CHAPTER 7
WIRED AND WIRELESS LOCAL AREA NETWORKS

Chapter Summary

This chapter examines the three major network architecture components that use Local Area Networks (LANs): the LANs that provide network access to users, the data center, and the e-commerce edge. We focus on the LANs that provide network access to users as these are more common. This chapter draws together the concepts from the first section of the book on fundamental concepts to describe how wired and wireless LANs work. We first summarize the major components of LANs and then describe the two most commonly used LAN technologies: wired and wireless Ethernet. The chapter ends with a discussion of how to design LANs and how to improve LAN performance.

Learning Objectives

After reading this chapter, students should:

- Understand the major components of LANs
- Understand the best practice recommendations for LAN design
- Be able to design wired Ethernet LANs
- Be able to design wireless Ethernet LANs
- Be able to improve LAN performance

Key Terms

access point (AP)
Active Directory Service (ADS)
association
beacon frame
bottleneck
bus topology
cable plan
cabling
Carrier Sense Multiple Access with Collision Detection (CSMA/CD)
Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)
channel
clear to send (CTS)
collision
collision detection (CD)
collision domain
cut-through switching
directional antenna
distributed coordination function (DCF)
domain controller
dual-band access point
Ethernet
fiber-optic cable
forwarding table
fragment-free switching
frame
hub
IEEE 802.3
IEEE 802.11
latency
layer-2 switch
lightweight directory access protocol (LDAP)
load balancer
logical topology
MAC address filtering
network-attached storage (NAS)
network interface card (NIC)
network operating system (NOS)
network profile
network segmentation
network server
omnidirectional antenna
overlay network
physical carrier sense method
physical topology
point coordination function (PCF)
port
power over Ethernet (POE)
probe frame
redundant array of inexpensive disks (RAID)
request to send (RTS)
server virtualization
shielded twisted-pair (STP)
site survey
small-office, home-office (SOHO)
storage area network (SAN)
store and forward switching
switching
switch
switched Ethernet
symmetric multi-processing (SMP)
topology
twisted-pair cable
unshielded twisted-pair (UTP) cable
virtual carrier sense
Warchalking
wardriving
Wi-Fi
Wi-Fi Protected Access (WPA)
WiGig
Wired Equivalent Privacy (WEP)
Wireless LAN (WLAN)
10Base-T
100Base-T
1000Base-T
10/100/1000 Ethernet
1 GbE
10 GbE
40 GbE
802.11ac
802.11ad
802.11a
802.11b
802.11g
802.11i
802.11n

Chapter Outline

1) Introduction
2) LAN Components
   a) Network Interface Cards
   b) Network Circuits
   c) Network Hubs, Switches, and Access Points
   d) Network Operating Systems
3) Wired Ethernet
   a) Topology
   b) Media Access Control
   c) Types of Ethernet
4) Wireless Ethernet
   a) Topology
   b) Media Access Control
   c) Wireless Ethernet Frame Layout
   d) Types of Wireless Ethernet
   e) Security
5) The Best Practice LAN Design
   a) Designing User Access with Wired Ethernet
   b) Designing User Access with Wireless Ethernet
   c) Designing the Data Center
   d) Designing the e-Commerce Edge
   e) Designing the SOHO Environment
6) Improving LAN Performance
   a) Improving Server Performance
   b) Improving Circuit Capacity
   c) Reducing Network Demand
7) Implications for Management
8) Summary

Answers to Textbook Exercises

Answers to End-of-Chapter Questions

1. Define *local area network*.

   A local area network is a group of microcomputers or other workstation devices located within a small or confined area and are connected by a common cable. A LAN can be part of a larger backbone network connected to other LANs, a host mainframe, or public networks.

2. Describe at least three types of servers.

   A LAN can have many different types of dedicated servers. Four common types are file servers, database servers, print servers, and communication servers. File servers allow many users to share the same set of files on a common, shared disk drive. A database server usually is more powerful than a file server. It not only provides shared access to the files on the server, but also can perform database processing on those files associated with client-server computing. The key benefit of database servers is that they reduce the amount of data moved between the server and the client workstation. They can also minimize data loss and prevent widespread data inconsistencies if the system fails.

   Print servers handle print requests on the LAN. By offloading the management of printing from the main LAN file server or database server, print servers help reduce the load on them and increase network efficiency in much the same way that front end processors improve the efficiency of mainframe computers. Communications servers are dedicated to performing communication processing. There are three fundamental types: fax servers, modem servers, and access servers.

   Fax servers manage a pool of fax-boards that enable LAN users to send or receive faxes. Access servers and modem servers allow users to dial into and out of the LAN by telephone. Dialing into the LAN is accomplished with an access server, whereas dialing out is accomplished with a modem server.

3. Describe the basic components of a wired LAN.

   The basic components of a wired LAN are the NICs, circuits, access points, and network operating system.
   The network interface card (NIC) allows the computer to be physically connected to the network cable, which provides the physical layer connection among the computers in the network.

   The circuits are the cables that connect devices together. In a LAN, these cables are generally twisted pair from the client to the hub or server. Outside the building, fiber optic is generally used.

   Network hubs and switches serve two purposes. First, they provide an easy way to connect network cables. In general, network cables can be directly connected by splicing two cables

together. Second, many hubs and switches act as repeaters or amplifiers. Signals can travel only so far in a network cable before they attenuate and can no longer be recognized.

The network operating system (NOS) is the software that controls the network. Every NOS provides two sets of software: one that runs on the network server(s), and one that runs on the network client(s). The server version of the NOS provides the software that performs the functions associated with the data link, network, and application layers and usually the computer's own operating system. The client version of the NOS provides the software that performs the functions associated with the data link and the network layers, and must interact with the application software and the computer's own operating system.

4. Describe the basic components of a wireless LAN.

The basic components of a wireless LAN are the NICs, circuits, access points, and network operating system.
The network interface card (NIC) allows the computer to be physically connected to the network cable, which provides the physical layer connection among the computers in the network.

The "circuit" is the air that connects the wireless clients to the access points. Between the access points and the switches or servers, twisted pair cable is typically utilized.

A wireless access point performs the same functions as a hub or switch in a wired environment.

The network operating system (NOS) is the software that controls the network. Every NOS provides two sets of software: one that runs on the network server(s), and one that runs on the network client(s). The server version of the NOS provides the software that performs the functions associated with the data link, network, and application layers and usually the computer's own operating system. The client version of the NOS provides the software that performs the functions associated with the data link and the network layers, and must interact with the application software and the computer's own operating system.

5. What types of cables are commonly used in wired LANs?

It is very common to see LANs built using traditional twisted pair cables (e.g., Cat 5, Cat 5e).

6. Compare and contrast category 5 UTP, category 5e UTP, and category 5 STP.

Category	Type	Max. Data Rate (Mbps)	Often Used By	Cost ($/foot)
5e	UTP	100	1,000Base-T Ethernet	.10
5	UTP	100	100Base-T Ethernet	.07
5	STP	100	100Base-T Ethernet	.18

7. What is a cable plan and why would you want one?

   A cable plan is a plan for the network layout, including how much cable is used, where the cables are, how many and where hubs are located, how many ports are available, what local city fire codes must be followed, and what are the identification labels of the cable. Most buildings under construction today have a separate LAN cable plan as they do for telephone cables and electrical cables. The same is true for older buildings in which new LAN cabling is being installed. It is common to install 20 to 50 percent more cable than you actually need to make future expansion simple.

   With today's explosion in LAN use, it is critical to plan for the effective installation and use of LAN cabling. The cheapest time to install network is during the construction of the building; adding cable to an existing building can cost significantly more. Indeed, the costs to install cable (i.e., paying those doing the installation and additional construction) are usually substantially more than the cost of the cable itself, making it expensive to re-install the cable if the cable plan does not meet the organization's needs.

8. What does a NOS do? What are the major software parts of a NOS?

   The network operating system (NOS) is the software that controls the network. Every NOS provides two sets of software: one that runs on the network server(s), and one that runs on the network client(s). The server version of the NOS provides the software that performs the functions associated with the data link, network, and application layers and usually the computer's own operating system. The client version of the NOS provides the software that performs the functions associated with the data link and the network layers, and must interact with the application software and the computer's own operating system.

9. How does wired Ethernet work?

   Ethernet is the most commonly used LAN in the world, accounting for almost 70 percent of all LANs. Ethernet uses a bus topology and a contention-based technique media access technique called Carrier Sense Multiple Access with Collision Detection (CSMA/CD). There are many different types of Ethernet that use different network cabling (e.g., 10Base-2, 10Base-5, 10Base-T, and 10Broad-36).

10. How does a logical topology differ from a physical topology?

    A logical topology illustrates how the network operates with the various protocols that may be running. A single network can have multiple protocols. A physical topology illustrates exactly where all the hardware and cabling are 'physically' located and connected.

11. Briefly describe how CSMA/CD works.

CSMA/CD, like all contention-based techniques, is very simple in concept: wait until the bus is free (sense for carrier) and then transmit. Computers wait until no other devices are transmitting, and then transmit their data. As long as no other computer attempts to transmit at the same time, everything is fine. However, it is possible that two computers located some distance from one another can both listen to the circuit, find it empty, and begin to simultaneously. This simultaneous transmission is called a collision. The two messages collide and destroy each other.

The solution to this is to listen while transmitting, better known as collision detection (CD). If the NIC detects any signal other than its own, it presumes that a collision has occurred, and sends a jamming signal. All computers stop transmitting and wait for the circuit to become free before trying to retransmit. The problem is that the computers which caused the collision could attempt to retransmit at the same time. To prevent this, each computer waits a random amount of time after the colliding message disappears before attempting to retransmit.

12. Explain the terms 100Base-T, 1000Base-T, 100Base-F,10GbE, and 10/100/1000 Ethernet.

Historically, the original Ethernet specification was a 10 Mbps data rate using baseband signaling on thick coaxial cable, called 10Base5 (or "Thicknet"), capable of running 500 meters between hubs. Following 10Base5 was 10Base2 or thinnet as we used to say. Thinnet or RG-58 coaxial cable, similar to what is used for cable TV was considerably cheaper and easier to work with, although it was limited to 185 meters between hubs. The 10Base-2 standard was often called "Cheapnet."

When twisted pair cabling was standardized for supporting Ethernet (app. 1988) the T replaced the 2 to represent "twisted-pair". Twisted pair is the most commonly used cable type for Ethernet. 10BaseT breaks down as 10 Mbps, baseband, and the "T" means it uses twisted pair wiring (actually unshielded twisted pair). It was the 10Base-T standard that revolutionized Ethernet, and made it the most popular type of LAN in the world.

Eventually the 10BaseT standard was improved to support Fast Ethernet or 100BaseT that breaks down as 100Mbps baseband over twisted-pair cable, and 100BaseF over fiber. This eventually was improved even further to 1000BaseT or 1 Billion BITs per second baseband. There is currently a revised standard evolving which makes Ethernet even faster. It is known as the 10GbE or 10 Billion BITs per second Ethernet. Though proven to work it has yet to reach the marketplace. But it would be astute to consider that it will be here in the near future.

Finally, 10/100Mbps Ethernet refers to the standard that can autosense which speed it needs to run at between the two speeds of 10Mbos or 100Mbps. It comes down to the type of NIC running at the individual node and the type of switch port that the node connects into. It is commonplace to run 10/100Mbps switches in LAN operating environments where there are older NICs already operating and no real business case requirements for upgrading these nodes.

13. How do Ethernet switches know where to send the frames they receive? Describe how switches gather and use this knowledge.

    Ethernet switches operate on the destination MAC address of each packet processed to determine which port to pass on each packet presented for transmission.

    Ethernet switches learn and store in memory in the form of a forwarding table, the specific port location of each MAC address for every device connected to any of its ports.

14. Compare and contrast cut-through, store and forward, and fragment-free switching.

    With cut through switching, the switch begins to transmit the incoming packet on the proper outgoing circuit as soon as it has read the destination address in the packet.

    With store and forward switching the switch does not begin transmitting the outgoing packet until it has received the entire incoming packet and has checked to make sure it contains no errors.

    Fragment-free switching lies between the extremes of cut through and store and forward switching. With fragment-free switching, the first 64 bytes and if all the header data appears correct, the switch presumes the rest of the packet is error free and begins transmitting.

15. Compare and contrast the two types of antennas.

    A directional antenna projects a signal only in one direction. Because the signal is concentrated in a narrower, focused area, it is a stronger signal and carries further. More popular is the omnidirectional antenna, which broadcasts in all directions except directly above itself.

16. How does Wi-Fi perform media access control?

    Media access control uses Carrier Sense Multiple Access with Collision Avoidance, or CSMA/CA, which is similar to the media access control used in Ethernet LANs. The computers "listen" before they transmit, and if there is not a collision, all is well. Wi-Fi does attempt to avoid a collision more than regular Ethernet LANs do, however, by using two techniques called Distributed Coordination Function and Point Coordination Function (refer to questions 12 and 13 for detailed descriptions of these two access control methods).

17. How does Wi-Fi differ from shared Ethernet in terms of topology, media access control, and error control, Ethernet frame?

    Wi-Fi is very similar to shared Ethernet in terms of the logical and physical topologies. The Wi-Fi approach uses a logical bus and a physical star arrangement, just like shared Ethernet. On the shared bus in Wi-Fi, the computers must take turns transmitting, which is not always so in shared Ethernet. For error control, Wi-Fi has a hidden node problem, where some

computers may not sense contention, and may therefore transmit when they should not, so Wi-Fi uses a slightly different technique for contention to try and cut down on collisions.

18. Explain how CSMA/CA DCF works.

    This technology relies on the ability of computers to physically listen before they transmit. With DCF, each frame in CSMA/CA is sent using stop and wait ARQ, and it is designed in such a way so that no other computer begins transmitting while the waiting period is going on.

19. Explain how CSMA/CA PCF works.

    Using PCF (also called the virtual carrier sense method), works in traditional Ethernet, and because every computer on thee shared circuit receives every transmission on the shared circuit. There can be a "hidden node problem" with CSMA/CA PCF because some computers at the edge of the network may not sense every transmission, increasing the likelihood of collisions.

20. Explain how association works in WLAN.

    Searching for an available AP is called scanning and NIC can engage in either active or passive scanning. During active scanning, a NIC transmits a special frame called probe frame on all active channels on its frequency range. When an AP receives a probe frame, it responds with a probe response that contains all the necessary information for a NIC to associate with it. A NIC can receive several probe responses from different APs. It is up to the NIC to choose with which AP to associate with. This usually depends on the speed rather than distance from an access point. Once a NIC associates with an access point they start exchanging packets over the channel that is specified by the access point.

    During passive scanning, the NIC listens on all channels for a special frame called beacon frame that is sent out by an access point. The beacon frame contains all the necessary information for a NIC to associate with it. Once a NIC detects this beacon frame it can decide to associate with it and start communication on the frequency channel set by the access point.

21. What is the best practice recommendation for wired LAN design?

    The best recommendations are based primarily on evaluating the trade-off between effective data rates and costs. Sometimes it is also interesting to evaluate LAN vs. WLAN as part of the process.

22. What are the best practice recommendations for WLAN design?

    The best recommendations are based primarily on evaluating the trade-off between effective data rates and costs. Sometimes it is also interesting to evaluate LAN vs. WLAN as part of the process.

23. What is a site survey, and why is it important?

    The site survey determines the feasibility of the desired coverage, the potential source of interference, the current locations of the wired network into which the WSAN will connect, and an estimate of the number of APs required to provide coverage.

24. How do you decide how many APs are needed and where they should be placed for best performance?

    The network manager will make a determination based off four factors: nominal data rates, error rates, efficiency of the data link layer protocols used, and efficiency of the media access control protocols.

25. How does the design of the data center differ from the design of the LANs intended to provide user access to the network?

    Data centers are designed to house significant number of servers because this is where most of the data on a network either comes from or goes to. Thus, the data center needs significant physical space and a significant amount of circuit capacity added to handle the data flow. The data center must also be built with other devices like load balancers and virtual servers, which the LAN does not have. Due to the physical space requirements and the large amount of data transferred, the design of the data center is different than that of a LAN for user access.

26. What are three special purpose devices you might find in a data center and what do they do?

    Three special purpose devices that the data center may contain include a load balancer, virtual servers, and storage area networks. The load balancer acts as a router at the front of the server farm to distribute any processing to an appropriate server. Logical servers are logically separate servers (e.g., a Web server, an email server, and a file server) on the same physical computer. The virtual servers run on the same physical computer but appear completely separate to the network. Lastly, the storage area network are LANs devoted solely to data storage.

27. What is a bottleneck and how can you locate one?

    In order to improve performance, the administrator must locate the bottleneck, the part of the network that is restricting the data flow. Generally speaking, the bottleneck will lie in one of two places. The first is the network server. In this case, the client computers have no difficulty sending requests to the network server, but the server lacks sufficient capacity to process all the requests it receives in a timely manner. The second location is the network circuit. The network server can easily process all the client requests it receives, but the network circuit lacks enough capacity to transmit all the requests to server. It is also possible that the bottleneck could also lie in the client computers themselves (e.g., they are receiving data to fast for them to process it), but this is extremely unlikely.

28. Describe three ways to improve network performance on the server.

    Improving server performance can be approached from two directions simultaneously: software and hardware.
    Software methods include changing the NOS and fine-tuning the NOS.
    Hardware methods include adding a second server and upgrading the server's hardware.

29. Describe three ways to improve network performance on circuits.

    Circuit performance can be improved by using faster technologies, by adding more circuits, and by segmenting the network into several separate LANs by adding more switches or access points.

30. Many of the wired and wireless LANs share the same or similar components (e.g., error control). Why?

    Wired and wireless LANs share the same or similar components because they are both generally based on the same Ethernet protocol. Thus, although some hardware components are different, the underlying foundation is the similar.

31. As WLANs become more powerful, what are the implications for networks of the future? Will wired LANS still be common or will we eliminate wired offices?

    Networks of the future will continue to become increasingly wireless due to the increased speed and portability wireless offers. Wired LANs will continue to be common partly due to better security and reliability. The best practice networks of the future will continue to be wired networks with added wireless capabilities.

Mini-Cases

I. Designing a New Ethernet

a) Let's assume that the smallest possible message is 64 bytes. If we use 100Base-T, how long (in feet or meters) is a 64-byte message? Hint: you can assume that the electricity in the cable travels at approximately the speed of light (186,000 miles per second).

This is a challenging question designed to get the students to think what is actually going on inside the cable. The length of a message can be calculated as its length in bits divided by the speed of the cable which results in how many nanoseconds it takes to transmit. This times the speed of light give the length in feet. Specifically:

$$\frac{64 \text{ bytes} \times 8 \text{ bits/byte}}{100 \text{ million bits per second}} = 0.0000052 \text{ seconds}$$

*0.0000052 seconds * 186,000 miles per second = .9672 miles or about 5,100 feet*

b) If we use 10 GbE, how long (in feet or meters) is a 64 -byte message?

10 GBE is 1000 times faster than 10Base-T, so the length is about 51 feet (or 15.5 meters).

c) The answer in part b is the maximum distance any single cable could run from a switch to one computer in a switched Ethernet LAN. How would you overcome the problem implied by this?

A message must be longer than the cable connecting two devices because if it is not, there could be a collision in the "middle" of the cable and neither device would detect it. With 10Base-T, we have a message that is longer than any cable we would care to make, but with 10GbE, we have a problem: cables can be no longer than 15 meters if we have messages as short as 64 bytes. This problem was addressed in the gigabit Ethernet specification by raising the minimum packet size to 512 bytes. In other words, no matter what it contains, all packets in gigabit Ethernet are a minimum of 512 bytes (or about 120 meters long), which means that the maximum cable length is 100 meters.

If data rates continue to increase, to say, 100 GbE, then we will have the same problem again, and the minimum packet size will have to increase to about 5K (which may significantly affect throughput) or we will have to change the fundamental Ethernet media access control protocol.

II. Pat's Petunias

You have been called in as a network consultant by your cousin Pat who operates a successful mail-order flower business. She is moving to a new office and wants to install a network for her telephone operators, who take phone calls and enter orders into the system. The number of operators working varies depending on the time of day and day of the week. On slow shifts, there are usually only 10 operators, whereas at peak times, here
are 50. She has bids from different companies to install (1) a shared Ethernet 100Base-T network, or (2) a switched Ethernet 100Base-T network. She wants you to give her some sense of the relative performance of the alternatives so the can compare that with their different costs. What would you recommend?

Shared 100Base-T is cheaper but all collisions would be shared by all devices connected to the network. This would have a negative effect on overall network performance. Switched 100Base-T would provide much better performance because there would be far less collisions and much greater speed.

III. Eureka!

Eureka has just leased a new office and are about to wire it. They have bids from different companies to install (a) 100Base-T network, or (c) a WI-FI network. What would you recommend? Why?

The 100BaseT solution is probably plenty for this application. The wireless network would provide good performance assuming they install enough wireless access points for the number of users. So, either option a or b would be good.

As Eureka scales to more staff members and additional services to customers it will need a robust system to accommodate the web traffic. It is in the firm's interest to develop an infrastructure that will be able to handle that growth.

IV. Tom's Home Automation

Should he continue to offer wiring services, which often cost $50 per room) or is wireless a better direction? What type of LAN would you recommend? Why?

Student answers may vary.
I would recommend that he continue to offer wired services. Although many users are wanting wireless due to the large number of wireless devices in use today, there is a lot to be said for wired connection because they always work and the speeds will be similar. For devices such as networked TVs, gaming systems (i.e. PS3, Wii, etc), and desktop systems, the wired connections are arguably better. In these cases, the cat 5 or cat 5e is adequate for the data transfer rates common to home users. I would also recommend a switch because these are more commonly sold in retail stores (as opposed to hubs). He can also recommend to the users to install wireless, because it is beneficial as well.

V. Sally's Shoes

Sally wants to network the computer with a LAN. What sort of LAN design would you recommend? Draw a picture.

Sally should use a simple, 10Base-T Ethernet or at most a 100Base-T network solution. The network can use either peer to peer or a dedicated server, but a dedicated server would be preferable for the long run. Cost will be an overriding concern, as this is probably a small business.

VI. South West State University

With a wireless network installed in the library of South West State University, how should they protect the network performance?

The network is suffering with a perfectly normal overload of users, so a good solution would be to segment the wireless LAN by adding more APs, as well as to explore more careful channel allocation to give better coverage. The university should try to accommodate the users in this fashion.

VII. Household Wireless

Is your sister's new house going to need any special wiring?

The sister will need to have a reliable and fast network for work reasons in the evening. Although a wired network might provide this, a wireless LAN would add convenience, and possibly worth the additional expense. Based on my own experiences, I suggest adding the Ethernet cable during the construction because it doesn't cost very much at that point. Then add the $40 switch and the three desktop computers are ready to be networked. The wired networks are very stable and fast.

Since they already have a laptop, I would also add an access point so that they could be wireless with the laptop anywhere in the house.

VIII. Ubiquitous Office

How many access points would you buy for the Ubiquitous Office? Where would you place them?

Depending on the bandwidth you require, the answers may vary. Using 802.11g with 54 Mbps (although they may consider 802.11n as well), you could install five access points to provide coverage throughout the office building.

As with any wireless networking, a site survey will need to be conducted with testing in each office to insure that each office has enough signal strength.

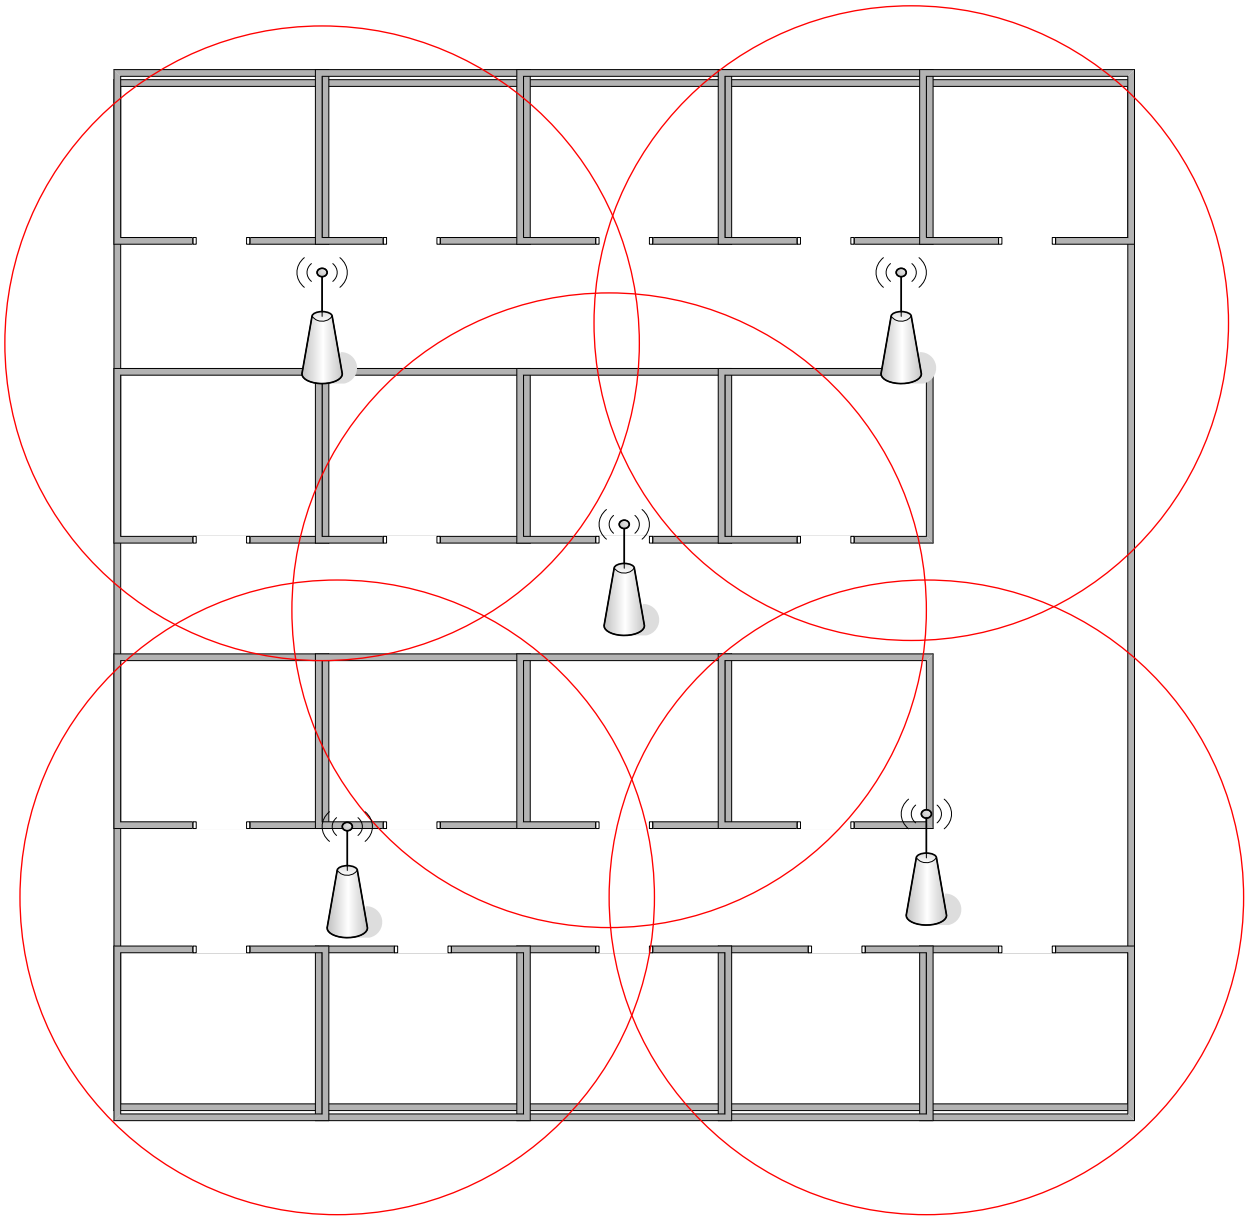

IX. ABC Warehouse

How many access points would you buy for the Ubiquitous Office? Where would you place them?

Depending on the bandwidth you require, the answers may vary. Using 802.11g with 5 Mbps, you could install 12 access points to provide coverage throughout the building. With a Z (length of one side of a square) of 60 feet for 5Mbps with obstacles, you would need two rows of six access points.

As with any wireless networking, a site survey will need to be conducted with testing in each office to insure that each office has enough signal strength.

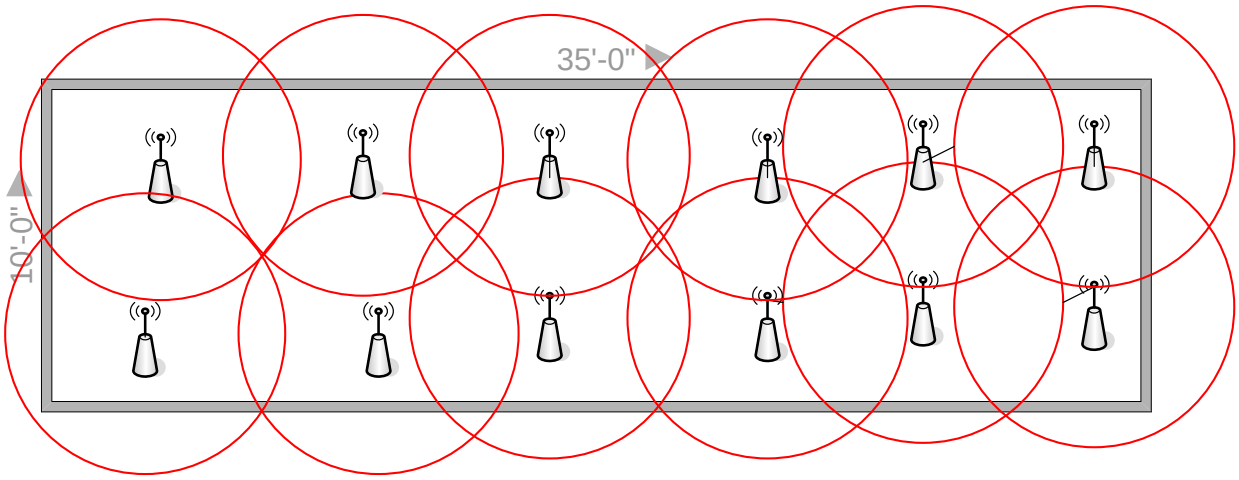

X. Metro Motel

How many access points would you buy for the Ubiquitous Office? Where would you place them?

Depending on the bandwidth you require, the answers may vary. Using 802.11g with 5 Mbps, you could install four access points per floor for a total of 16 in the building to provide coverage throughout the building. With a Z (length of one side of a square) of 60 feet for 5Mbps with obstacles, you would need one row of four access points per floor.

As with any wireless networking, a site survey will need to be conducted with testing in each office to insure that each office has enough signal strength.

Next Day Air Service Case Study

1. Which is best for the International Services Division, a dedicated-server network or peer-to-peer LAN? Explain your choice.

   The best choice for the International Services Division is a dedicated-server network. The reasons for this are:

   a. A peer-to-peer LAN is too small for a network with the goals of the International Service Division. Many remote offices may need to call in to access the LAN. This would not be possible with a peer-to-peer LAN, because the volume of traffic would overload the LAN.
   b. A peer-to-peer LAN might work initially, but it would not be suitable if NDAS continues with its growth projections.
   c. A dedicated-server network is best, because it has the flexibility to handle dial-up call inquiries today, and future growth. It can be interconnected to all of NDAS's remote offices as needed.

2. Draw a network plan using Microsoft Visio that includes the general layout of the LAN (computers, servers, cables, hubs/switches) and recommend what type of LAN to install (e.g., 10Base-2, 10Base-T, switched 10Base-T, wireless Ethernet 802.11b). Justify your recommendation.

   a. Figures for the bus, and wireless topology for a local area network should be similar to those in the text (see figure 6-6 for bus, figure 6-7 for wireless). Note in both figures that both of these appear to be star networks. But internal to the hardware circuitry it is operating like an Ethernet BUS.

   b. Any of these (Ethernet, using coaxial cable, 10base-T or wireless) would be correct, as long as the students explain their choices using reasonable logic. Reasons should include factors like cost, flexibility, growth potential, standards, reliability, and serviceability. Grade the three to five reasons on how logical they are for the NDAS situation. You can simulate an on-the-job learning experience by having the students read their reasons for the class to evaluate. In a real-life situation, they would have to justify their selection to management. The students can have a lively discussion by having to defend their choices.

3. Sally Wong has heard "horror stories" about LAN bottlenecks. Prepare a brief discussion of LAN bottlenecks and what can be done to improve LAN performance.

   a. LAN bottlenecks: Normally, bottlenecks will be in one of two places. The first is in the computers attached to the network and the second location in the network circuit. To find the bottleneck, watch the utilization of the network server during periods of poor performance. If the server utilization is high (e.g., 60-100%), then the bottleneck is the server; it cannot process all the requests it receives in a timely manner. If the server utilization is low during periods of poor performance (e.g., 10-40%), then the problem lies with the network circuit; the circuit cannot transmit requests to the server as quickly as needed. Things become more difficult if utilization is in the mid-range (e.g., 40-60%). This suggests that the bottleneck may shift between the server and circuit depending upon

the type of request, and suggests that both need to be upgraded to provide the best performance.

b. You can improve LAN performance by:

a. Improving server performance
   (1) Software
      (a) Fine tune the current NOS
      (b) Replacing the NOS with a faster NOS
      (c) Disk caching
      (d) Disk elevating
   (2) Hardware
      (a) Add servers
      (b) Faster CPU
      (c) Add memory
      (d) Faster disks
      (e) Faster network interface card
b. Improving circuit capacity
   (1) Bigger circuit
   (2) Change token ring
   (3) Segment network
c. Reduce network demand
   (1) Move files off the network
   (2) Revise schedules to allow off-peak processing

4. What safeguards do you recommend for NDAS to control the use of illegal copies of software on the LANs?

Students recommendations to safeguard against illegal copies of software on NDAS LANs should include at least the following points:

a. Not allowing any personal software to be used on the LAN. Informing all employees that using personal software is grounds for immediate dismissal.
b. Training all users about the importance of controlling illegal (unauthorized) copies of software on a local area network because
   (1) the organization may be sued for copyright infringement;
   (2) the organization may acquire a computer virus, resulting in significant monetary losses;
   (3) software publishers may go out of business if users do not pay for their software.
c. Careful monitoring by the LAN supervisor to detect any illegal software on the LAN with LAN metering software.
d. Management is committed to continuously ensure that procedures to safeguard against illegal copies of software on NDAS LANs are enforced, that software audits are conducted regularly, and whistle-blowers are protected if they report violations to management.

Additional Content

Teaching Notes

This Chapter will require 5-6 hours of class time.

My goals in this chapter are to ensure that the students understand the basic components of wired and wireless LANs and to ensure that they understand Ethernet. Many students already have a good understanding of LAN components, but many do not, so I discuss these in some detail to ensure an even understanding on the topics. But I warn those with experience that they will be bored. I also talk a bit about the need to develop a cable plant diagram and how such a tool can be used to operate and maintain a LAN.

The focus on Ethernet and switched Ethernet in this chapter is because Ethernet-based LANs have replaced token ring for the most part, although a very few organizations to a still use token ring.

The chapter contains technical detail on how the three major WLAN technologies work, but frankly it is not necessary to go overboard on teaching the technical details of this chapter. I just focus on the basics, plus the key elements of WLAN design and improving performance. Because WLANs have increased in significance and are more likely to come under the student's job descriptions in the future, it is important that they understand the connection between a wired LAN and the concept of using a WLAN in an overlay arrangement in organizations.

Improving performance is a good topic to reinforce some of the basic concepts and also can provide a practical checklist for students interviewing or entering the job market. Toward the end of the chapter I summarize some concepts in this subject matter. This chapter is an important part of the management of networks, and covers some practical aspects to managing networks.

The section comparing the different types of Ethernet is technically somewhat challenging, but is a good way to help students put together the concepts in the previous chapters and show how they shape the fundamental performance of networks.

A Technical focus box highlights the fact that Ethernet as deployed today typically does not use error control, although error control is built-in.

War Stories

In 1988 I accepted the awesome challenge of building a campus-like Local Area Network for the Bureau of Mines, Pittsburgh Energy Technology Center (PETC). The center was located in the South Hills of Pittsburgh and occupied a 252-acre campus of 28 buildings used in various aspects of health and safety research. There were a total of 394 people working at PETC, as we referred to it. At the time, the center had no means of networking beyond that of 9600bps MODEMs. We had a VAX Cluster of six Digital Equipment Corporation (DEC) processors, two IBM mainframes and three PRIME mainframe computers.

All the processing was centralized mainframe with the users either dialing in via MODEM or actually being hardwired directly into the back plane of these processors running at 9600bps. For users exceeding the distance limitations of twisted-pair copper wiring, a pair of devices known as "line drivers" were used to amplify and forward the signaling up to the processor's back plane.

There was a need to integrate newer applications running under newer processor technologies among several of the groups in the various buildings. Cost economies dictated that an Ethernet based network was needed in order to provide high-speed connectivity and support resource sharing across all buildings on the campus over the long-term.

Over the course of the next 6 years I was to play the principal role of managing the design, installation, operation and evolution of this campus-like environment to an Ethernet-based LAN running not only the systems mentioned above, but also adding on in many of the physically dispersed buildings: four SUN subnets running the SUN-OS and Solaris Operating Systems, six Novell Netware subnets running Netware versions 2.x up through the first ever Netware 4.x Server Operating System in the Pittsburgh Region and eventually the WINDOWS NT Operating System. As the network evolved users from any building on campus that was networked could be given authorized access by the Network Administrator to any of the servers or processors located in any of the other buildings networked on campus.

My staff consisted of 3 Network Technician's who completed much of the physical layer requirements. When cabling projects were too large for staff, I would contract out with nearby cabling plant installation companies. Staff also included 2 Client-Server Programmers who were responsible for installing and configuring server and client PC/workstations and network related software, 2 Computer Operators who were primarily responsible for the centralized mainframe computers, and 2 Network Specialists who were cross-trained to support the help desk and perform Network Administrator functions. As the network evolved and the end-user groups acquired their own client-server systems these groups hired people to support their end systems within their groups and to coordinate operations with the larger network.

PETC became known for establishing some of the first multi-protocol stack client PCs that could actually do many client-server applications without having to reboot and reload specific protocol types. For example on startup the client PC could load TCP/UDP/IP, IPX/SPX, IBM-3270 for connectivity to the processors or servers. We also loaded on some PCs, Xerox Network Services (XNS) and Local Area Transport (LAT) protocols for connectivity to terminal servers that connected to the back plane of the VAX processors→ for those users who for whatever reasons could not be completely connected to the evolving Ethernet. For instance, when their PC would not support installation of a NIC or they could not be rewired to connect to a HUB or Switch.

We eventually integrated Internet and a Wide Area Network interfaces into PETC's LAN. PETC became the Internet gateway site for 17 other sites nationwide. Though wireless Ethernet did not arrive prior to the end of my tenure, I'm sure I would have integrated this technology into PETC's architecture had I remained there. However, I needed to move on to start my own company in the growing areas of Networking & Telecommunications. But PETC became a notable site in the region for demonstrating the promise of Local Area Networking and client-server computing over a physically distributed environment.

War Story
Wireless Snooping
(Objective: Illustrate the dangers of low security wireless)
A friend of mine took his laptop with a wireless sniffer to a public access point in a local store. The store provided free wireless Internet access – no signing up for access, no buying daypasses or annual contracts – just plain and simple free access. Anyone could bring their laptop in and use the Internet.

My friend didn't connect into the network, but simply used the sniffer to monitor traffic flowing between the access point and the laptops that customers brought in to work on as they sipped their coffee. After a couple of hours, he had collected many thousands of packets as he sat there, watching the traffic flow. Some of the packets were encrypted, some weren't. He poked through the non-encrypted packets and found a dozen or so that contained user logins to Web sites. The web sites failed to use proper encryption for the login and thus he had usernames and passwords for supposedly private Web sites. One was a free e-mail site, so he now had access to the individual's e-mail and could read, erase or send e-mail as this person.

And to this point, he had broken no laws. Monitoring public radio broadcasts – even data packets – is not a crime. As long as you take no active measures, it is the same as listening to a radio.

At this point, he started using his decryption software to see if he could break the encrypted packets, because they likely contained the most "interesting" stuff. When he did this action, he crossed the line and started to break the law, and I asked him not to tell me any more.

Wardriving
(Objective: Illustrate the dangers of low security wireless)
A friend of mine went wardriving around our small college town (Bloomington IN) in January 2004. He found about 130 wireless access points as he drove around town. About 40% of them had no security and another 40% used SSID. Only about 20% used WEP or WPA.

He was planning to try to break into unsecured or SSID-protected networks, but there were so many that he gave up trying – it was just too simple.

