II UTP Repeater - 200 meters max
100 m UTP
100 m UTP
Two Class II UTP Repeaters - 205 meters max
100 m UTP
5 m UTP
100 m UTP
One Class II UTP/Fiber Repeater - 308 meters max
100 m UTP
208m Fiber
One Class II Fiber Repeater - 320 meters max
160m Fiber
160 m Fiber

Other topologies are possible as long as the round-trip propagation delay does not exceed 5.12 microseconds (512 bit times). When the delay does exceed 5.12 microseconds, the network experiences illegal (late) collisions and CRC errors.

Checking the Propagation Delay

To determine whether configurations other than the standard ones shown in Figure 4-15 will work, use the following information from the IEEE 802.3u specification.

To check a path to make sure the path delay value (PDV) does not exceed 512 bit times, add up the following delays:

- All link segment delays
- All repeater delays
- DTE delay
- A safety margin (0 to 5 bit times)

Use the following steps to calculate the PDV:

1. Determine the delay for each link segment; this is the link segment delay value (LSDV), including interrepeater links, using the following formula. (Multiply by two so it is a round-trip delay.)

   LSDV = 2 × segment length × cable delay for this segment.

   For end-node segments, the segment length is the cable length between the physical interface at the repeater and the physical interface at the DTE. Use your two farthest DTEs for a worst-case calculation. For interrepeater links, the segment length is the cable length between the repeater physical interfaces.

   Cable delay is the delay specified by the manufacturer if available. When actual cable lengths or propagation delays are not known, use the delay in bit times as specified in Table 4-4.

   Cable delay must be specified in bit times per meter (BT/m).

2. Add together the LSDVs for all segments in the path.

3. Determine the delay for each repeater in the path. If model-specific data is not available from the manufacturer, determine the class of repeater (I or II).

4. MII cables for 100BaseT should not exceed 0.5 meters each in length. When evaluating system topology, MII cable lengths need not be accounted for separately. Delays attributed to the MII are incorporated into DTE and repeater delays.

5. Use the DTE delay value shown in Table 4-4 unless your equipment manufacturer defines a different value.

6 Decide on an appropriate safety margin from 0 to 5 bit times. Five bit times is a safe value.

7 Insert the values obtained from the preceding calculations into the formula for calculating the PDV:

PDV = link delays + repeater delays + DTE delay + safety margin

8 If the PDV is less than 512, the path is qualified in terms of worst-case delay.

Round-Trip Delay

Table 4-4 shows round-trip delay in bit times for standard cables and maximum round-trip delay in bit times for DTEs, repeaters, and maximum-length cables.

NOTE Note that the values shown in Table 4-4 have been multiplied by two to provide a round-trip delay. If you use these numbers, you need not multiply by two again in the LSDV formula (LSDV = 2 × segment length × cable delay for this segment).

Table 4-4 *Network Component Delays*

Component	Round-Trip Delay in Bit Times per Meter	Maximum Round-Trip Delay in Bit Times
Two TX/FX DTEs	N/A	100
Two T4 DTEs	N/A	138
One T4 DTE and one TX/FX DTE	N/A	127
Category 3 cable segment	1.14	114 (100 meters)
Category 4 cable segment	1.14	114 (100 meters)
Category 5 cable segment	1.112	111.2 (100 meters)
STP cable segment	1.112	111.2 (100 meters)
Fiber-optic cable segment	1.0	412 (412 meters)
Class I repeater	N/A	140
Class II repeater with all ports TX or FX	N/A	92
Class II repeater with any port T4	N/A	67

Source: IEEE 802.3u—1995, “Media Access Control (MAC) Parameters, Physical Layer, Medium Attachment Units, and Repeater for 100 Mb/s Operation, Type 100BASE-T.”

Example Network Cabling Implementation

See Figure 4-16 for this example. Company ABC has all UTP Category 5 cabling. Two Class II repeaters are separated by 20 meters instead of the standard 5 meters. The network administrators are trying to determine whether this configuration will work.

Figure 4-16 *An Example Network Cabling Implementation for Company ABC (Showing the Two Most Distant DTEs)*

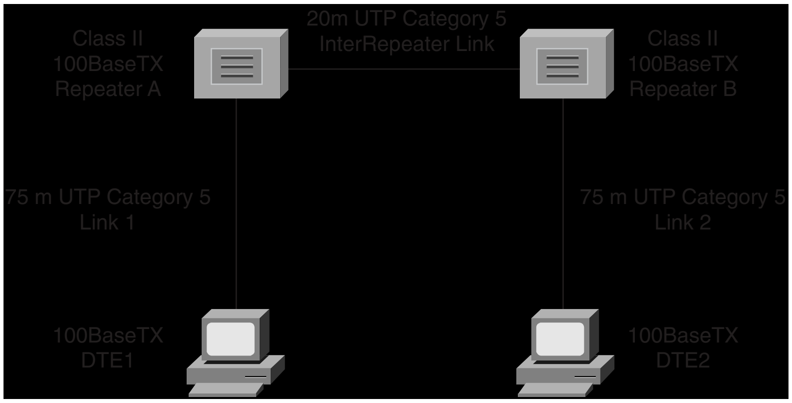

To ensure that the PDV does not exceed 512 bit times, the network administrators must calculate a worst-case scenario using DTE 1 and DTE 2, which are 75 meters from their repeaters.

Assume that DTE 1 starts transmitting a minimum-sized frame of 64 bytes (512 bits). DTE 2 just barely misses hearing DTE 1's transmission and starts transmitting also. The collision happens on the far-right side of the network and must traverse back to DTE 1. These events must occur within 512 bit times. If they take any longer than 512 bit times, then DTE 1 will have stopped sending when it learns about the collision and will not know that its frame was damaged by the collision. To calculate the link delays for the Category 5 cable segments, the repeaters, and DTEs, the administrators use the values from Table 4-4. (Remember that Table 4-4 uses round-trip delay values, so you need not multiply by two.)

To test whether this network will work, the network administrators filled in Table 4-5.

Table 4-5 *Delays of Components in Company ABC's Network*

Delay Cause	Calculation of Network Component Delay	Total (Bit Times)
Link 1	75m × 1.112 bit times/m	83.4
Link 2	75m × 1.112 bit times/m	83.4
Interrepeater link	20m × 1.112 bit times/m	22.24
Repeater A	92 bit times	92

Table 4-5 *Delays of Components in Company ABC's Network (Continued)*

Delay Cause	Calculation of Network Component Delay	Total (Bit Times)
Repeater B	92 bit times	92
DTE 1 and 2	100 bit times	100
Safety margin	5 bit times	5
Grand Total	**Add Individual Totals**	**478.04**

The grand total in Table 4-5 is fewer than 512 bit times, so this network will work.

Calculating Cable Delays

Some cable manufacturers specify propagation delays relative to the speed of light or in nanoseconds per meter (ns/m). To convert these values to bit times per meter (BT/m), use Table 4-6.

Table 4-6 *Conversion to Bit Times per Meter for Cable Delays*

Speed Relative to Speed of Light	Nanoseconds per Meter (ns/m)	Bit Times per Meter (BT/m)
0.4	8.34	0.834
0.5	6.67	0.667
0.51	6.54	0.654
0.52	6.41	0.641
0.53	6.29	0.629
0.54	6.18	0.618
0.55	6.06	0.606
0.56	5.96	0.596
0.57	5.85	0.585
0.58	5.75	0.575
0.5852	5.70	0.570
0.59	5.65	0.565
0.6	5.56	0.556
0.61	5.47	0.547
0.62	5.38	0.538

continues

Table 4-6 *Conversion to Bit Times per Meter for Cable Delays (Continued)*

Speed Relative to Speed of Light	Nanoseconds per Meter (ns/m)	Bit Times per Meter (BT/m)
0.63	5.29	0.529
0.64	5.21	0.521
0.65	5.13	0.513
0.654	5.10	0.510
0.66	5.05	0.505
0.666	5.01	0.501
0.67	4.98	0.498
0.68	4.91	0.491
0.69	4.83	0.483
0.7	4.77	0.477
0.8	4.17	0.417
0.9	3.71	0.371

Source: IEEE 802.3u — 1995, “Media Access Control (MAC) Parameters, Physical Layer, Medium Attachment Units, and Repeater for 100 Mb/s Operation, Type 100BASE-T.”

Token Ring Design Rules

Table 4-7 lists some scalability concerns when designing Token Ring segments. Refer to IBM’s Token Ring planning guides for more information on the maximum segment sizes and maximum diameter of a network.

Table 4-7 *Scalability Constraints for Token Ring*

	IBM Token Ring	IEEE 802.5
Topology	Star	Not specified
Maximum Segment Length (Meters)	Depends on type of cable, number of MAUs, and so on	Depends on type of cable, number of MAUs, and so on
Maximum Number of Attachments per Segment	260 for STP, 72 for UTP	250
Maximum Network Diameter	Depends on type of cable, number of MAUs, and so on	Depends on type of cable, number of MAUs, and so on

CH01.book Page 131 Friday, January 14, 2000 2:07 PM

Gigabit Ethernet Design Rules

The most recent development in the Ethernet arena is Gigabit Ethernet. Gigabit Ethernet is specified by two standards: IEEE 802.3z and 802.3ab. The 802.3z standard specifies the operation of Gigabit Ethernet over fiber and coaxial cable and introduces the Gigabit Media Independent Interface (GMII). The 802.3z standard was approved in June 1998.

The 802.3ab standard specifies the operation of Gigabit Ethernet over Category 5 UTP. Gigabit Ethernet still retains the frame formats and frame sizes and it still uses CSMA/CD. As with Ethernet and Fast Ethernet, full duplex operation is possible. Differences can be found in the encoding; Gigabit Ethernet uses 8B/10B coding with simple nonreturn to zero (NRZ). Because of the 20 percent overhead, pulses run at 1250 MHz to achieve a 1000 Mbps. Table 4-8 covers Gigabit Ethernet scalability constraints.

Table 4-8 *Gigabit Ethernet Scalability Constraints*

Type	Speed	Maximum segment length	Encoding	Media
1000BaseT	1000 Mbps	100m	5-level	Cat 5 UTP
1000BaseLX (long wave)	1000 Mbps	550m	8B/10B	Single/multiple mode fiber
1000BaseSX (short wave)	1000 Mbps	62.5 micrometers: 220m 50 micrometers: 500m	8B/10B	Multimode fiber
1000BaseCX	1000 Mbps	25m	8B/10B	Shielded balanced copper

FDDI Design Rules

The FDDI specification does not actually specify the maximum segment length or network diameter. It specifies the amount of allowed power loss, which works out to the approximate distances shown in Table 4-9.

Table 4-9 *Scalability Constraints for FDDI*

	Multimode Fiber	Single-Mode Fiber	UTP
Topology	Dual ring, tree of concentrators, and others	Dual ring, tree of concentrators, and others	Star
Maximum Segment Length	2km between stations	60km between stations	100m from hub to station

continues

Table 4-9 *Scalability Constraints for FDDI (Continued)*

	Multimode Fiber	Single-Mode Fiber	UTP
Maximum Number of Attachments per Segment	1000 (500 dual-attached stations)	1000 (500 dual-attached stations)	2 (hub and station or hub-hub)
Maximum Network Diameter	200km	200km	200km

LAN Hardware

The CCDA objectives covered in this section are as follows:

13	Describe the advantages, disadvantages, scalability issues, and applicability of standard internetwork topologies.
15	Recognize scalability constraints and issues for standard LAN technologies.

This section covers the following hardware technologies as they can be applied to LAN design:

- Repeaters
- Hubs
- Bridges
- Switches
- Routers
- Layer 3 switches
- Combining hubs, switches, and routers

Repeaters

Repeaters are the basic unit used in networks to connect separate segments. Repeaters take incoming frames, regenerate the preamble, amplify the signals, and send the frame out all other interfaces. Repeaters operate in the physical layer of the OSI model. Because repeaters are not aware of packets or frame formats, they do not control broadcasts or collision domains. Repeaters are said to be protocol transparent because they are not aware of upper-layer protocols such as IP, IPX, and so on.

One basic rule of using repeaters is the 5-4-3 Rule. The maximum path between two stations on the network should not be more than 5 segments with 4 repeaters between those segments and no more than 3 populated segments. Repeaters introduce a small amount of latency, or delay, when propagating the frames. A transmitting device must be able to detect a collision with another device within the specified time after the delay introduced by the cable segments and repeaters is factored in. The 512 bit-time specification also governs segment lengths. A more detailed explanation of the specification can be found at www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ethernet.htm. Figure 4-17 illustrates an example of the 5-4-3 Rule.

Figure 4-17 *Repeater 5-4-3 Rule*

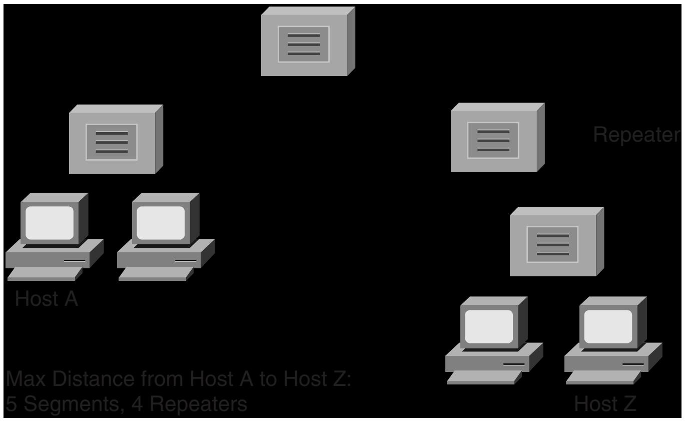

Hubs

With the increasing density of LANs in the late 80s and early 90s, hubs were introduced to concentrate Thinnet and 10BaseT networks in the wiring closet. Traditional hubs operate on the physical layer of the OSI model and perform the same functions as basic repeaters.

Bridges

Bridges are used to connect separate segments of a network. They differ from repeaters in that bridges are intelligent devices that operate in the data link layer of the OSI model. Bridges control the collision domains on the network. Bridges also learn the MAC layer addresses of each node on each segment and on which interface they are located. For any incoming frame, bridges forward the frame only if the destination MAC address is on another port or if the bridge is not aware of its location. The latter is called *flooding*. Bridges filter any incoming frames with destination MAC addresses that are on the same segment from where the frame arrives; they do not forward the frame on.

Bridges are store and forward devices. They store the entire frame and verify the CRC before forwarding. If a CRC error is detected, the frame is discarded. Bridges are protocol transparent;

they are not aware of the upper-layer protocols like IP, IPX, and AppleTalk. Bridges are designed to flood all unknown and broadcast traffic.

Bridges implement the Spanning-Tree Protocol to build a loop free network topology. Bridges communicate with each other, exchanging information such as priority and bridge interface MAC addresses. They select a root bridge and then implement the Spanning-Tree Protocol. Some interfaces are placed in a hold state, while other bridges will have interfaces in forwarding mode. Looking at Figure 4-18, note that there is no load sharing or dual paths with bridge protocols as there is in routing.

Figure 4-18 *Spanning-Tree Protocol*

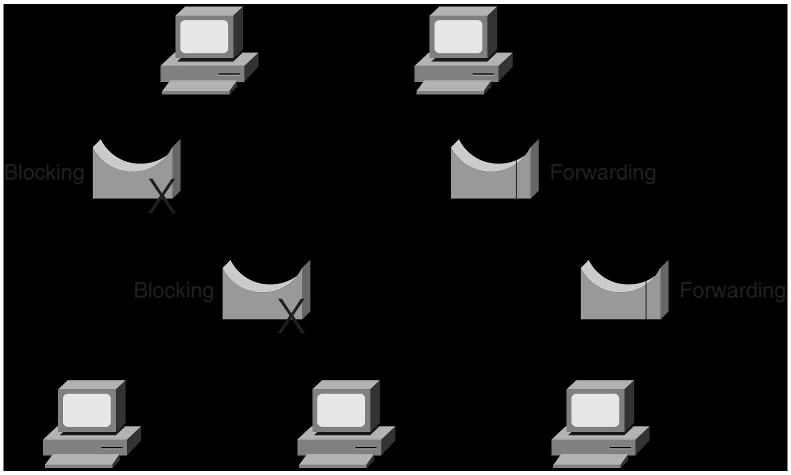

Switches

Switches are the evolution of bridges. Switches use fast integrated circuits that reduce the latency that bridges introduce to the network. Switches also enable the capability to run in cut-through mode. In cut-through mode, the switch does not wait for the entire frame to enter its buffer; instead, it forwards the frame after it has read the destination MAC address field of the frame. Cut-through operation increases the probability that error frames are propagated on the network, which increases CRC and runt frames on the network. Because of these problems, most switches today perform store-and-forward operation with CRC check as bridges do. Figure 4-19 shows a switch; note that it controls collision domains but not broadcast domains.

Figure 4-19 *Switches Control Collision Domains*

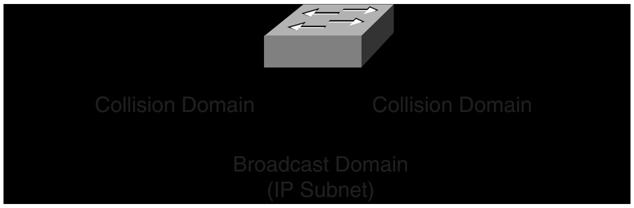

Switches have characteristics similar to bridges; however, they have more ports and run faster. Switches keep a table of MAC addresses per port, and they implement Spanning-Tree Protocol. Switches also operate in the data link layer and are protocol transparent. Each port on a switch is a separate collision domain but part of the same broadcast domain. Switches do not control broadcasts on the network.

Routers

Routers make forwarding decisions based on network layer addresses. In addition to controlling collision domains, routers control broadcast domains. Each interface of a router is a separate broadcast domain defined by a subnet and a mask. Routers are protocol aware, which means they are capable of forwarding packets of routed protocols such as IP, IPX, Decnet, and AppleTalk. Figure 4-20 describes a router; each interface is a broadcast and a collision domain.

Figure 4-20 *Routers Control Broadcast and Collision Domains*

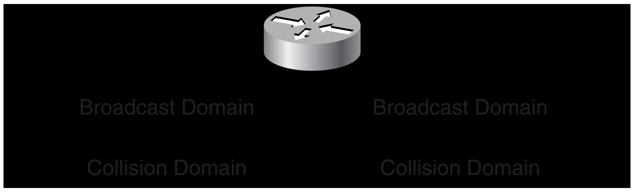

Routers exchange information about destination networks by using one of several routing protocols. The following are lists of routing protocols. The lists are divided by the protocols that can be routed.

For routing TCP/IP:

- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP)

- Intermediate System-to-Intermediate System (ISIS)
- Protocol Independent Multicast (PIM)

For routing Novell:

- Novell Routing Information Protocol (Novell RIP)
- NetWare Link Services Protocol (NLSP)
- Enhanced Interior Gateway Routing Protocol (EIGRP)

For routing AppleTalk:

- Routing Table Maintenance Protocol (RTMP)
- Enhanced Interior Gateway Routing Protocol (EIGRP)

Routing protocols are discussed in further detail in Chapter 6, “Designing for Specific Protocols.”

Routers are the preferred method of forwarding packets between networks of differing media, such as Ethernet to Token Ring, Ethernet to FDDI, or Ethernet to Serial. They also provide methods to filter traffic based on the network layer address, route redundancy, load balancing, hierarchical addressing, and multicast routing.

Layer 3 Switches

LAN switches that are capable of running routing protocols are *Layer 3 switches*. These switches are capable of running routing protocols and communicating with neighboring routers. An example is a Catalyst 5500 with a Routing Switch Module (RSM). Layer 3 switches have LAN technology interfaces that perform network layer forwarding; legacy routers provide connectivity to WAN circuits. The switches off-load local traffic from the WAN routers.

Layer 3 switches perform the functions of both data link layer switches and network layer routers. Each port is a collision domain. Interfaces are grouped into broadcast domains (subnets) and a routing protocol is selected to provide network information to other Layer 3 switches and routers.

Combining Hubs, Switches, and Routers

Available in Ethernet and Fast Ethernet, hubs are best used in small networks where there are few nodes on the segment. Hubs do not control the broadcasts nor do they filter collision domains on the network. If higher bandwidth is required, use 100 Mbps hubs. When the number of nodes on the network grows, move to switches.

With the cost of switch ports comparable to hubs, use switches as the basic network connectivity devices on the network. Switches reduce collisions and resolve media contention on the network by providing a collision domain per port. Replace hubs with switches if the utilization is over 40 percent on Ethernet networks or above 70 percent on Token Ring and FDDI networks. Switches cannot resolve broadcast characteristics of protocols; use routing to resolve protocol-related problems. As you can see in the sample in Figure 4-21, the repeaters are pushed to the outer layer of the design, connecting to switches. Switches control the collision domains. Fast Layer 3 switches are used for routing between LAN segments, and the router provides access to the WAN.

Figure 4-21 *Combining Routers, Switches, and Hubs*

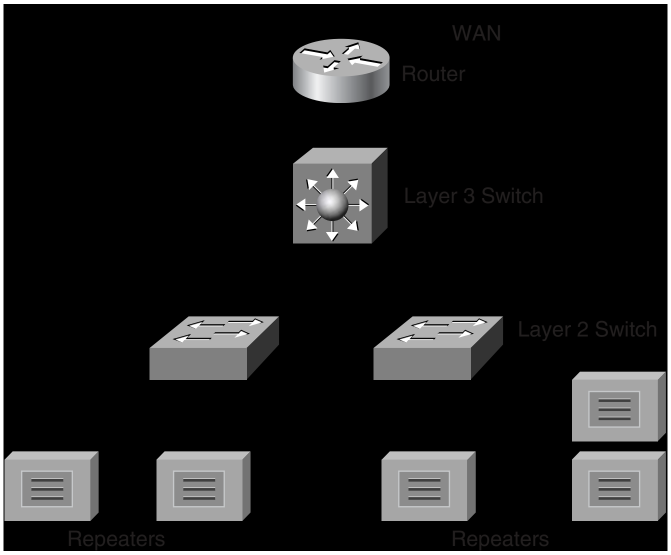

Use routers for segmenting the network into separate broadcast domains, security filtering, and access to the WAN. If broadcast traffic on the network is over 20 percent, use routing.

Cisco LAN Equipment

The CCDA objectives covered in this section are as follows:

2	Assemble Cisco product lines into an end-to-end networking solution.
16	Recommend Cisco products and LAN technologies that will meet a customer’s requirements for performance, capacity, and scalability in small- to medium-sized networks.
17	Update the network topology drawing you created in the previous section to include hardware and media.

A CCDA must be familiar with Cisco products, product capabilities, and how to best apply the products to meet performance, scalability, redundancy, and cost requirements. This section lists and explains Cisco equipment for LAN requirements. A complete list of Cisco products can be found at the CCO web site.

FastHub 400

The FastHub 400 10/100 series is a full line of products that includes 12- and 24-port 10/100 Fast Ethernet repeaters in managed and manageable versions. The FastHub 400 10/100 series provides low-cost 10/100 autosensing desktop connectivity where dedicated bandwidth is not required. The Cisco 412 provides 12 UTP ports of 10/100 Fast Ethernet. The Cisco 424M provides 24 UTP ports of 10/100 Fast Ethernet in a SNMP-managed version.

Cisco Catalyst 1900/2820 Series

The Catalyst 1900 and 2820 series provide 12- or 24-switched, 10-Mbps 10BaseT ports. Different models provide Fast Ethernet uplinks in 100BaseT and 100BaseF media. Different models can keep 1KB, 2KB, or 8KB storage of MAC addresses. The specifications of the various models in these series are presented in Table 4-10.

Table 4-10 *Catalyst 1900 and 2820 Series Specifications*

Model	Specifications
WS-C1912-EN	• 12 10BaseT • Two 100BaseTX • 1KB MAC • Enterprise Edition
WS-C1912C-EN	• 12 10BaseT • One 100BaseTX • One 100BaseFX • 1KB MAC • Enterprise Edition
WS-C1924-EN	• 24 10BaseT • Two 100BaseTX • 1KB MAC • Enterprise Edition

Table 4-10 *Catalyst 1900 and 2820 Series Specifications (Continued)*

Model	Specifications
WS-C1924C-EN	• 24 10BaseT • One 100BaseTX • One 100BaseFX • 1KB MAC • Enterprise Edition
WS-C1924F-EN	• 24 10BaseT • Two 100BaseFX • 1KB MAC • Enterprise Edition
WS-C1924-EN-DC	• 24 10BaseT • Two 100BaseTX • 48-volt DC Dual-Feed Power System • 1KB MAC • Enterprise Edition
WS-C2822-EN	• 24 10BaseT • Two slots • 2KB MAC • Enterprise Edition
WS-C2828-EN	• 24 10BaseT • Two slots • 8KB MAC • Enterprise Edition

Catalyst 2900

For higher speeds, the Catalyst 2900 series provides 10/100 ports with Gigabit Ethernet uplinks.

Catalyst 2948G offers 48 ports of 10/100 Ethernet with two Gigabit Ethernet uplinks.

Catalyst 3000 Series Stackable Switches

The Catalyst 3100 switch is designed for networks that require flexibility and growth with minimal initial investment. This switch contains 24 fixed 10BaseT Ethernet ports, one StackPort slot for scalability, and one expansion FlexSlot for broad media support. It is designed for a variety of campus LAN and enterprise WAN solutions; the Catalyst 3100 switch fits well in a wiring closet and branch office applications.

The Catalyst 3200 is a high port density stackable switch chassis with a modular Catalyst 3000 architecture supervisor engine and seven additional media expansion module slots. The expansion slots are backward compatible with all existing Catalyst 3000 media expansion modules. The seventh slot, called *FlexSlot*, is an expansion slot that accepts either a standard Catalyst 3000 expansion module or new doublewide expansion modules providing forward and backward investment protection.

The 3011 WAN access module for the Catalyst 3200 and Catalyst 3100 provides WAN interconnect integrated with the switch backplane. The 3011 WAN access module was the first FlexSlot module to be introduced. Based on the Cisco 2503 router, the 3011 provides two high-speed serial ports, an ISDN BRI port, and an auxiliary (AUX) port.

Catalyst 3900 Token Ring Stackable Switch

The Catalyst 3920 switch provides 24 Token Ring ports. With the Catalyst 3920 switch, you can start with a single 24-port switch and add capacity as you need it, while still managing the entire stack system as one device.

Catalyst 3500 10/100 Autosensing Switch

The Catalyst 3500 XL architecture is designed to meet the technical requirements of autosensing 10/100BaseT Ethernet interfaces. *Autosensing* enables each port to self-configure to the correct bandwidth upon determining whether it is connected to a 10- or 100-Mbps Ethernet channel. This feature simplifies setup and configuration and provides flexibility in the mix of 10 and 100 Mbps connections the switch supports. Network managers can alter connections without having to replace port interfaces.

GBIC-Based Gigabit Ethernet Ports

Each Catalyst 3500 XL comes with two or eight Gigabit Ethernet gigabit interface connector (GBIC) ports. Customers can use any of the following IEEE 802.3z-compliant GBICs based on their connection needs: 1000BaseSX, 1000BaseLX/LH, or the Cisco GigaStack stacking GBIC. These GBIC ports support standards-based, field-replaceable media modules and provide unprecedented flexibility in switch deployment while protecting customers' investments.

Catalyst 4000

The Catalyst 4912G is a 12-port dedicated Gigabit Ethernet switch featuring high-performance Layer 2 switching and intelligent Cisco OSI network (Layer 3) services for high-speed network aggregation.

The Catalyst 4003 offers 24 Gbps of switching bandwidth and provides expansion to 96 ports of 10/100 Ethernet or 36 ports of Gigabit Ethernet. Up to 96 10/100 Ethernet ports, or up to 36 Gigabit Ethernet ports, can be installed into one managed unit.

The Catalyst 4000 series provides an advanced high-performance enterprise switching solution optimized for wiring closets with up to 96 users and data center server environments that require up to 36 Gigabit Ethernet ports. New FlexiMod uplinks support up to eight 100BaseFX riser connections with EtherChannel benefits. The Catalyst 4000 series provides intelligent Layer 2 services leveraging a multiGigabit architecture for 10/100/1000-Mb Ethernet switching. The modular three-slot Catalyst 4003 system leverages the software code base from the industry-leading Catalyst 5500/5000 series to provide the rich and proven feature set that customers demand in the wiring closet for true end-to-end enterprise networking.

Catalyst 5000 Switch Series

The Cisco Catalyst 5000 series features modular chassis in 2-, 5-, 9-, and 13-slot versions. All chassis share the same set of line cards and software features, which provides scalability while maintaining interoperability across all chassis.

The Catalyst 5002 is positioned to deliver a consistent architecture and features set in a smaller package that addresses the needs of smaller wiring closets. The Catalyst 5002 switches at the 1 Mpps (million packets per second) range. The Catalyst 5002 is a fully modular, two-slot Catalyst 5000 series member, using the same architecture and software as the Catalyst 5000. The switch can deliver more than one million packets per second throughput across a 1.2-Gbps, media-independent backplane that supports Ethernet, Fast Ethernet, FDDI, Token Ring, and ATM.

The Catalyst 5000 will continue to address the needs of switched 10BaseT and group switched wiring closets with performance in the 1–3 Mpps range.

The Catalyst 5505, a five-slot chassis like the Catalyst 5000, is designed for a high-end wiring closet and data applications with performance in the 1–25 Mpps range. The Catalyst 5505 combines the size of the original Catalyst 5000 with the performance boost and added features of the Catalyst 5500 series.

The Catalyst 5509 supports high-density 10/100 Ethernet for the wiring closet, or high-density Gigabit Ethernet for backbone applications, delivering over 25-Mpps switching performance. The Catalyst 5509 provides dedicated switching for up to 384 users, making this chassis an ideal platform for wiring closet solutions. The Catalyst 5509 also supports high-density Gigabit Ethernet for switched intranet backbones and data centers.

The Catalyst 5500 is the most versatile switch in the Catalyst family, able to support LightStream 1010 ATM switching or Catalyst 8500 Layer 3 switching line cards in addition to all the Catalyst 5000 family line cards. The Catalyst 5500 is positioned as a high-capacity wiring closet or data center switch, delivering over 25-Mpps switching performance.

The Catalyst 5500 is a 13-slot chassis that is rack-mountable using the rack-mount kit. All functional components, including power supplies, fan trays, supervisors, ATM switch processors (ASPs), and interface modules are accessible and hot-swappable from the network side of the chassis. This setup ensures ease of use in tight wiring closets.

Foundation Summary

Foundation Summary is a section presented in a concise format to provide quick reference information relating to the objectives covered in this chapter.

Table 4-11 *Ethernet CSMA/CD Based Media*

Specification	Speed	Max Segment Size	Encoding	Media
10Base5	10 Mbps	500m	Manchester	0.4in 50ohm Coax (Thicknet)
10Base2	10 Mbps	185m	Manchester	0.2in 50ohm Coax (Thinnet)
10BaseT	10 Mbps	100m	Manchester	UTP
100BaseT	100 Mbps	100m	4B/5B	UTP
1000BaseT	1000 Mbps	100m	5-level	Cat 5 UTP
1000BaseLX (long wave)	1000 Mbps	550m	8B/10B	Single/multimode fiber
1000BaseSX (short wave)	1000 Mbps	62.5 micrometers: 220m 50 micrometers: 500m	8B/10B	Multimode fiber
1000BaseCX	1000 Mbps	25m	8B/10B	Shielded balanced copper

Table 4-12 *Token Access Based Media*

Type	Speed	Ring types	Encoding	Media
Token Ring	4/16 Mbps	Unidirectional single ring	Differential Manchester	UTP, STP
FDDI	100 Mbps	Dual counter rotation rings	4B/5B with nonreturn to zero inverted (NRZI)	Fiber

Table 4-13 *Network Devices*

Device	OSI Layer	Protocol	Domains	Understands
Repeaters	Layer 1: Physical	Transparent	Amplify signal	Bits
Hubs	Layer 1: Physical	Transparent	Amplify signal	Bits
Bridges	Layer 2: Data link	Transparent	Collision domain	Frames
Switches	Layer 2: Data link	Transparent	Collision domain	Frames
Routers	Layer 3: Network	Aware	Broadcast domain	Packets
Layer 3 Switches	Layer 3: Network	Aware	Broadcast domain	Packets

Table 4-14 *LAN Types*

LAN Type	Characteristics
Large building LAN	Large number of users, data center, floor closet switches
Campus LAN	High-speed backbone switching
Small/remote LAN	Small number of users, small hubs/switches

Table 4-15 *Cisco Devices*

Device	Characteristics
FastHub 400 repeater	12/24 ports of 10/100 Fast Ethernet
1900/2820 switch	12/24 ports of 10BaseT, 100 Mbps uplinks
2948G switch	48 ports of 10/100 Ethernet, 2 Gigabit Ethernet uplinks
3000 series switches	24 ports of 10BaseT stackable switches with expansion slots
3500 switch	10/100 autosensing, 2 Gigabit GBIC ports
3900 switch	24 Token Ring ports
4000 switch	Up to 24 Gigabit switched ports, plus expansion slots for up to 96 10/100 ports of 36 Gigabit Ethernet ports
5002 switch	2 slot modular chassis, 1 Mpps
5000 switch	5 slot modular chassis, 1–3 Mpps
5505 switch	5 slot modular chassis, 25 Mpps
5509 switch	9 slot modular chassis, 25 Mpps
5500 switch	13 slot modular chassis, supports Layer 3 line cards, and ATM modules

Q&A

The following questions are designed to test your understanding of the topics covered in this chapter. When you have answered the questions, you can find the answers in Appendix A, "Answers to Quiz Questions." After you identify the subject matter you missed, review those sections in the chapter until you feel comfortable with this material.

1 What is the maximum segment size in 10BaseT?

2 What is the maximum segment size in 10Base2?

3 What is the maximum segment size in 10Base5?

4 What is the maximum segment size in 100BaseT?

5 What is the maximum segment size in 1000BaseT?

6 What does the acronym DIX stands for?

7 What are the three layers of hierarchical design?

8 At what percent utilization are Ethernets over-utilized?

9 At what percent utilization are Token Ring networks over-utilized?

10 At what percent utilization are FDDI networks over-utilized?

11 What is the maximum recommended percentage of broadcasts on the network?

12 What is the standard(s) for Gigabit Ethernet?

13 What standard governs Token Ring?

14 What media implements a dual-ring and forwards tokens?

15 Routers operate on which layer of the OSI model?

16 Switches operate on which layer of the OSI model?

17 Repeaters operate on which layer of the OSI model?

18 Bridges operate on which layer of the OSI model?

19 Transceivers operate on which layer of the OSI model?

20 Are bridges protocol transparent?

21 When switches implement cut-through switching mode, what is not verified to check for frame errors?

____

____

____

22 True or False: A repeater keeps a table of each MAC address on its ports and forwards frames accordingly.

____

23 True or False: Routers forward frames based on the destination MAC address.

____

24 True or False: Bridges forward frames based on the source MAC address.

____

25 What LAN media uses dual counter rotating rings?

____

____

____

26 Which Cisco device provides 48 ports of 10/100 Ethernet with 2 Gb uplinks?

____

____

____

27 What are the components of the three-part firewall system?

____

____

____

28 What is the encoding scheme of 10-Mbps Ethernet?

____

____

____

29 What is the encoding scheme of Token Ring?

30 100BaseT forwards frames at what speed?

Case Studies

The following case study questions are based on the ongoing scenarios that are presented in the "Case Studies" section of Chapter 1, "Design Goals." If you want to familiarize yourself with the entire scenario, refer to that section before working through the following questions. The answers to these questions can be found in the "Case Study Answers" section at the end of this chapter.

Case Study #1: GHY Resources

1. What issues does GHY Resources have on its Ethernet segments that may cause packet loss on the network?
2. Is the LAN in Kansas City running over the recommended maximum utilization?
3. Draw the current LAN network at the headquarters in St. Louis.
4. If each sales office has less than 20 nodes but may grow to over 30, what switch would you recommend to meet the current requirements?
5. The headquarters segments have 30 users each and will not grow to over 48 nodes. Draw out the topology for a possible solution. Update the topology with Cisco products for this building LAN that uses switched ports and Fast Ethernet media.

Case Study #2: Pages Magazine, Inc.

1. What type of media is used at Pages Magazine, Inc., locations?
2. Ms. Phillips mentions that their Ethernet segments are running over 40 percent utilization during peak hours of the day. What would you suggest?
3. Is the number of nodes at the remote sites too large?

Case Study Answers

Case Study #1: GHY Resources

1 What issues does GHY Resources have on its Ethernet segments that may cause packet loss on the network?

Segments are running at 45 percent utilization.

Broadcast storms on the network.

2 Is the LAN in Kansas City running over the recommended maximum utilization?

No, the maximum is around 40 percent.

3 Draw the current LAN network at the headquarters in St. Louis.

Figure 4-22 shows a representation of the current LAN at the headquarters in St. Louis.

Figure 4-22 *Current Headquarters LAN*

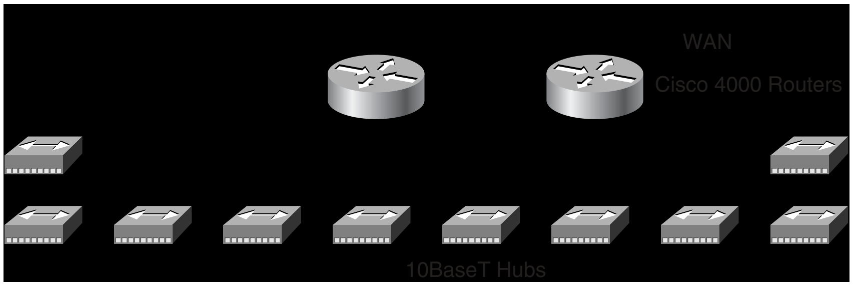

4 If each sales office has less than 20 nodes but may grow to over 30, what switch would you recommend to meet the current requirements?

Catalyst 3200 stack switch

Catalyst 2948G switch

5 The headquarters segments have 30 users each and will not grow to over 48 nodes. Draw out the topology for a possible solution. Update the topology with Cisco products for this building LAN that uses switched ports and Fast Ethernet media.

Figure 4-23 shows a strong topological solution to the customer's needs.

Figure 4-23 *GHY Headquarters LAN Solution*

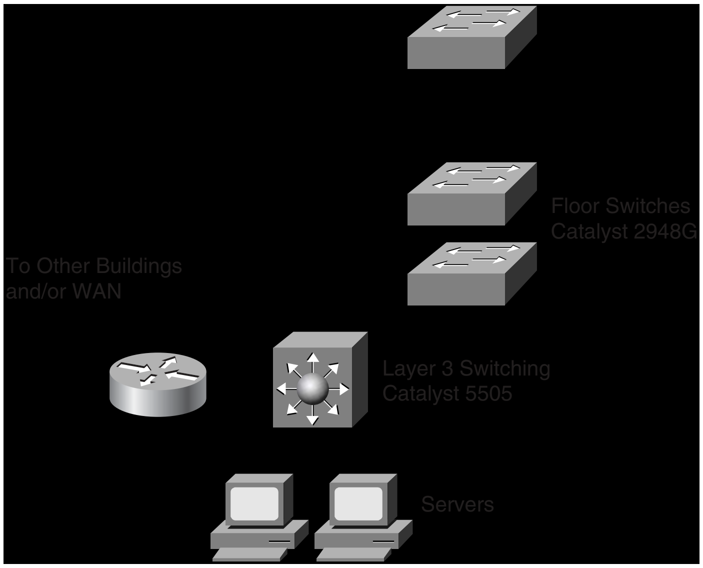

Case Study #2: Pages Magazine, Inc.

1 What type of media is used at Pages Magazine, Inc., locations?

Unshielded twisted pair

2 Ms. Phillips mentions that Pages Magazine's Ethernet segments are running over 40 percent utilization during peak hours of the day. What would you suggest?

A good suggestion would be to replace the 10BaseT hubs with 100BaseT switches.

3 Is the number of nodes at the remote sites too large?

No. The number of nodes at each remote site is under the recommended maximum of 200 for multiprotocol networks.