- Prepare for CCDP certification while mastering the intricacies of internetwork design
- Build your campus LAN design skills with information on the technical, business, hardware, topology, VLAN, LANE, and design issues affiliated with campus LANs
- Master the skills needed for effective TCP/IP network design, including addressing and routing decisions, OSPF, IGRP, and EIGRP
- Learn to construct effective IPX, AppleTalk, and Windows Networking-based internetworks
- Understand how to create effective WANs using Frame Relay, X.25, and ATM, as well as how to contend with design using dedicated lines and remote access
- Refine your knowledge of complex SNA internetworking concepts
- Apply your internetworking knowledge by redesigning networks presented in six exhaustive case studies

This volume is in the Certification and Training Series offered by Cisco Press. Books in this series provide technology-specific information to help networking professionals prepare for the Cisco Career Certifications examinations.

Category: Cisco Press—Cisco Certification
Covers: Network Design

Cisco Systems
Cisco Press
www.ciscopress.com

CISCO CAREER CERTIFICATIONS

4

Cisco Internetwork Design

Based on the Cisco Systems instructor-led course available worldwide, *Cisco Internetwork Design* teaches you how to plan and design a network using various internetworking technologies. Created for those seeking to attain CCDP certification, this book presents the fundamental, technical, and design issues associated with campus LANs; TCP/IP networks; IPX, AppleTalk, and Windows-based networks; WANs; and SNA networks.

By using this book, you will be able to identify internetwork requirements, determine appropriate infrastructure and routing issues within an internetwork, and construct a viable plan to deploy or upgrade to a more effective network topology. Filled with invaluable foundation information on various internetworking technologies and supported with useful design examples, *Cisco Internetwork Design* will aid you in determining the best internetworking technology for your needs. Self-assessment through chapter-ending questions starts you down the path for attaining your CCDP certification.

Cisco Systems, Inc., is the leading global supplier of internetworking solutions, including routers, LAN and ATM switches, dial-up access servers, and network management software. The core of this book is a collaboration of multiple authors and experts over time to develop the definitive course in Cisco internetwork design.

Matthew H. Birkner is a network consulting engineer and works for Cisco Systems in the NSA (Network Supported Accounts) Program. Currently, he supports and designs enterprise networks. He has been a network design engineer, network operations center engineer, and technical support specialist. Matt holds a B.S. degree from Tufts University, where he majored in electrical engineering. Additionally, he is a Cisco Certified Internetwork Expert (CCIE), Cisco Certified Design Professional (CCDP), Certified Netware Engineer(CNE), and Bay Networks Certified Specialist (BNCS).

$60.00 USA
$89.95 CAN

ISBN 1-57870-171-6

6 19472 70171 3
9 781578 701711 96000

CISCO CAREER CERTIFICATIONS

CCDP

Cisco Internetwork Design

Revised and Edited by Matthew H. B... E

Cisco Press

CISCO CAREER CERTIFICATIONS
CISCO CERTIFIED DESIGN PROFESSIONAL

91

Edited and Revised by
Matthew H. Birkner
CCIE #3719

Prepare for CCDP certification with the official CID coursebook

Cisco Internetwork Design

Cisco Systems
Cisco Press
www.ciscopress.com

I see designs by technologists who fall in love with specific technologies and cause problems because of overemphasizing them. I like candy, but candy for breakfast, lunch, and dinner can become unappealing, not to mention unhealthy. Networks are the synthesis of both Layer 2 and Layer 3 technologies, and the best designs are a balance of these technologies.

—Bill Kelly, Chief Network Architect, Cisco Systems

Upon completion of this chapter, you will be able to do the following:

- Describe cabling topologies used in campus LAN designs
- Describe and implement distributed backbone designs in campus LANs
- Describe and implement collapsed backbone designs in campus LANs
- Describe deployment of VLANs and LANE in campus LANs
- Determine where to use switches, where to use bridges, and where to use routers in segmenting campus LANs

CHAPTER 5

Campus LAN Design Models

Campus LANs have their own set of design ground rules, which, when understood and implemented, can produce positive results. However, I have seen many network designers ignore key areas of campus LAN design, such as cabling topologies, or perhaps unknowingly let VLAN management get out of control. The result is always the same—unplanned outages and unnecessary network downtime. The techniques presented in Chapters 1–5 are the network design building blocks and the litmus test for any network designer, regardless of the protocols, topologies, or devices used. The fundamentals discussed in these chapters are paramount to network success, yet oddly, these are the very areas that are most frequently skipped.

Sample Building Layout

Figure 5-1 shows a typical building layout that is referenced throughout the rest of the chapter.

This sample building layout assumes that user end stations (clients) must access servers and other end stations. The servers may be identified as those that are local to a workgroup or a floor, and those that are used more widely (building, campus, or enterprise-wide). End-station connections will be required for the workgroup, which, as shown in Figure 5-1, is isolated to a floor. A wiring closet is available on each floor to connect the client workstations and floor (local) servers. The riser goes between floors and provides a cable path to interconnect workgroups located throughout the building. The basement houses a data center, where the most heavily used servers are located, along with cable connections to other buildings and sites as required.

Figure 5-1 *Sample Building Layout*

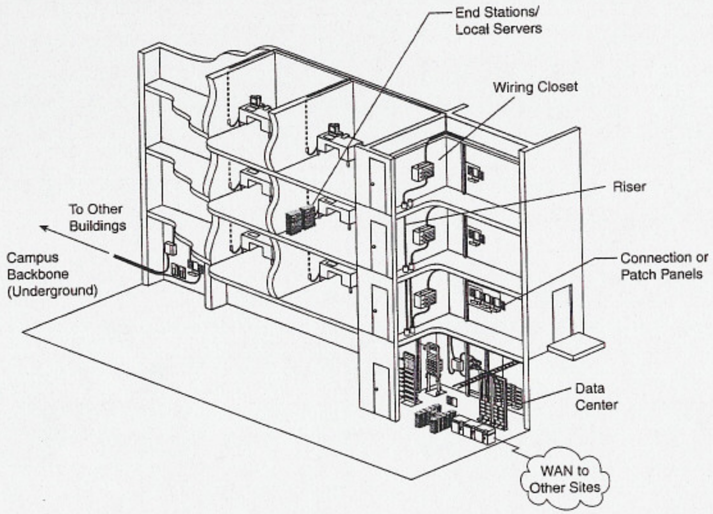

Cabling Issues

Cabling recommendations in general call for multimode fiber in the risers between floors and in the tunnels that typically connect buildings in a campus environment. Fiber allows the bandwidth to be scaled up as far as necessary (FDDI, 100BaseFX, fiber-based ATM, Gigabit Ethernet) while providing reliable links that are resistant to sources of electromagnetic interference (EMI). As shown in Figure 5-2, network interconnection devices (hubs, bridges, switches, routers) are typically located in the network closet on each floor, and in the data center in the basement. Category 5 unshielded twisted-pair (UTP) in most cases should be used to link desktop clients and local servers to the active components in the network closets. A Category 5 UTP segment is generally considered to be usable at speeds up to 100 Mbps, over a maximum distance of 100 meters.

Figure 5-2 *Building Cabling Schematic*

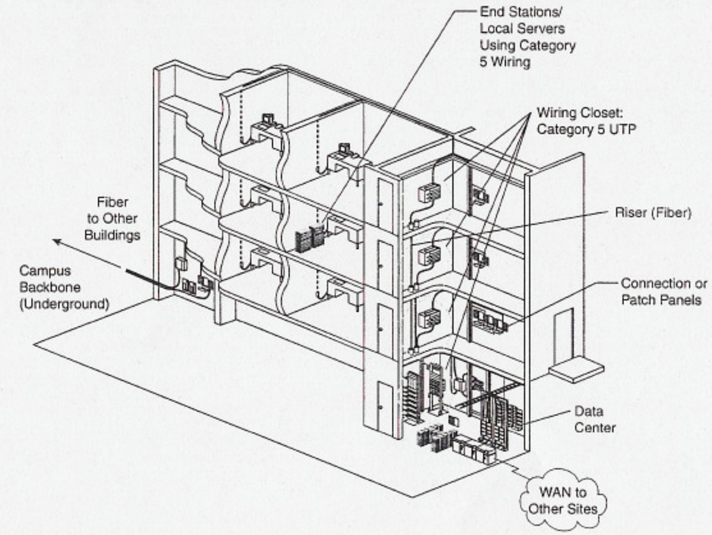

Distributed Backbones

Distributed backbones are a somewhat less-flexible approach to wiring a building, but they also spread your risk of a problem over several devices or topologies. Therefore, distributed backbones generally do not contain a single point of failure. Just like any other design approach, however, if not selected and designed carefully, there can be several points of failure. (This situation is often corrected by deploying a second router running Hot Standby Router Protocol [HSRP] at each backbone location because HRSP automatically detects a network or router failure and subsequently switches to the alternate router without the end-user systems noticing that a problem has ever occurred.)

DESIGN RULE Use HSRP to eliminate a single point of failure at the router.

Distributed Backbones in Buildings

Figure 5-3 shows the distributed backbone model where routers on each floor connect to a backbone in the riser. In this case, the backbone is typically an FDDI ring. This design has the benefit of distributing the connections to the backbone and thereby eliminating any single point of failure.

Figure 5-3 *Distributed Backbone in a Building*

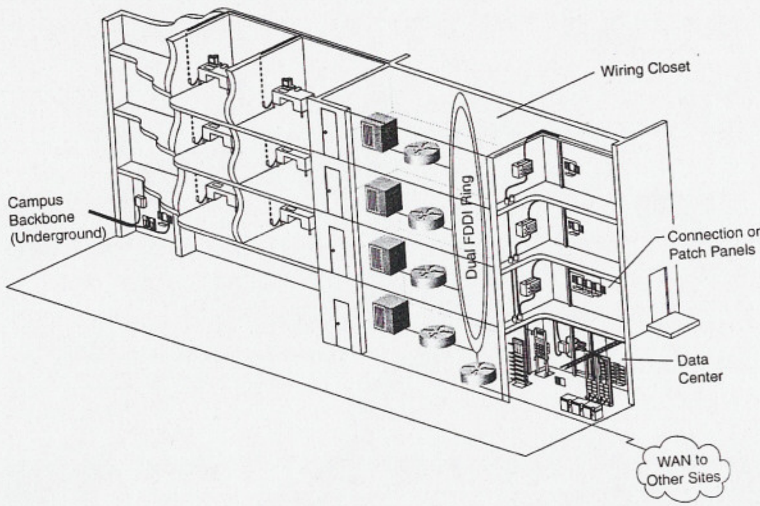

This design also has drawbacks. Multiple IP network numbers within a building reduce the ease with which user adds, moves, and changes can be made. Furthermore, this design tends to be more expensive and does not always migrate easily to switching.

Note that in this design, no end stations are attached to the backbone. This is an important design guideline that we should continue to follow. The backbone should be used only as a transit path between local networks, and not as a host-based network itself. This guideline keeps the backbone more stable, facilitates traffic management and capacity planning, and enhances overall scalability for future redesign. Possible exceptions include heavily utilized servers (such as servers providing Domain Name System [DNS] or Simple Mail Transfer Protocol [SMTP]); however, these servers need to be closely monitored. If you have a bad server NIC on the backbone, for example, it could potentially cause a network outage, with far-reaching impact. This is because many servers, unlike routers and switches, do not have advanced system utilities to quickly correct and isolate network problems. Therefore, servers on the backbone are not generally recommended.

NOTE Note that we already discussed this principle in Chapter 2, "Hierarchical Design," when we stressed that distribution layer services should be kept separate from the access layer.

Distributed Backbones on the Campus

The distributed backbone on the campus is a more resource-efficient solution than in a building. As depicted in Figure 5-4, this solution involves a single router per building, typically located in the basement, with a combination of hubs and switches providing user access throughout the building. Using fewer logical networks per building increases the ease of user adds, moves, and changes. The only drawback is the lack of flexibility in connecting to other buildings on the campus. Switching could easily be deployed in the building, but not across the campus.

Figure 5-4 *Distributed Backbone on the Campus*

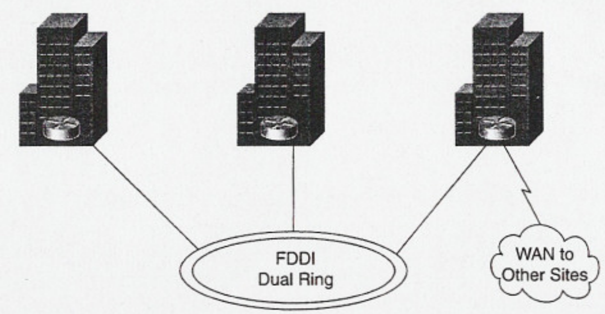

Collapsed Backbones

Collapsed backbones generally represent a much more flexible and cost-effective approach to wiring a building. For example, Figure 5-5 shows the traditional ring or bus backbone collapsed inside a single switch or router (Figure 5-5 shows a router), making the device backplane act as the backbone network. Hubs on each floor connect to the end stations in the workgroup, and each hub attaches to a separate LAN interface on the router.

Figure 5-5 *Collapsed Backbone—Router/Hub Design*

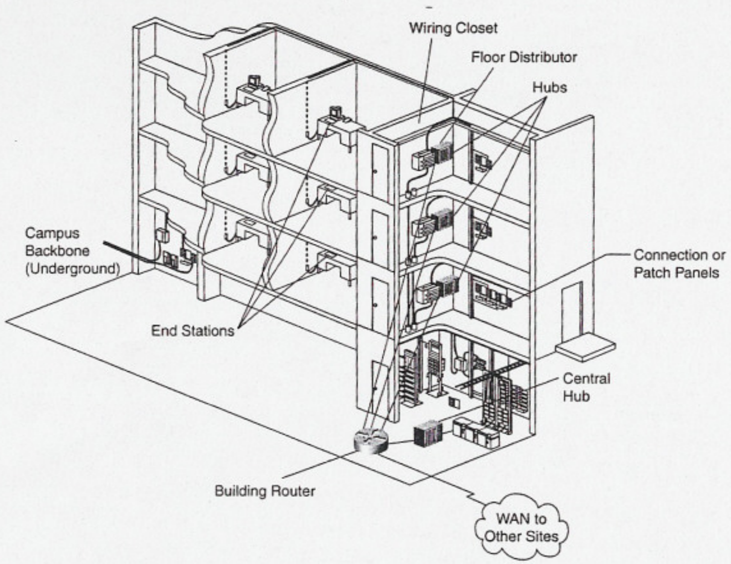

This design would make moving users a little easier than in the distributed backbone model, but the solution is not yet ideal. Also, the router represents a single point of failure. This situation could be corrected by deploying a second router that attaches to the hub and uses the HSRP and redundant hubs running Multigroup HSRP (MHSRP).

In another option, as shown in Figure 5-6, the hub in the preceding design could easily be replaced with an Ethernet switch to provide more bandwidth to the workgroup. However, the use of a single router represents a bottleneck in the traffic flow and a single point of failure.

Figure 5-6 *Collapsed Backbone—Router/Switch Design*

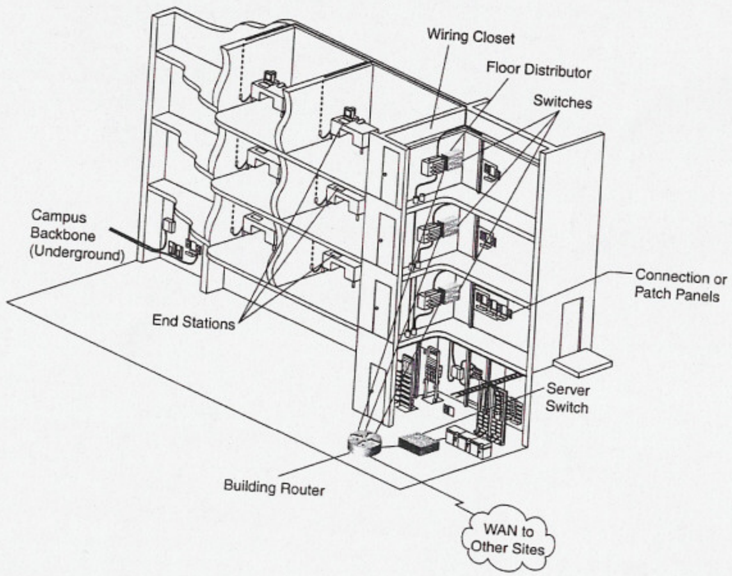

The collapsed backbone design can easily be extended to accommodate VLANs, just by adding another Ethernet switch in the data center. Figure 5-7 shows the backplane of the switch now becoming the collapsed backbone for the building. In this case, the router has a separate Ethernet port for each VLAN, and communication between VLANs can occur only through the router. The router also provides the edge connection point to access other building networks, as before. Again, it is recommended to use multiple hubs and multiple routers for best redundancy.

Figure 5-7 *Collapsed Backbone—Building VLAN*

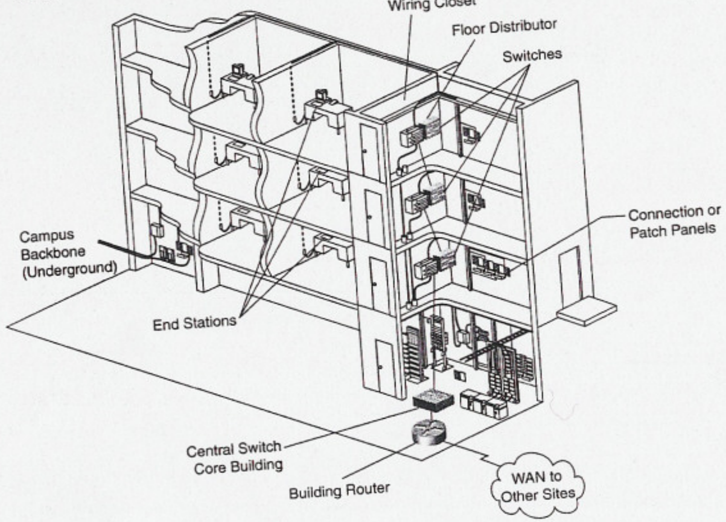

Deploying VLANs Within a Building

The use of VLANs provides more flexibility in the physical positioning of end stations and servers. User workstations can be physically placed anywhere in the building and still remain in the same LAN. In addition, a user workstation can remain in its current physical location and still move to a new logical LAN (VLAN) assignment. All the servers in the building can be placed in one physical location (server farm) but still remain logically in separate LANs. Figure 5-8 shows a typical collapsed backbone VLAN.

Figure 5-8 *Collapsed Backbone—VLAN*

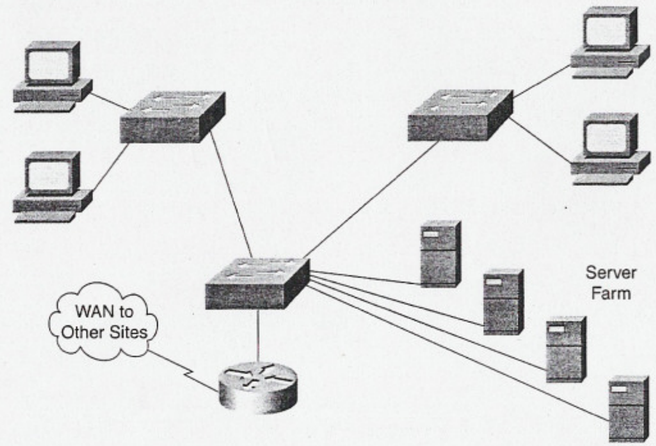

The trunk links between the floors should then be FDDI or Fast Ethernet, and all frames would carry the VLAN ID. The router connection to the backbone switch could also be FDDI or Fast Ethernet and would then receive the VLAN ID. Separate subinterfaces could then be set up for each VLAN, and the router would route between them accordingly.

Deploying VLANs Across a Campus

The collapsed backbone concept could be taken further to include the entire campus. As Figure 5-9 shows, one switch would act as the backbone for the entire campus, which would allow maximum flexibility in moving users around the campus, either physically or logically. All servers on the campus could then be placed in one location for ease of administration.

Figure 5-9 *Collapsed Backbone—Campus VLAN*

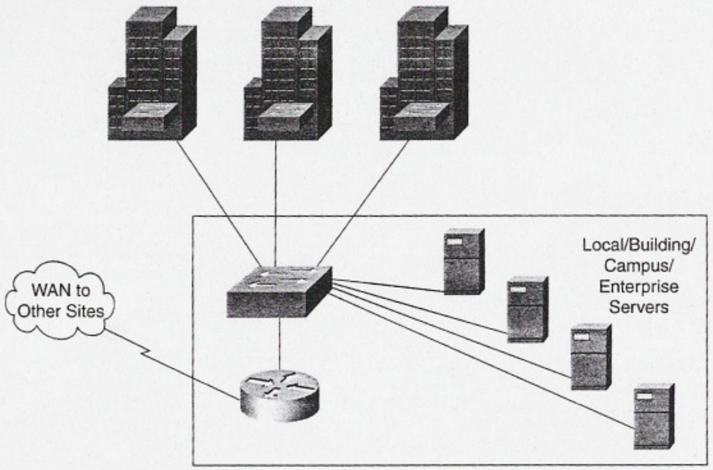

The potential for growth should always be considered an important variable in any network design, so you should choose your VLAN assignments carefully and follow a naming convention that is easy to troubleshoot, as discussed in Chapter 1, "Internetwork Design Overview." Also, note that the potential for wide-scale problems increases if you do not manage your VLAN implementation carefully. I have recently seen cases where VLAN trunking problems have caused a ripple effect of network instability across an entire VLAN domain, which could be a global enterprise network. This type of problem can be avoided, however, with proper planning, regular design reviews, and effective change-management processes and procedures.

DESIGN RULE Choose your VLAN assignments carefully. Whenever possible, assign stations to VLANs such that only 20 percent of their traffic is destined to other VLANs.

133

Migrating to ATM in the Campus LAN

Migration to ATM would be needed when routinely transferring large amounts of data, and will typically follow one of two strategies:

- Using a distributed campus backbone
- Using the LANE model

As Figure 5-10 shows, in the case where a distributed campus backbone is used, the ATM switch (or switches) replaces the FDDI ring. In this case, the ATM network generally is fully meshed and represents one logical subnet, just as the FDDI ring did in Figure 5-4.

Figure 5-10 *Migrating to ATM—Distributed Campus Backbone Model*

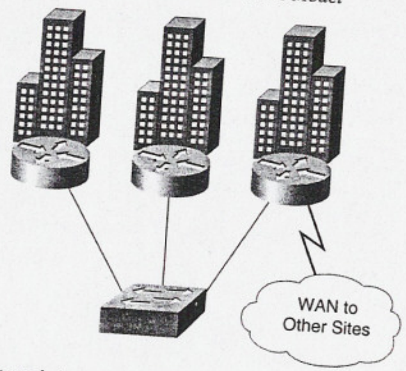

The second strategy is to use a LANE model. You can achieve increased flexibility by going to the LANE model. In essence, LANE describes an architecture of VLANs deployed in ATM environments. The router still would act as the interconnection point for traffic traveling between VLANs. Figure 5-11 illustrates a typical LANE network diagram.

NOTE You can find more detailed information about ATM LANE at www.atmforum.com.

Figure 5-11 *Migrating to ATM with the LANE Model*

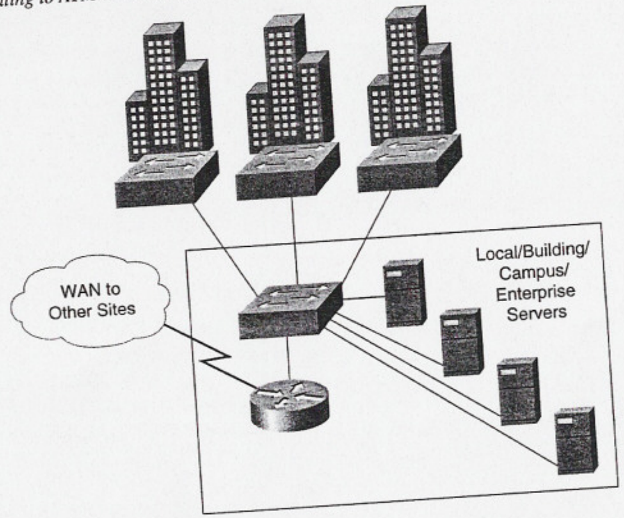

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) is a reliable method for automatically assigning IP addresses to hosts on your network. Rather than assigning static IP addresses to hosts, they are dynamically configured by a DHCP Server. The DHCP Server provides a master listing of IP addresses and host information for all registered hosts. DHCP is discussed further in Chapter 7, "TCP/IP Addressing Design," but it is important to note that subsequent to DHCP, the impact of simplifying the add/move/change process with the benefits of small broadcast domains could be had. DHCP completely conceals the complexity of the addressing structure.

In a nutshell, as shown in Chapter 7, DHCP permits the ease of configuration of a Layer 2 network with the protections and performance of a Layer 3 design.

DESIGN RULE If possible, implement DHCP in your network today.

Summary

Much like the other design techniques discussed in previous chapters, campus LANs have their own set of guidelines, which, when properly followed, increase the likelihood of long-term network stability. You will face issues such as cabling topologies and whether to use distributed or collapsed backbone design strategies. You will also need to consider the deployment of VLANs and LANE in campus LANs. In many cases, you'll find that your final design uses a balance of Layer 2 and Layer 3 technologies. Without such a balance, your network could get a major "stomachache" from overloading on one particular technology or approach for something other than technical reasons.

Chapter Review Questions

1 An organization requires a network that will permit any user within a building to be in any logical workgroup. It also requires that all servers in the building be centrally located. What design would you recommend to meet these needs?

2 How would you extend the design in Question 1 so that a user could be located anywhere within a campus?

3 What is DHCP?

4 Why do cabling recommendations in general call for multimode fiber in the risers between floors and in the tunnels that typically connect buildings in a campus environment?

5 What is LANE?

6 What is Hot Standby Router Protocol (HSRP), and how can it be used in a VLAN design?

7 What is a widely deployed high-speed upgrade path for existing legacy FDDI distributed backbones?