

Networks

Computer networks are everywhere. It's impossible to escape them in the modern world in which we live and work. We use them at work, at home, and even in between, in places like our cars, the park, and the coffee shop. We have come to take them for granted in the same way we treat electricity and hot water.

But a lot is going on behind the scenes when we use these networks. Cisco routers and switches play a critical role in networks' successful operation.

This opening chapter lays the foundation required to understand all the details that make networks function. Specifically, this chapter covers the following topics:

- ▶ **Describing network components**
- ▶ **Classifying networks by function**
- ▶ **Defining network architectures**

Describing Network Components

To understand how networks work, it helps to have an appreciation of why they exist in the first place. As incredible as it may seem now, for a number of years, when computers first came into use, very few computers were networked. They operated as little islands of information with no connection to one another. Data had to be transferred between computers by copying it to a floppy disk, physically taking that floppy disk to the other computer, and copying the data to the destination machine. This process is now sometimes jokingly referred to as the *sneakernet*.

Modern networks can include many components. Some of the most basic components are computers, routers, and switches. Figure 1.1 shows some Cisco routers and switches. *Routers* are used in a network to transfer information between computers that are not on the same network. Routers are

capable of doing this by maintaining a table of all networks and the routes (directions) used to locate those networks. *Switches* come in two varieties. Layer 2 switches simply connect computers or devices that are in the same network. Layer 3 switches can do that but are capable of acting as routers as well. Two models of routers are depicted in Figure 1.1, with a switch in the middle of the stack. Routers and switches are covered in depth in Chapter 10, “Network Devices.”



FIGURE 1.1 Cisco routers and switches

In this section, the benefits of networking are covered as well as the components required to constitute a network.

Defining the Benefits of Networks

There are many benefits to networks, one of which was touched on in the introduction to this section: using a network makes sharing resources possible (without putting on your sneakers and leaving your seat). When connected by networks, users can share files, folders, printers, music, movies, you name it!

If it can be put on a hard drive, it can be shared. Additional benefits are included in the following list:

Resource Sharing Resource sharing is less earthshaking at home, but in the workplace it was a key element that drove the adoption of PCs. Other computer types such as mainframe computers and dumb terminals were already in use, but were seen as specialized pieces of equipment to be used only by guys in lab coats and some other geeky types. There were other reasons for the PC revolution, but resource sharing helped to increase productivity. As an example, 10 coworkers could access a file on the network at the same time, which eliminated the time and effort spent burning, labeling, transporting, and storing 10 floppies.

Reduced Cost and Easier Installation of Software Another advantage for business that didn't become apparent as quickly as resource sharing was a reduced cost of software. Many software products are sold to organizations on a network basis. For example, instead of buying 25 retail versions of word processing software, a single copy can be purchased for the network and then a number of seat licenses can be added to the bundle. The result is a significant savings to the company.

Taking that idea a step further, the network also makes it possible to place the installation files (from the CD containing the software) on a server and to then install the software over the network (as shown in Figure 1.2). This capability relieves IT staff from having to physically visit each machine with CD in hand to perform the installation. Moreover, the software could be installed on all five machines at once over the network by using those same files.

The term *resource* is used extensively when discussing networking and simply refers to anything that a user on one computer may want to access on a different computer. Examples include files, folders, printers, and scanners.




FIGURE 1.2 Network installation

Improved Security All this peace, love, and sharing doesn't mean that everything is available to everyone. Shared resources can be secured with restrictions on who can access them and what exact access each individual possesses. So you can share a file on your computer but share it with only two of your coworkers, and not all of them. Moreover, you could allow one coworker to only read the document while the other could be allowed to edit or even delete the document.

This type of control was difficult when files were shared on floppies. After the floppy left your hand, it was out of your control. Computer networks can enforce security controls among the computers and users.

Improved Communications It's hard to even imagine today's workplace without email, instant messaging, video chatting, and conferencing, but 25 years ago, these tools did not exist. In today's world, almost *no* communication can wait for regular postal mail (this service that we once depended on is now often called *snail mail*). Even more impressive is that distance is no obstacle. You can chat online with someone in India or China as easily as with a fellow worker sitting in the next cubical!

Now think of all the paper that is being saved that used to be consumed by companies sending regular mail to one another. The problem was multiplied by the need to keep multiple copies of the documents sent through the regular mail. Email systems can be configured to maintain a copy of every email sent, and documents that used to exist in multiple physical copies now reside as a single digital copy on a server (and probably also on a backup tape).

Meetings that used to require plane trips and hotel stays can now be held online with all participants able to see one another, share documents, view slides or documents from the presenter, and even hold votes and surveys. The only consideration is time zones!

More Workplace Flexibility Users are no longer physically tied to the same computer. If resources are stored on servers, as they are in most organizations, a computer problem no longer renders a user unable to work. In a domain-based network (more on that later in this chapter in the section "Understanding Client-Server Networks"), the user can move to any other computer that is a member of the domain, access his files on the server, and continue to work while his computer is repaired or replaced.

Building on this idea, workers are increasingly telecommuting as they can use the Internet to connect to the work network and operate as if physically present in the office.

Telecommuting means working from another physical location, usually from home. It saves gas, time, and in many cases results in more productivity on the part of the worker.



Reduced Cost of Peripherals When users can share printers, scanners, and fax machines, usually fewer devices are needed. This reduces costs for the organization. Sharing these devices also offloads the responsibility for managing and maintaining these shared devices.

Centralized Administration Although not possible in a peer-to-peer network, in a domain-based network, all computer administration is centralized. This means that the LAN administrator is responsible for maintaining the security of the network, and this work is done from a special type of server called a *domain controller*. Domain controllers do more than provide security. They also serve as the directory of the resources available on the network. This is why these services are called *directory services*. (Peer-to-peer networks, domain-based networks, and LANs are explained throughout the rest of this chapter.)



Peripherals are any devices that operate in conjunction with the computer yet reside outside the computer's box. Examples include the display, mouse, keyboard, printer, camera, speakers, and scanners.

DIRECTORY ASSISTANCE, PLEASE!

Directory services, such as Active Directory by Microsoft, help users to locate files, folders, and other resources in the network.

Identifying the Requirements for a Network

A network cannot be called a network if it does not meet certain requirements. At their simplest, those requirements include the following:

- ▶ At least two computers
- ▶ A resource that needs to be shared
- ▶ A transmission medium
- ▶ A communications agreement

Each requirement is detailed in the following list. The coverage of the last two bullet points is somewhat brief as transmission mediums are discussed in Chapter 9, "Cabling," and protocols (communications agreements) are covered in detail in Chapter 4, "Protocols."

At Least Two Computers It seems obvious, but if there are not at least two computers, there is no need for a network. A single computer doesn't need a network to access the information on its own hard drive. Getting information

from computer A to computer B without using the sneakernet is what drove the development of networks.

A Resource That Can Be Shared You already know from our earlier discussion that resources are anything that needs to be shared. This can include physical entities such as printers and scanners, or it can be files and folders located on another computer, as shown in Figure 1.3. If it can be shared and moved from one computer to another, it can be considered a resource.




FIGURE 1.3 Sharing resources

A Transmission Medium Some form of communications medium is also required. The most common form is a cable, but wireless communications are becoming increasingly widespread because of certain advantages to this approach. Both methods are shown in Figure 1.4.




FIGURE 1.4 Transmission mediums

MEDIUM? DO I NEED A OUIJA BOARD?

A communications *medium* is any process that can be used by two computers to transfer data. It can be bounded (via a cable) or boundless (wireless).

A Communications Agreement One of the main stumbling blocks present when computers were first being networked was a language problem. As you know, two people who need to converse cannot do so unless they speak a common language. Likewise, computers have to be speaking the same language in order to have a communications agreement. Networking languages are called *protocols*. In Figure 1.5, workstation 2 is able to communicate with workstation 3 because they are both using TCP/IP, but cannot communicate with workstation 1 because it is using IPX/SPX, a different networking protocol.

Protocols are discussed in Chapter 4.

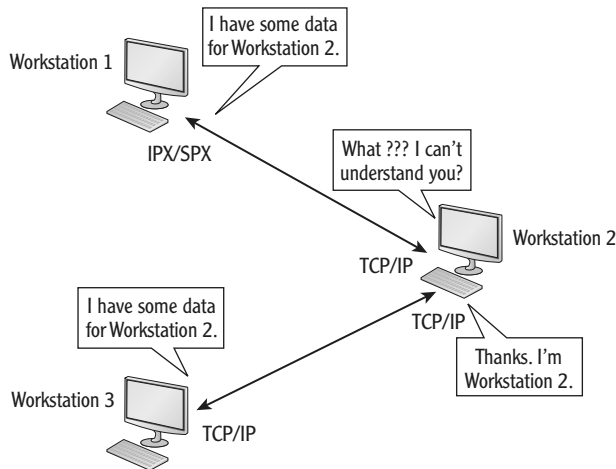


FIGURE 1.5 Protocol agreement

Before the standardization of network protocols, brought about by the explosion of the Internet and the introduction of reference models such as the OSI and the DoD models, computers from different vendors could not be networked together because they used proprietary and dissimilar network protocols.

In addition to the minimum requirements for a network, additional components are usually present in varying combinations. *Repeaters* are devices designed to regenerate or strengthen transmission signals to avoid attenuation

The OSI and DoD network models are covered in Chapter 2, "The OSI Model."

or weakening of the signal, which leads to data corruption. *Hubs* are junction boxes with no intelligence that are used to connect devices together on the same physical network. Switches can act as hubs but provide vastly improved performance and offer additional functions not available in hubs. Routers, as discussed earlier, are used to connect networks and allow computers located on different networks to communicate. Cisco routers and switches are intelligent because of the Cisco Internetwork Operating System (IOS), which is included in and is used to manage the functions of these products. The Cisco IOS is discussed in Chapter 12, “Managing the Cisco IOS.” Routers, switches, and hubs are covered in detail in Chapter 10.

PROPRIETARY VS. STANDARD

The term *proprietary*, used often in the IT world, refers to any process or way of doing something that works only on a single vendor’s equipment. The opposite of this is a *standard*, which is any way of carrying out a function that the industry has agreed upon. An everyday example of a standard is the ubiquitous wall socket. A standard was developed so that consumers could be assured that any electrical device would match this standard outlet type.

As the next few chapters unfold, you will gain new perspectives about these requirements as you learn more about the details of each. Now let’s look at some characteristics of various types of networks.

Classifying Networks by Function

Networks can be classified according to a number of different characteristics. They can differ based on location, and they can differ in the security relationship that the computers have with another. These are not the only ways networks can differ, but they are commonly used distinctions. In this section, the distance factor is examined in a discussion of LANs and WANs. After examining LANs and WANs, you will take a closer look at defining networks by security relationships in the “Defining Network Architectures” section.

Understanding LANs

If you survey networking books, you will find that the distinction between a local area network (LAN) and a wide area network (WAN) differs from one text to the next. In some treatments of this subject, the difference lies in physical

location, while in others the distinction is discussed in terms of the speed of the connection. Because this text is designed to prepare you to manage Cisco routers and switches, a Cisco perspective is appropriate.

Cisco defines a *LAN* as a high-speed data network covering a small geographical area. For the purposes of this discussion, a LAN is a single physical location, which could be a part of a building, an entire building, or a complex of buildings.

In the vast majority of cases, the network will use a networking technology called Ethernet. Other technologies do exist (such as one called Token Ring), but Ethernet has become the de facto standard technology that is used for connecting LANs.

Ethernet is discussed in more detail in Chapter 2 and Chapter 5, “Physical and Logical Topologies.”

STANDARDS

As stated earlier in this chapter, a standard is an agreed upon way of doing things. In the networking world, there are two types: official and de facto. An *official standard* is one that all parties agree to and is usually adopted by a body formed to create standards, such as the International Organization for Standardization (ISO). A *de facto standard*, on the other hand, is one that becomes the standard simply by being the method that all parties gradually choose to use over a period of time, without a formal adoption process.

Ethernet networks are typically built, owned, and managed by an organization. It is impractical for the organization to connect offices in two cities with Ethernet cabling (for many reasons that will be discussed later, one of which is a limit on cable length of about 100 ft.).

In a LAN, all of the computers are connected with a high-speed connection. *High speed* is a relative term, but in this case, it indicates at least 10 Mbps. In most cases today, the connection will be either 100 Mbps or 1,000 Mbps. The location may contain multiple buildings; it could even be an entire complex, but if the buildings are connected with a high-speed connection, they would still collectively be considered a single LAN.

Cables are discussed in Chapter 9.

Understanding WANs

A *wide area network (WAN)* is a collection of LANs connected to one another with a WAN technology or with the Internet, allowing it to function as one large network. In the previous section, the impracticality of a company strung together by private Ethernet lines from one office to another was mentioned. Above and beyond the cable length issue, there would be issues of where to place the cables and how to maintain them.

The solutions that are available are as follows:

- ▶ Leasing a WAN connection from a telecommunications company
- ▶ Using the Internet

When a WAN connection is leased from a telecommunications provider, the company offloads all maintenance and simply uses the existing network that the telecommunication provider built. The advantage to this approach is that your connection is dedicated, meaning there is no other traffic on it. WAN technologies do not use Ethernet. There are a variety of WAN connection types, such as Frame Relay, Integrated Services Digital Network (ISDN), and Point-to-Point Protocol (PPP), and each has advantages and disadvantages.

Another available option is to use the Internet. When this approach is taken, the company creates a logical connection called a *virtual private network (VPN)* between the offices by using the Internet as the physical medium. It is called *private* because the information that crosses the Internet from one office to another is typically encrypted so that if it is intercepted, it cannot be read.

Regardless of the underlying details, a WAN is used to connect LANs. The relationship between the two network types is illustrated in Figure 1.6. The figure depicts three LANs in different cities using the wide area connection to form a WAN.




FIGURE 1.6 A wide area network (WAN)

▶
WAN technologies are beyond the scope of this book. For more information, simply search for WAN methods on the Internet.

Defining Network Architectures

The *architecture* (or structure) of a network can be discussed from both a physical and a logical viewpoint. For example, in the previous section you looked at how distance can be used to differentiate networks into architectures called LANs and WANs. The architecture of a network can also describe the rules and processes used on the network. The security relationships that exist among the computers on the network can define different architectures. In this section, the difference between peer-to-peer and client-server architectures is explored.

Understanding Peer-to-Peer Networks

Peer-to-peer networks were the first type of networks to appear. This type of network is often referred to as a *workgroup*. In a peer-to-peer network, each computer is in charge of its own security, and the computers have no security relationship with one another. This does *not* mean that the users on the computers cannot share resources; otherwise, it wouldn't be a network!

There are certain shortcomings to this paradigm. In a workgroup, a user can access resources on another computer only if that user has an account on the computer where the resource resides. Moreover, depending on how the sharing is set up, she may also have to identify herself and provide a password to access the resource.

The ramifications of this can be illustrated with an example. Suppose you have four computers in an office that are used by four different users. If your goal is to allow all users to access resources located on all four computers, you would have to create an account for each person on all four computers. That means you would be creating 16 accounts in all (4 computers \times 4 people). That's a lot of work! (I guess it's a form of job security!)

Figure 1.7 illustrates this situation. Each computer is named after its user, and as you can see, all users must have an account on all computers. Also note each user can be given different levels of access. Note that the passwords that a user has been assigned on any two computers have no relationship to each other. A user can have the same password on all computers, or a different password on each computer, with no effect on functionality because they are not related to each other in any way in a peer-to-peer network.

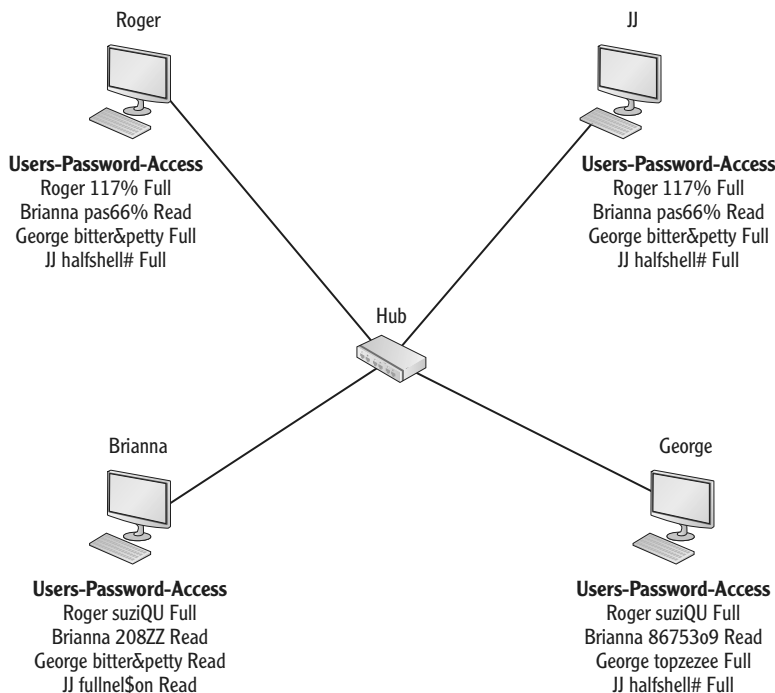


FIGURE 1.7 Peer-to-peer architecture

Another challenge with workgroups is that after the number of computers gets beyond 10, two problems occur. First, because of the nature of the communication process that occurs in a workgroup, traffic overwhelms the physical infrastructure, and the network gets very slow. This occurs because in order to locate each other, the computers must broadcast for one another. A broadcast is akin to a person calling out in a crowded room, “Who is Joe?” Then, when Joe answers, you send him the data. In Figure 1.8, workstation 10 is seeking to connect to a computer named Bannarama, so a broadcast is sent out to every computer. Then Bannarama answers with its IP address.

Moreover, unlike humans, the computers can remember who is who for only a minute or so, and then they must broadcast again.

The second problem that occurs when more than 10 computers are present in a peer-to-peer network has to do with the design of client operating systems. Most client operating systems (meaning any operating system that is not a server operating system) can host only 10 concurrent connections from other computers at a time. So if a popular file is located on a computer in a workgroup, and 10 computers are already connected, the 11th computer won’t be able to access the resource until a computer disconnects!

An *IP address* is a number in a specific format that is used to identify a computer. This topic is covered in detail in Chapter 7, “Classful IP Addressing.”

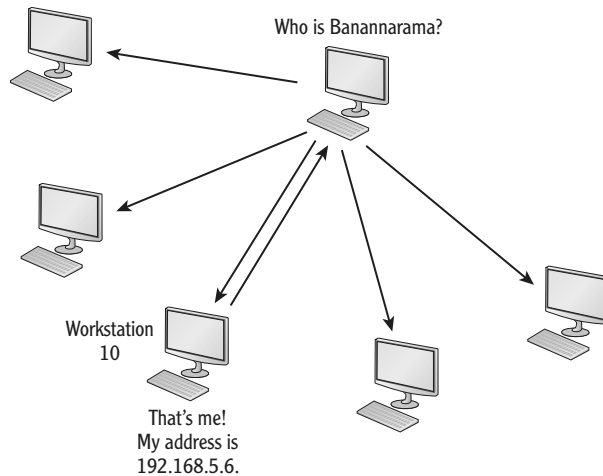


FIGURE 1.8 Broadcasting

Workgroups still have their place and their advantages. One is low cost when compared with a client-server network. Obviously, no servers (which cost more than client computers) need to be purchased. Workgroups are also quite simple to set up when compared with client-server networks. Home networks are usually peer-to-peer, and many small office and home office (SOHO) networks function well as workgroups.

However, in medium to large networks, the management of security becomes an administrative nightmare. As discussed earlier, each user must have an account on every computer that he will use or access over the network. Also, peer-to-peer networks are not scalable. When a network can be grown (with respect to the number of computers) without causing additional network traffic or additional administrative effort, it is said to be scalable.

In summary, the advantages of a peer-to-peer network are as follows:

- ▶ Low cost
- ▶ Easy to set up
- ▶ No server required

The disadvantages of a peer-to-peer network are as follows:

- ▶ No centralized control of security
- ▶ Administrative burden of maintaining accounts on all computers
- ▶ Not scalable

Understanding Client-Server Networks

The most obvious difference between a client-server network and a peer-to-peer network is the presence of at least one server. This brings up an issue that needs to be addressed before you encounter it. There are two explanations of a *client-server network* that are commonly used. Both are applicable, so let's cover both.

First, a client-server network can be explained in terms of resource access. When viewed from this perspective, it means that the shared data is centralized on a device called a file server.

WHAT'S THE DIFFERENCE BETWEEN A CLIENT AND A SERVER, ANYWAY?

Which computer is the client and which is the server is simply a matter of perspective. If the computer is seeking to access a resource on another computer, it is acting as a *client*. If it possesses a resource that another computer accesses, it is acting as a *server*. Consequently, computers in a peer-to-peer network will be acting as either at various times, depending on whether they are accessing a resource or allowing access to a resource.

A directory server or domain controller maintains the location of all resources in the network (including the computers themselves) and the locations of each. The computers in the network use this server to find things. Instead of broadcasting to find resources, the computers check with the directory server, which results in a great reduction of traffic!

A *file server* is a computer that contains resources (files) that users in the network need. A server's operating system is designed differently than one that will be used on client computers. It is *not* bound by a limit to the number of connections. Hundreds of computers can connect. The advantage is that the security surrounding the resources can be centralized on that server.

Using our example from Figure 1.7, if there was a file server in that network, we would not have to create an account for every user on all computers. We would have to do that only one time, on the server where the resources are located.

The other explanation of a client-server network takes this a step further. These networks are sometimes called *domain-based networks*. In this case, the server is a special type of server called a *directory server* or *domain controller*.

The domain controller creates a group security association between the computers that are members of what is commonly called a *domain* (or a *realm* in Unix). After a user is made a member of the domain, the user will have two types of user accounts: a local account on her computer, as she had in the peer-to-peer network, and a domain account. The domain account will be created on the domain controller where it will be stored.

This domain account will allow the user to log into the domain from any computer that is a member of the domain. This simplifies the account creation process in the same way illustrated in the explanation of using a file server. The accounts are created one time on the domain controller, and then the account will work on any computer in the domain.

The domain controller, rather than the individual computers, is responsible for validating the credentials of users. Whenever a user logs into the domain from a member computer, the login request is sent to the domain controller, which verifies the name and password and then sends the user an access token. An *access token* is a file that lists the resources that the user is allowed to access in the network, regardless of where the resource is located.

The benefit of this security paradigm is a feature called *single sign-on*. After logging into the domain, a user will not be prompted for a password again, even when accessing resources. It doesn't even matter which computer the resource is on!

On other hand, there are disadvantages to implementing a client-server network. The hardware and software required to deploy servers is significantly more expensive than client software found in a peer-to-peer network. Configuring and maintaining these servers also requires a much higher degree of skill.


Moreover, when a single domain controller is in use, a single point of failure has been introduced to the operation of the network. If something happens to the domain controller, such as a hardware failure, all access to resources can be interrupted. For these reasons, most networks deploy multiple domain controllers to eliminate this single point of failure, further adding to the cost of deploying a client-server network.

In summary, these are the advantages of a client-server network:

- ▶ Centralized administration
- ▶ Single sign-on
- ▶ Reduced broadcast traffic
- ▶ Scalability

Disadvantages of a client-server network are as follows:

- ▶ Higher cost for server software and hardware
- ▶ More challenging technically to implement
- ▶ Single point of failure with a single domain controller or single file server



Scalability means that the network can grow without the congestion problems that arise when a peer-to-peer network grows larger.

Figure 1.9 compares the peer-to-peer and client-server networks.

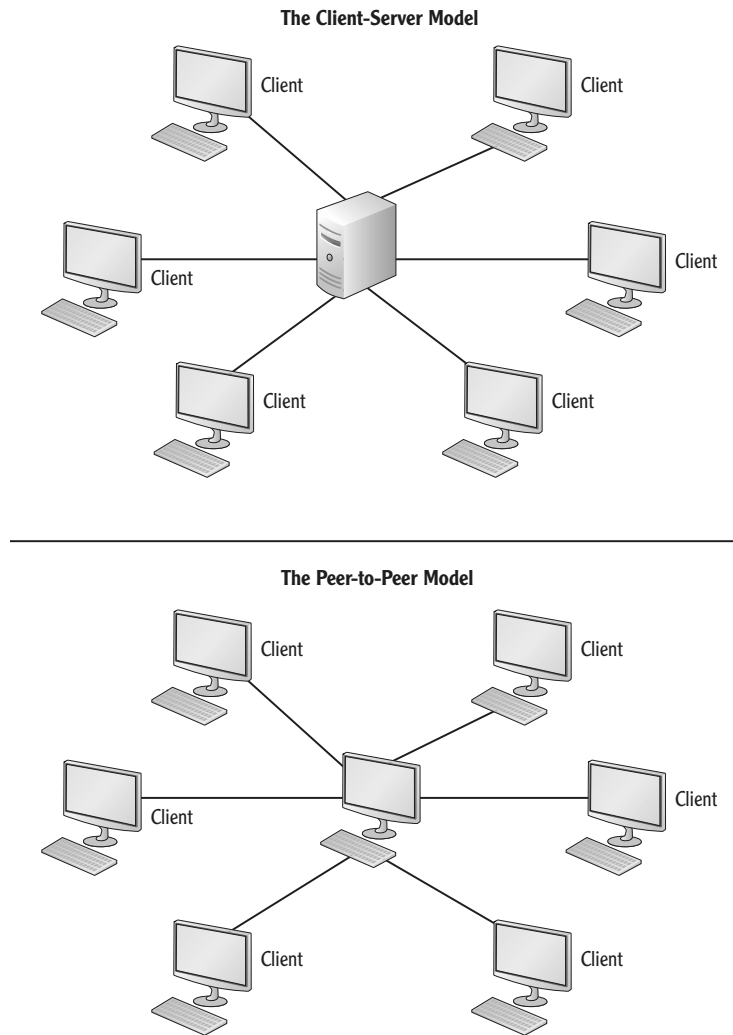


FIGURE 1.9 The client-server model (top) and the peer-to-peer model (bottom)

THE ESSENTIALS AND BEYOND

Networks allow computers to communicate and share resources. At their simplest, the requirements are two computers connected by communications media sharing a resource. The advantages of networks are resource sharing, lower software and peripheral costs in the enterprise, workplace flexibility, improved communications and security, and centralized administration.

A LAN is a network of computers connected with a high-speed connection and located in one physical location. A WAN is a group of geographically distributed LANs joined by a WAN connection. A LAN can be either a peer-to-peer network or a client-server network. Resource access and security are distributed in a peer-to-peer network, while both are centralized in a client-server network.

ADDITIONAL EXERCISES

You are a consultant specializing in network design. Consider the following scenarios and propose a design using the principles discussed in this chapter (LAN, WAN, peer-to-peer, client-server). Be prepared to discuss and defend your answer.

- ▶ An auto parts chain with 75 locations in five states
- ▶ A doctor's office with three computers
- ▶ A call center in which the users work in three shifts using a single set of computers

REVIEW QUESTIONS

1. Which of the following is *not* an advantage of networking computers?
 - A. Resource sharing
 - B. Reduced security for data
 - C. Potential for increased productivity
 - D. Improved communications
2. A _____ server is one that forms a security association between network members and helps to locate resources.
 - A. File
 - B. Directory services
 - C. Security controller
 - D. Network browser
3. What is the minimum number of computers required to form a network?
 - A. One
 - B. Two
 - C. Three
 - D. Four
4. True or False: Telecommuting is when a user works from another physical location.

(Continues)

THE ESSENTIALS AND BEYOND *(Continued)*

5. What is a protocol?
 - A. A type of transmission medium
 - B. A security agreement
 - C. A communications agreement
 - D. A suggested best practice
6. _____ refers to any process or way of doing something that works only on a single vendor's equipment.
 - A. Proprietary
 - B. Standard
 - C. De facto
 - D. Registered
7. Which statement is true with regard to a LAN?
 - A. Distributed across a large geographical area
 - B. High speed
 - C. Leased from a telecommunications company
 - D. Requires a server
8. True or False: A de facto standard is one that all parties agree to and is usually adopted by a body formed to create standards.
9. A peer-to-peer network is also sometimes called a _____ .
 - A. Realm
 - B. Domain
 - C. Workgroup
 - D. Organizational unit
10. Which of the following are shortcomings of a peer-to-peer network?
 - A. Difficult to implement
 - B. Requires server
 - C. High cost
 - D. Network congestion