chapter 5

Visible Networks

"Information at your fingertips."

—Bill Gates

In this chapter, you will learn how to

- **Describe the basic functions of a network, including identifying common devices and connectors**
- **Discuss the differences between a LAN and a WAN and the importance of TCP/IP**
- **Perform basic resource sharing**

It's hard to find a computer that's not connected to a network. Whether your system is part of a large enterprise network or a single PC with an Internet connection, every computer has some form of network connection. CompTIA has wisely added quite a bit of networking coverage to the CompTIA A+ exams, and a number of chapters toward the back of this book cover networking in great detail. By covering the basic or "user level" networking topics early in the book, however, you'll see how networking impacts every facet of modern computing.

I didn't name this chapter "Visible Networks" simply as a clever follow-up to the titles of the previous chapters. The primary goal of this chapter is to cover the more visible parts of the network: the connections and settings that enable you and me to set up simple networks. Let's begin by answering a big question: Why do we network?

Historical/Conceptual

Take a moment to think about what you do on a network. Most of us, when asked, would say, "surf the Internet," or "watch YouTube videos," or maybe "print to the printer downstairs." These are all good reasons to use a network, but what ties them together? In each of these situations, you are using your computer (the local computer) to access "stuff" stored on a remote computer (not your local computer). So what do remote computers have that you might want (see Figure 5.1)?

A remote computer called a **Web server** stores the files that make up a Web site. The Web server uses server programs to store and share the data. The two most famous Web server programs are Apache HTTP Server and Internet Information Server (IIS). When you access a Web site, your **Web browser** (likely Internet Explorer, Mozilla Firefox, or Google Chrome) asks the Web server to share the Web page files and then displays them (see Figure 5.2). Because your computer asks for the Web page, we call it the **client**. The remote computer that serves the Web site is a **server**.

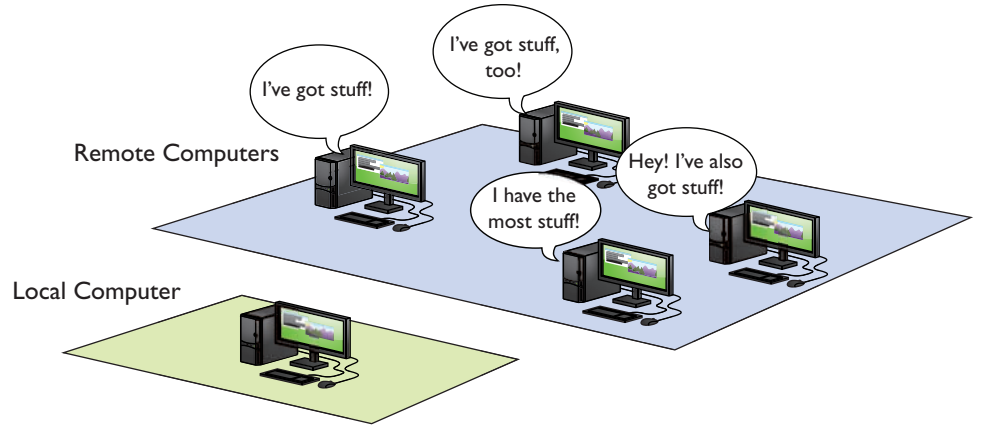

• **Figure 5.1** Accessing remote computers

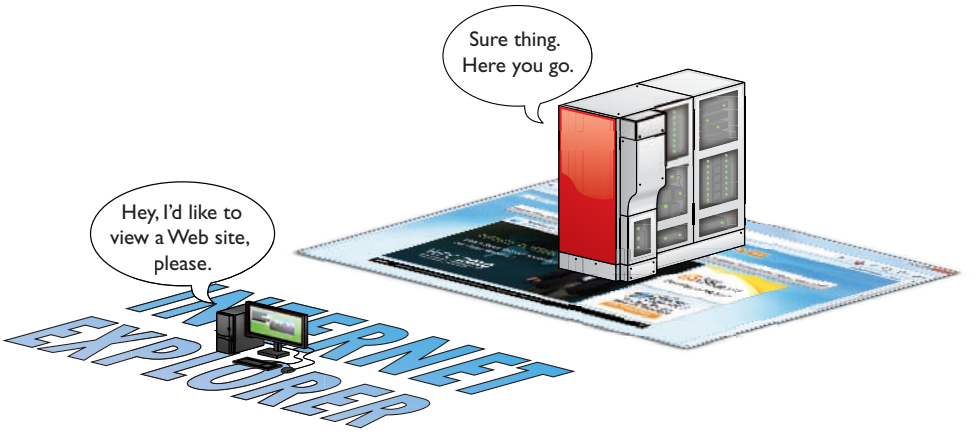

• **Figure 5.2** Accessing a Web page

Any computer that's running a sharing program is by definition a server.

But what about YouTube? YouTube also uses Web servers, but these Web servers connect to massive video databases. Like a normal Web server, these remote computers share the videos with your client PC, but they use special software capable of sending video fast enough that you can watch it without waiting (see Figure 5.3).

But we don't need the Internet to share stuff. Figure 5.4 shows a small home network with each computer running Windows 7. One of the computers on the network has a printer connected via a USB port. This computer has enabled a printer-sharing program built into Windows so that the other computers on the network can use the printer.

No matter how big the network, we use networks to share stuff. This stuff might be Web pages, videos, printers, folders, e-mail messages, music—what you can share is limited only by your ability to find a server program capable of sharing it and a client program that can access it. Network people

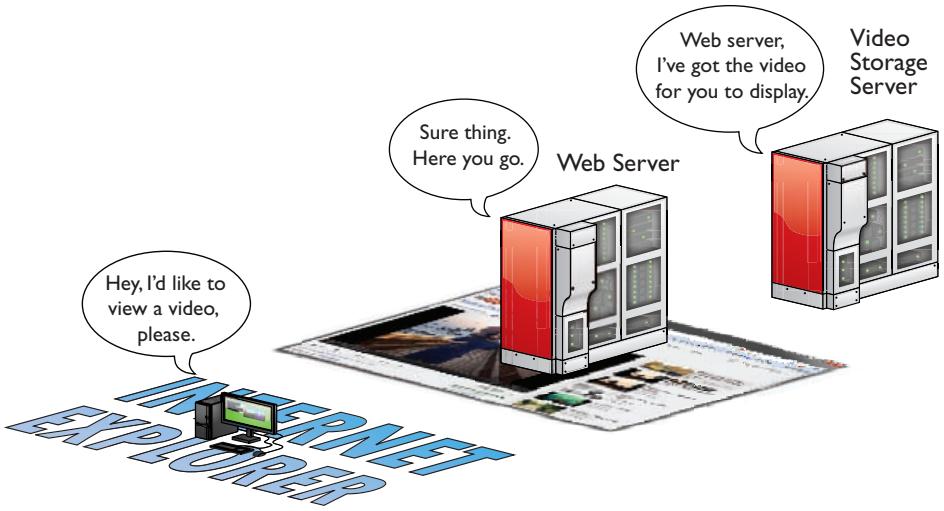

Figure 5.3 Accessing a YouTube page

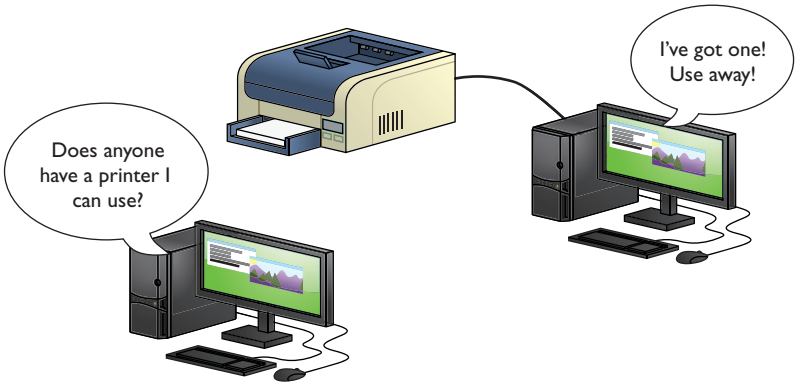

Figure 5.4 Sharing a printer in Windows 7

call anything that one computer might share with another a **resource**. Therefore, the goal of networking is to connect computers so that they can share resources or access other shared resources.

To share and access resources, a network must have the following:

1. Something that defines and standardizes the design and operation of cabling, network cards, and the interconnection of multiple computers
2. An addressing method that enables clients to find servers and enables servers to send data to clients, no matter the size of the network
3. Some method of sharing resources and accessing those shared resources

Networking Technologies

When the first network designers sat down at a café to figure out how to get two or more PCs to share data and peripherals, they had to write a lot of notes on little white napkins to answer even the most basic questions. The first question was: *How*? It's easy to say, "Well, just run a wire between them!" But that doesn't tell us how the wire works or how the computers connect to the wire. Here are some more big-picture questions:

- How will each computer be identified?
- If two or more computers want to talk at the same time, how do you ensure that all conversations are understood?
- What kind of wire? What gauge? How many wires in the cable? Which wires do what? How long can the cable be? What type of connectors?

Clearly, making a modern PC network entails a lot more than just stringing up some cable! As you saw a bit earlier, most networks have one or more *client machines*, PCs that request information or services, and a *server*, the machine that hosts and shares the data. Both clients and servers need **network interface controllers (NICs)** that define or label the machine on the network. A NIC also breaks files into smaller data units, called **frames**, to send across the network and reassembles the frames it receives into whole files. You also need some medium for delivering the frames between two or more PCs—most often this is a wire that can carry electrical pulses; sometimes it's radio waves or other wireless methods. Finally, your PC's operating system has to be able to communicate with its own networking hardware and with other machines on the network. Figure 5.5 shows a typical network layout.

This section of the chapter looks at the inventive solutions network engineers found to handle frames and cabling. After a brief look at the core technology, the chapter dives into four specific types of networks. You'll dig into the software side of things later in the chapter.

Not too many years ago, every NIC came on an expansion card that you added to a motherboard. Most techs called that card a *network interface card* or *NIC*. Now that just about every motherboard has the networking feature built in, the acronym has shifted to network interface *controller*. You're likely only to see the term *NIC* on the exams, though I call them *network cards*, too.

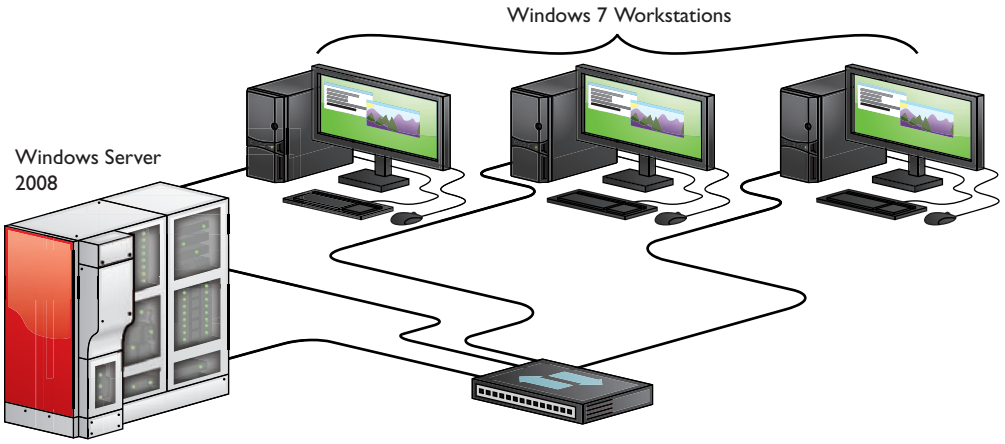

• **Figure 5.5** A typical network

Topology

If a bunch of computers connect together to make a network, some logic or order must influence the way they connect. Perhaps each computer connects to a single main line that snakes around the office. Each computer might have its own cable, with all of the cables coming together at a central point. Or maybe all of the cables from all of the computers connect to a main loop that moves data along a track, picking up and dropping off data like a circular subway line.

A network's **topology** describes the way that computers connect to each other in that network. The most common network topologies are called *bus*, *ring*, *star*, and *mesh*. Figure 5.6 shows the four types: a **bus** topology, where all computers connect to the network via a main line called a *bus cable*; a **ring** topology, where all computers on the network attach to a central ring of cable; a **star** topology, where the computers on the network connect to a central wiring point (usually called a *switch*); and a **mesh** topology, where each computer has a dedicated line to every other computer. There are also **hybrid** topologies, such as star bus or star ring, which combine aspects of the other topologies to capitalize on their strengths and minimize their weaknesses. You'll look at the most important hybrid topology, star bus, in a moment, but for now, make sure you know the four main topologies!

Look at Figure 5.6. A mesh topology looks amazingly resilient and robust, doesn't it? And it is, at least on paper. Every computer physically connects to every other computer on the network, so even if half of the PCs crash, the network functions as well as ever (for the survivors). In a practical sense, however, implementing a mesh topology network would be an expensive mess. For example, a tiny network with only 10 PCs would need *45* distinct pieces of cable to connect every PC to every other PC. What a mesh mess! Because of this, mesh topologies have never been practical in a cabled network.

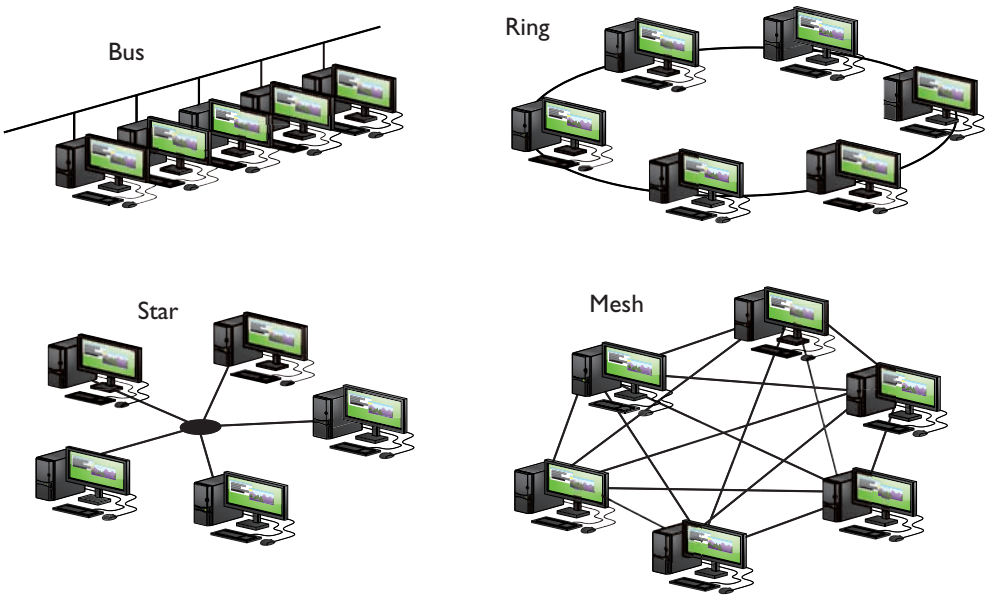

Figure 5.6 Clockwise from top left: bus, ring, mesh, and star topologies

Although a topology describes the method by which systems in a network connect, the topology alone doesn't describe all of the features necessary to make a cabling system work. The term *bus topology*, for example, describes a network that consists of some number of machines connected to the network via the same piece of cable. Notice that this definition leaves a lot of questions unanswered. What is the cable made of? How long can the cable be? How do the machines decide which machine should send data at a specific moment? A network based on a bus topology can answer these questions in a number of different ways.

Most techs make a clear distinction between the *logical topology* of a network—how the network is laid out on paper, with nice straight lines and boxes, similar to an electronic schematic—and the physical topology. The *physical topology* describes the typically messy computer network, with cables running diagonally through the ceiling space or snaking their way through walls. If someone describes the topology of a particular network, make sure you understand whether they're talking about the logical topology or the physical topology.

Over the years, manufacturers and standards bodies created several specific network technologies based on different topologies. A **network technology** is a practical application of a topology and other critical standards to provide a method to get data from one computer to another on a network. It defines many aspects of a network, from the topology, to the frame type, to the cabling and connectors used. A network technology defines everything necessary to get data from one computer to another.

801

Frames and NICs

Data is moved from one PC to another in discrete chunks called *frames*. You'll sometimes hear the word *packet* used instead of frames—this is incorrect. Packets are a part of the frame. You'll find more information about packets in Chapter 22.

Every NIC in the world has a built-in identifier, an address unique to that single network card, called a **media access control (MAC) address**. You read that right—every network card in the world has its own unique MAC address!

A MAC address is a *binary number*, meaning it's a string of 1s and 0s. Each 1 or 0 is called a *bit*. (You'll learn more about binary in Chapter 6.)

The MAC address is 48 bits long, providing more than 281 *trillion* MAC addresses, so there are plenty of MAC addresses to go around. Because people have trouble keeping track of that many 1s and 0s, we need another way to display the addresses. *Hexadecimal* is shorthand for representing strings of 1s and 0s. One hex character is used to represent four binary characters. Here's the key:

0000 = 0	0100 = 4	1000 = 8	1100 = C
0001 = 1	0101 = 5	1001 = 9	1101 = D
0010 = 2	0110 = 6	1010 = A	1110 = E
0011 = 3	0111 = 7	1011 = B	1111 = F

Even though MAC addresses are embedded into the NIC, some NICs allow you to change the MAC address on the NIC. This is rarely done.

So, MAC addresses may be binary, but we represent them by using 12 hexadecimal characters. These MAC addresses are burned into every NIC, and some NIC makers print the MAC address on the card. Figure 5.7 shows the System Information utility description of a NIC, with the MAC address highlighted.

Hey! I thought we were talking about frames! Well, we are, but you need to understand MAC addresses to understand frames. The many varieties of frames share common features (see Figure 5.8). First, frames contain the MAC address of the network card to which the data is being sent. Second, they have the MAC address of the network card that sent the data. Third is the data itself (at this point, we have no idea what the data is—certain software handles that question), which can vary in size depending on the type of frame. Finally, the frame must contain some type of data check to verify that the data was received in good order. Most frames use a

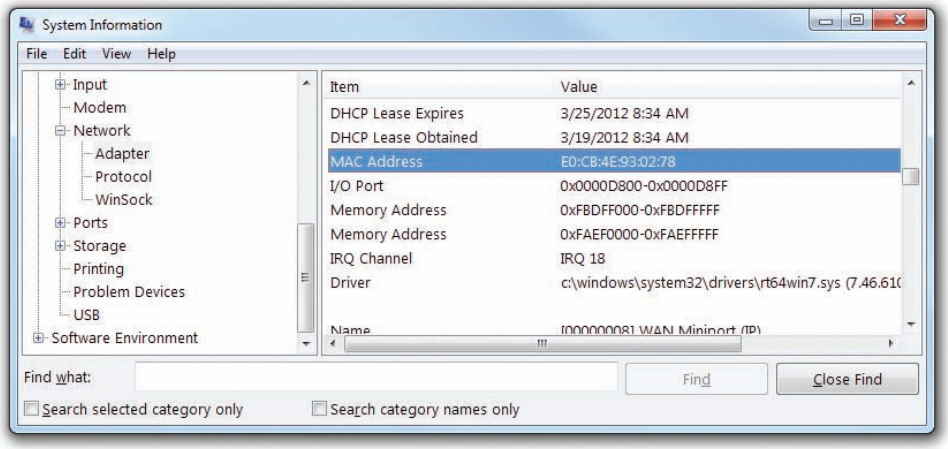

• **Figure 5.7** MAC address

clever mathematical algorithm called a **cyclic redundancy check (CRC)**.

This discussion of frames raises the question, how big is a frame? Or more specifically, how much data do you put into each frame? How do you ensure that the receiving PC understands the *way* the data was broken down by the sending machine and can thus put the pieces back together? The problem in answering these questions is that they encompass so many items. When the first networks were created, *everything* from the frames to the connectors to the type of cable had to be invented from scratch.

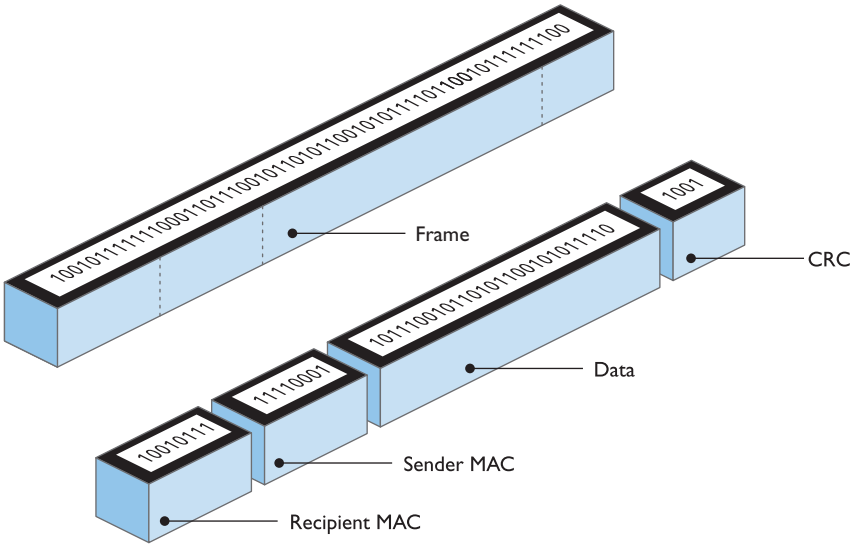

Figure 5.8 Generic frame

To make a successful network, you need the sending and receiving PCs to use the same network technology. Over the years, many hardware protocols have been implemented, with such names as Token Ring, FDDI, and ARCnet, but today only one hardware protocol dominates the modern PC computing landscape: *Ethernet*.

Introducing Ethernet

A consortium of companies centered on Digital Equipment Corporation, Intel, and Xerox invented the first network in the mid-1970s. More than just create a network, they wrote a series of standards that defined everything necessary to get data from one computer to another. This series of standards was called **Ethernet**. Over the years, Ethernet has gotten faster and has used different types of cabling. As a result, the Ethernet folks have a number of Ethernet versions, often called *Ethernet flavors*. Even though there are various speeds and cable types, all flavors of Ethernet use the same frame. This is very important: you can have any combination of hardware devices and cabling using different Ethernet flavors on a single Ethernet network and, in most cases, the PCs will be able to communicate just fine.

Most modern Ethernet networks employ one of three speeds (and sometimes all three), **10BaseT**, **100BaseT**, or **1000BaseT**. As the numbers in the names suggest, 10BaseT networks run at 10 Mbps, 100BaseT networks run at 100 Mbps, and 1000BaseT networks—called Gigabit Ethernet—run at 1000 Mbps, or 1 Gbps. All three technologies—sometimes referred to collectively as *10/100/1000BaseT* or just plain Ethernet—use a *star bus* topology and connect via a type of cable called *unshielded twisted pair (UTP)*.

Ethernet developers continue to refine the technology. 1000BaseT might be the most common standard now, but 10-Gigabit Ethernet is starting to make inroads today.

Star Bus

Imagine taking a bus network (where every computer connects to a common wire) and shrinking the bus down so it fits inside a box. Then, instead of attaching each PC directly to the wire, you attach them via cables to special ports on the box (see Figure 5.9). The box with the bus takes care of all of the tedious details required by a bus network. The bus topology would look a lot like a star topology, wouldn't it? Modern Ethernet networks all use **star bus** topology.

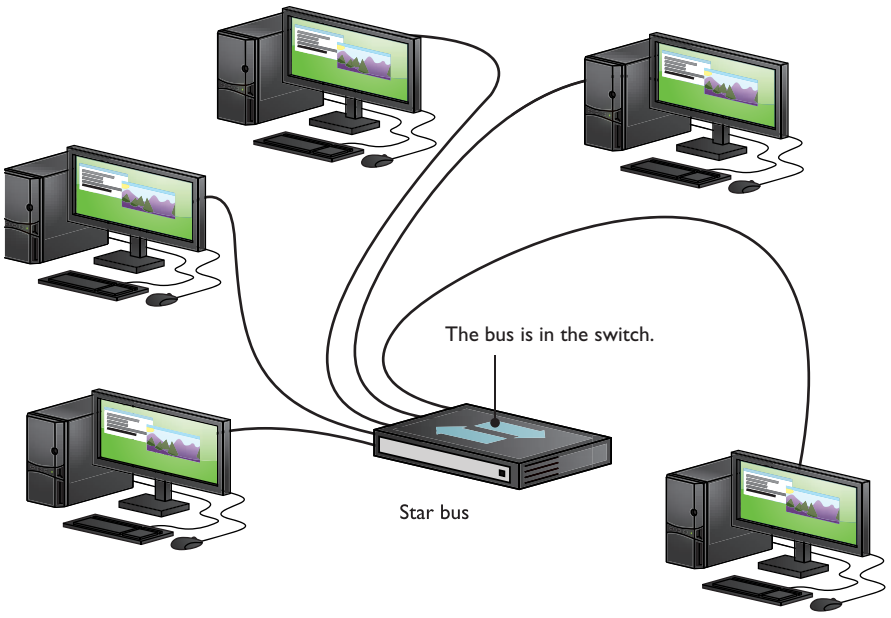

• **Figure 5.9** Star bus

The central box with the bus—a **switch**—provides a common point of connection for network devices. Switches can have a wide variety of ports. Most consumer-level switches have 4 or 8 ports, but business-level switches can have 32 or more ports.

Early Ethernet networks used a **hub**. A switch is a far superior and far more common version of a hub. Figure 5.10 shows a typical consumer-level switch.

A simple example demonstrates the difference between hubs and switches. Let's say you have a network of 32 PCs, all using 100-Mbps NICs attached to a 100-Mbps hub or switch. We would say the network's **bandwidth** is 100 Mbps. If you put the 32 PCs on a 32-port 100-Mbps hub, you have 32 PCs sharing the 100 Mbps of bandwidth. A switch addresses this problem by making each port its own separate network. Each PC gets to use the full bandwidth. The bottom line? Once switches became affordable, hubs went away.

• **Figure 5.10** A switch

The connection between a computer and a switch is called a **segment**. With most cable types, Ethernet segments are limited to 100 meters or less.

Cheap and centralized, a star bus network does not go down if a single cable breaks. True, the network would go down if the switch failed, but that is rare. Even if a switch fails, replacing a switch in a closet is much easier than tracing a bus running through walls and ceilings and trying to find a break!

Unshielded Twisted Pair

Unshielded twisted pair (UTP) cabling is the specified cabling for 10/100/1000BaseT and is the predominant

cabling system used today. Many types of twisted pair cabling are available, and the type used depends on the needs of the network. Twisted pair cabling consists of AWG 22–26 gauge wire twisted together into color-coded pairs. Each wire is individually insulated and encased as a group in a common jacket.

CAT Levels UTP cables come in categories that define the maximum speed at which data can be transferred (also called *bandwidth*). The major categories (CATs) are outlined in Table 5.1.

Table 5.1	CAT levels
CAT 1	Standard telephone line
CAT 3	Designed for 10-Mbps networks; a variant that used all four pairs of wires supported 100-Mbps speeds
CAT 5	Designed for 100-Mbps networks
CAT 5e	Enhanced to handle 1000-Mbps networks
CAT 6	Supports 1000-Mbps networks at 100-meter segments; 10-Gbps networks up to 55-meter segments
CAT 6a	Supports 10-Gbps networks at 100-meter segments

Although these days you'll only find CAT 3 installed for telephones and in very old network installations, CompTIA traditionally enjoys tripping up techs who don't know it could handle 100-Mbps networks.

The CAT level should be clearly marked on the cable, as Figure 5.11 shows.

The *Telecommunication Industry Association/Electronics Industries Alliance* (*TIA/EIA*) establishes the UTP categories, which fall under the TIA/EIA 568 specification. Currently, most installers use CAT 5e or CAT 6 cable.

• **Figure 5.11** Cable markings for CAT level

Shielded Twisted Pair

Shielded twisted pair (STP), as its name implies, consists of twisted pairs of wires surrounded by shielding to protect them from EMI, or electromagnetic interference. STP is pretty rare, primarily because there's so little need for STP's shielding; it only really matters in locations with excessive electronic noise, such as a shop floor area with lots of lights, electric motors, or other machinery that could cause problems for other cables.

Implementing 10/100/1000BaseT

The 10BaseT and 100BaseT standards require two pairs of wires: a pair for sending and a pair for receiving. 10BaseT ran on an ancient CAT version called CAT 3, but typically used at least CAT 5 cable. 100BaseT requires at least CAT 5 to run. 1000BaseT needs all four pairs of wires in a CAT 5e or CAT 6 cable. These cables use a connector called an **RJ-45** connector. The *RJ* (*registered jack*) designation was invented by Ma Bell (the phone company, for you youngsters) years ago and is still used today.

There are CAT levels for connectors as well as cables. Don't even try to use a CAT 5e RJ-45 connector with a CAT 6 cable.

Currently only two types of RJ connectors are used for networking: RJ-11 and RJ-45 (see Figure 5.12). **RJ-11** connects your telephone to the telephone jack in the wall of your house. It supports up to two pairs of wires, though most phone lines use only one pair. The other pair is used to support a second phone line. RJ-11 connectors are primarily used for telephone-based

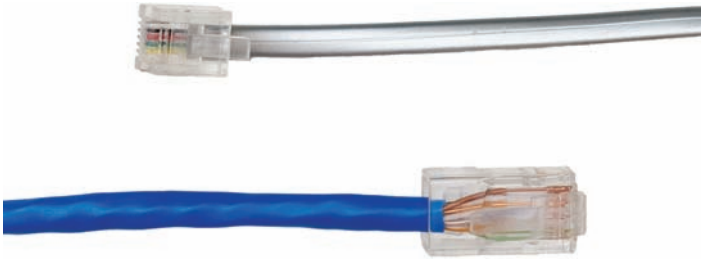

• **Figure 5.12** RJ-11 and RJ-45

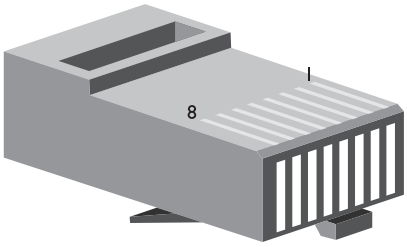

• **Figure 5.13** RJ-45 pin numbers

Internet connections (see Chapter 24) and are not used in any common LAN installation, although a few weird (and out of business) "network in a box" companies used them. RJ-45 is the standard for UTP connectors. RJ-45 has connections for up to four pairs and is visibly much wider than RJ-11. Figure 5.13 shows the position of the #1 and #8 pins on an RJ-45 jack.

The TIA/EIA has two standards for connecting the RJ-45 connector to the UTP cable: the TIA/EIA 568A (**T568A**) and the TIA/EIA 568B (**T568B**). Both are acceptable. You do not have to follow any standard as long as you use the same pairings on each end of the cable; however, you will make your life simpler if you choose a standard. Make sure that all of your cabling uses the same standard and you will save a great deal of work in the end. Most importantly, *keep records*!

Like all wires, the wires in UTP are numbered. A number does not appear on each wire, but rather each wire has a standardized color. Table 5.2 shows the official TIA/EIA Standard Color Chart for UTP.

Understanding Ethernet is critical in your understanding of how networks function. But when we talk about networking, there are two interconnected but very different worlds: the small local area networks where nearby users connect their computers via switches, and the Internet. Let's make sure we understand the differences by exploring the world of LANs and WANs in the next section.

Tech Tip

Plenum Versus PVC Cabling

Most workplace installations of network cable go up above the ceiling and then drop down through the walls to present a nice port in the wall. The space in the ceiling, under the floors, and in the walls through which cable runs is called the plenum *space. The potential problem with this cabling running through the plenum space is that the protective sheathing for networking cables, called the jacket, is made from plastic, and if you get any plastic hot enough, it creates smoke and noxious fumes.*

Standard network cables usually use PVC (polyvinyl chloride) *for the jacket, but PVC produces noxious fumes when burned. Fumes from cables burning in the plenum space can quickly spread throughout the building, so you want to use a more fire-retardant cable in the plenum space.*

Plenum-grade cable is simply network cabling with a fire-retardant jacket and is required for cables that go in the plenum space. Plenum-grade cable costs about three to five times more than PVC, but you should use it whenever you install cable in a plenum space.

Network Protocols, LANs, and WANs

The whole idea of networking is often confusing to new techs simply because there are two seemingly different ways to look at networking: small networks that share resources like documents, music, and printers, and big networks like the Internet where you share Web pages, e-mail, and just about anything you want. Let's take a moment to clarify how these two very different networks are very similar because today both use the same language or protocol.

Table 5.2 UTP Cabling Color Chart

Pin	T568A	T568B	Pin	T568A	T568B
1	White/Green	White/Orange	5	White/Blue	White/Blue
2	Green	Orange	6	Orange	Green
3	White/Orange	White/Green	7	White/Brown	White/Brown
4	Blue	Blue	8	Brown	Brown

A Short History of the War of the Network Protocols and Why TCP/IP Won

Ethernet does a fine job of moving data from one machine to another, but Ethernet alone isn't enough to make a complete network; many other functions need to be handled. For example, an Ethernet frame holds a maximum of 1500 bytes. What if the data being moved is larger than 1500 bytes? Something has to chop up the data into chunks on one end and something else needs to reassemble those chunks on the other end so the file can be properly reassembled.

Another issue arises if one of the machines on the network has its network card replaced. Up to this point, the only way to distinguish one machine from another was by the MAC address on the network card. To solve this, each machine must have a name, an identifier for the network, which is "above" the MAC address. Something needs to keep track of the MAC addresses on the network and the names of the machines so that frames and names can be correlated. If you replace a PC's network card, the network will, after some special queries, update the list to associate the name of the PC with the new network card's MAC address.

Network protocol software takes the incoming data received by the network card, keeps it organized, sends it to the application that needs it, and then takes outgoing data from the application and hands it to the NIC to be sent out over the network. All networks use some protocol. Over the years there have been many network protocols, most combining multiple simple protocols into groups, called *protocol stacks*. This lead to some crazily named network protocols, such as NetBIOS/NetBEUI and TCP/IP.

NetBIOS/NetBEUI

During the 1980s, IBM developed the **NetBIOS Extended User Interface (NetBEUI)**, the default protocol for early versions of Windows. NetBEUI offered small size, easy configuration, and a relatively high speed. The underlying protocol stack was called *NetBIOS/NetBEUI*. The NetBIOS protocol handled naming conventions, while NetBEUI chopped up data for delivery via frames.

NetBIOS names were very simple. You could call your computer TIMMY or MIKESPC. NetBIOS didn't allow names to include anything but letters (uppercase only), numbers, and a very few special characters. NetBIOS/NetBEUI was great for little networks, but it relied on individual computers to send out frames addressed to the MAC address FF-FF-FF-FF-FF-FF—which meant *everybody*. The official term for this is a **broadcast**. Broadcasts eat up bandwidth, but they're great for a node that's trying to get a MAC address for another node, as shown in Figure 5.14.

Broadcasts are useful, but the larger the network is, the more bandwidth broadcasts eat up. In general, NetBIOS could handle about 300 computers on a single network before the broadcast became too much. By the mid-1980s, it was clear NetBIOS wasn't going to work for really large networks, so a new network protocol was in the works. Plus, scaling required new hardware as well.

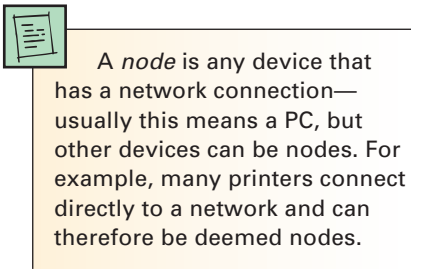

A *node* is any device that has a network connection—usually this means a PC, but other devices can be nodes. For example, many printers connect directly to a network and can therefore be deemed nodes.

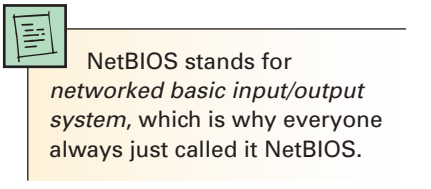

NetBIOS stands for *networked basic input/output system*, which is why everyone always just called it NetBIOS.

• **Figure 5.14** A broadcast in action

LANs, Routing, and WANs

For the CompTIA A+ exams, remember that a LAN is a group of networked computers that are close to each other. Also, remember that a LAN is almost always a broadcast domain.

A **local area network (LAN)** is a group of computers that are located physically close to each other—no more than a few hundred meters apart at most. A LAN might be in a single room, on a single floor, or in a single building. But I'm going to add that a LAN is almost always a group of computers that are able to "hear" each other when one of them sends a broadcast. A group of computers connected by one or more switches is a broadcast domain (see Figure 5.15).

A **wide area network (WAN)** is a widespread group of computers connected using long-distance technologies. You connect LANs into a WAN with a magical box called a **router** (see Figure 5.16). The best example of a WAN is the Internet.

You can connect multiple smaller networks into a bigger network, turning a group of LANs into one big WAN, but this raises a couple of issues

• **Figure 5.15** Two broadcast domains—two separate LANs

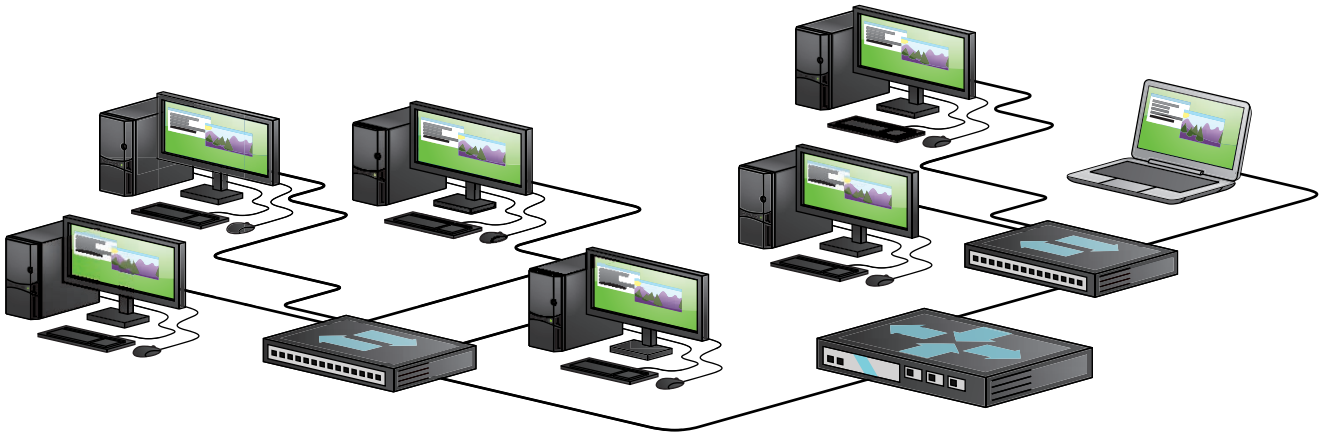

• **Figure 5.16** Two broadcast domains connected by a router—a WAN

with network traffic. A computer needs some form of powerful, flexible addressing to address a frame so that it goes to a computer within its own LAN or to a computer on another LAN on the same WAN. Broadcasting is also unacceptable, at least between LANs. If every computer saw every frame, the network traffic would quickly spin out of control! Plus, the addressing scheme needs to work so that routers can sort the frames and send them along to the proper LAN. This process, called *routing*, requires routers and a routing-capable protocol to function correctly.

Routers destroy any incoming broadcast frames, by design. No broadcast frame can ever go through a router. This makes broadcasting still quite common within a single broadcast domain, but never anywhere else.

NetBIOS/NetBEUI was great for a single LAN, but it lacked the extra addressing capabilities needed for a WAN. A new protocol was needed, one that could handle routing.

TCP/IP

Transmission Control Protocol/Internet Protocol (TCP/IP) was originally developed for the Internet's progenitor, the *Advanced Research Projects Agency Network* (*ARPANET*) of the U.S. Department of Defense. In 1983, TCP/IP became the built-in protocol for the popular BSD (Berkeley Software Distribution) UNIX, and other flavors of UNIX quickly adopted it as well. The biggest network of all, the Internet, uses TCP/IP as its protocol. All versions of Windows (in fact, all operating systems today) use TCP/IP as the default protocol.

The reason TCP/IP has a slash in the middle is to reflect that TCP/IP isn't a single network protocol. It's a number of protocols that all work together. TCP handles getting the data between computers, while IP handles the addressing scheme that gives us something more powerful and flexible than MAC addresses. Let's cover IP addresses in this chapter and save some of the other protocols for Chapter 22.

Chapter 22 covers TCP/IP in much greater detail.

IP Addresses and Subnet Masks In a TCP/IP network, the systems don't have names but rather have IP addresses. The **IP address** is the unique identification number for your system on the network. Part of the address identifies the network, and part identifies the local computer (host) address on the network. IP addresses consist of four sets of eight binary numbers (octets), with each set separated by a period. The numbers range from 0 to 255. This is called *dotted-decimal notation*. Every computer on a network running TCP/IP gets an IP address like so:

202.34.16.11

Every computer in the same broadcast domain as this computer will have some numbers in common. If the network is small, then all the computers will share the first three octets. In this case, a computer with the IP address of 202.34.16.123 is in the same broadcast domain as 202.34.16.11. The part of the IP address that is common for all the computers in the same broadcast domain is called the **network ID**. The network ID for this example is 202.34.16.

Are these example computers part of the same broadcast domain as a device with an IP address of 202.34.15.33? No. In this example, only

computers that are part of 202.34.16 are part of the same broadcast domain as 202.34.16.11. But how does every computer know which part of its IP address identifies the network ID? The **subnet mask** tells the computer which part of its IP address is the network ID. A typical subnet mask looks like this:

255.255.255.0

The network ID is determined by the number of 255 octets in the subnet mask. If there is a 255, that part of your address is the network ID. If your IP address were 190.24.16.11 and your subnet mask were 255.255.0.0, your network ID would be 190.24. If the subnet mask were 255.255.255.0, the same IP address would have a network ID of 190.24.16.

So far, I've only discussed Internet Protocol version 4 (IPv4) addresses. A newer version, known as Internet Protocol version 6 (IPv6), doesn't follow the same conventions. You'll learn a lot more about IPv6 in Chapter 22.

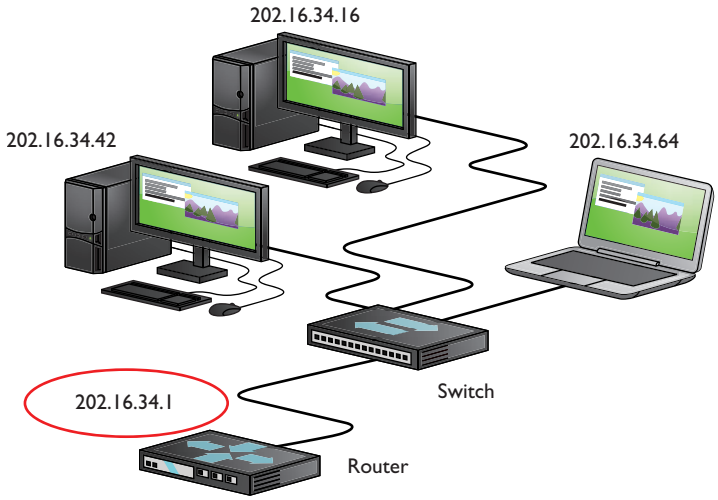

• **Figure 5.17** Default gateway

Default Gateway Sometimes you'll want to talk to computers that are outside your network. In that case, you'll need to connect to a router. I can now give another description of a router. A router has at least two IP addresses: one that connects to your LAN's switch and one that connects to the "next network." That next network could be your Internet service provider (ISP) or another router at your company—who knows (and more importantly, who cares, as long as it gets there)? The port on your router that connects to your LAN is given an IP address that's part of your network ID. In most cases, this is the first address shown in Figure 5.17.

The IP address of the "LAN" side of your router (the port connected to your LAN) is the address your computer uses to send data to anything outside your network ID. This is called the **default gateway**.

Domain Name Service (DNS) Knowing that users could not remember lots of IP addresses, early Internet pioneers came up with a way to correlate those numbers with more human-friendly designations. Special computers, called **domain name service (DNS)** servers, keep databases of IP addresses and their corresponding names. For example, let's say a machine with the IP address 209.34.45.163 hosts a Web site and we want it to be known as www.totalsem.com. When we set up the Web site, we pay money for a DNS server to register the DNS name www.totalsem.com to the IP address 209.34.45.163. So instead of typing "http://209.34.45.163" to access the Web page, you can type "www.totalsem.com." Your system will then query the DNS server to get www.totalsem.com's IP address and use that to find the right machine. Unless you want to type in IP addresses all the time, you'll need to use DNS servers (see Figure 5.18).

The Internet has regulated domain names. If you want a domain name that others can access on the Internet, you must register your domain name and pay a small yearly fee. Originally, DNS names all ended with

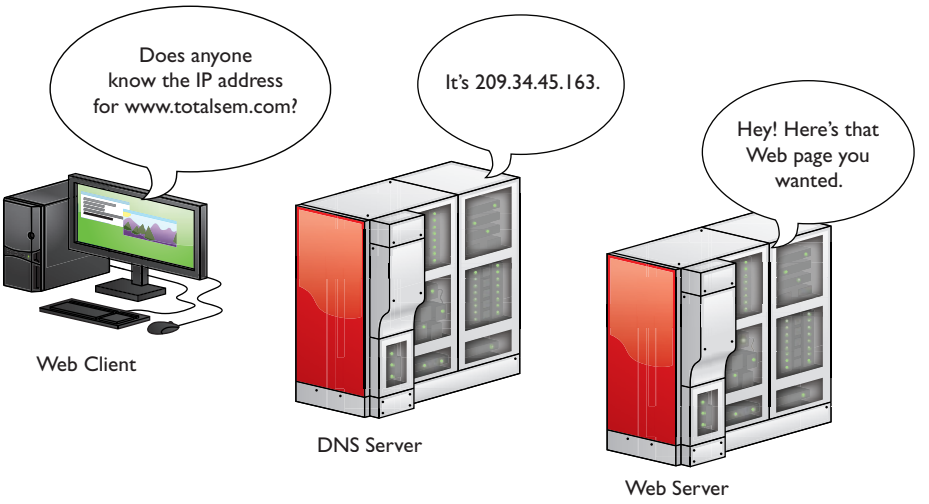

• **Figure 5.18** Domain name service

one of the following seven domain name qualifiers, called *top-level domains* (*TLDs*):

.com	General business	**.mil**	Military organizations
.edu	Educational organizations	**.net**	Internet organizations
.gov	Government organizations	**.org**	Nonprofit organizations
.int	International		

As more and more countries joined the Internet, a new level of domains was added to the original seven to indicate a DNS name from a particular country, such as .uk for the United Kingdom. It's common to see DNS names such as www.bbc.co.uk or www.louvre.fr. The *Internet Corporation for Assigned Names and Numbers* (*ICANN*) has added several new domains over the years, including .name, .biz, .info, .tv, and others. Given the explosive growth of the Internet, these are unlikely to be the last ones! For the latest developments, check ICANN's Web site at www.icann.org.

Entering the IP Information When you configure a computer to connect to a network, you must enter the IP address, the subnet mask, the default gateway, and at least one DNS server. Let's review:

- **IP address** A computer's unique address on the network
- **Subnet mask** Identifies your network ID
- **Default gateway** IP address or the LAN side of your router
- **DNS server** Tracks easy-to-remember DNS names for IP addresses

CompTIA uses the term *client-side* DNS to describe DNS configuration done on a local machine.

Configuring the IP address differs between each version of Windows. Figure 5.19 shows the IP settings on a Windows 7 system.

As you look at Figure 5.19, note the radio button for *Obtain an IP address automatically*. This is a common setting for which you don't need to enter any information. You can use this setting if your network uses Dynamic Host Configuration Protocol (DHCP). If you have DHCP (most networks

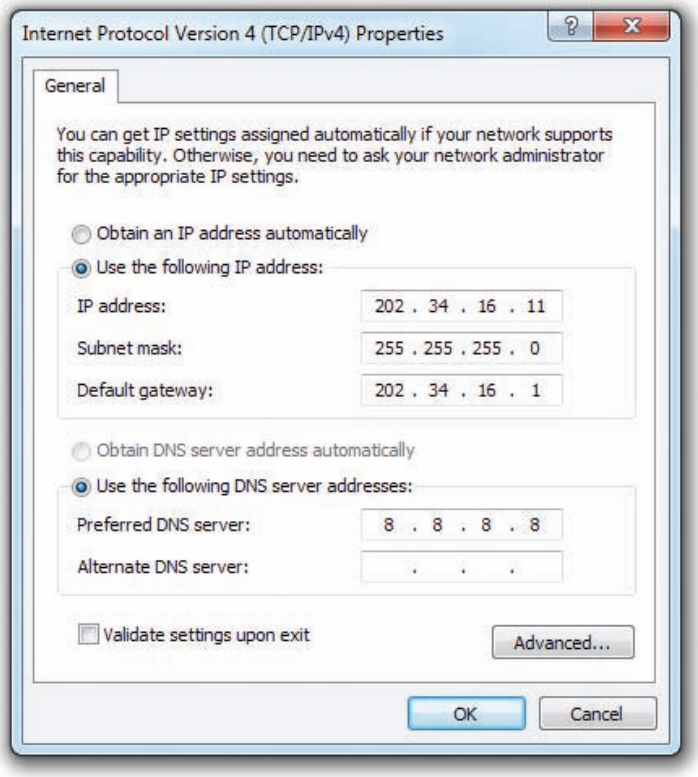

Figure 5.19 IP settings on a Windows 7 system

do) and your computer is configured to obtain an IP address automatically, your computer boots up and will broadcast a DHCP request. The DHCP server provides your computer with all the IP information it needs to get on the network (see Figure 5.20).

Refer to Chapter 22 for more details on IP addresses, subnet masks, default gateways, and DNS servers. For now, let's continue to tour the visible network.

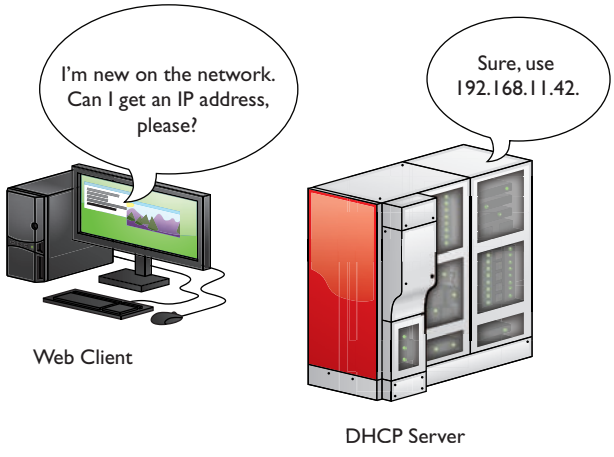

Figure 5.20 A DHCP server handing out an IP address

802

Network Organization

Once a network is created using appropriate network technology like Ethernet, users need to be able to share resources in some organized fashion. Resources such as folders and printers need a way to determine who can and cannot use them and how they can be used. Microsoft designed Windows networks to work in one of three categories: workgroups, domains, or homegroups. (These are the Microsoft terms, but the concepts have been adopted by the entire computer industry and apply to Mac OS X and other operating systems.) These three organizations differ in control, number of machines needed, compatibility, and security.

Let's start with the oldest network organization: workgroups.

Workgroups

Workgroups are the most basic and simplistic of the three network organizations. They are also the default for almost every fresh installation of Windows. Workgroups have been around since the ancient Windows for Workgroups came out back in the early 1990s.

By default, all computers on the network are assigned to a workgroup called WORKGROUP. You can see your workgroup name by opening the System applet (press WINDOWS KEY-PAUSE or go to Start | Control Panel | System applet), as shown in Figure 5.21.

There's nothing special about the name WORKGROUP, except that every computer on the network needs the same workgroup name to be able to share resources. If you want to change your workgroup name, you need to use the System applet. Click the *Change settings* button to open the System Properties dialog box. Then click the Change button to change your workgroup name (see Figure 5.22).

Most workgroup-based Windows networks keep the default name of WORKGROUP.

Workgroups lack centralized control over the network; all systems connected to the network are equals. This works well for smaller networks because there are fewer users, connections, and security concerns to think about. But what do you do when your network encompasses dozens or hundreds of users and systems? How can you control all of that?

User Names and Passwords

As you'll recall from Chapter 4, when you log on to a Windows computer, you need to enter a user name and password. Windows 7 makes this easy by giving you a pretty logon interface, as shown in Figure 5.23.

The user names and their passwords are stored in an encrypted format on your computer. User names have a number of jobs on your computer, but at this point the job most interesting to us is to give a user access to the computer. User names work well when you access your own computer, but these same user names and passwords are used to access shared resources on other computers in the network—and that's where we run into trouble.

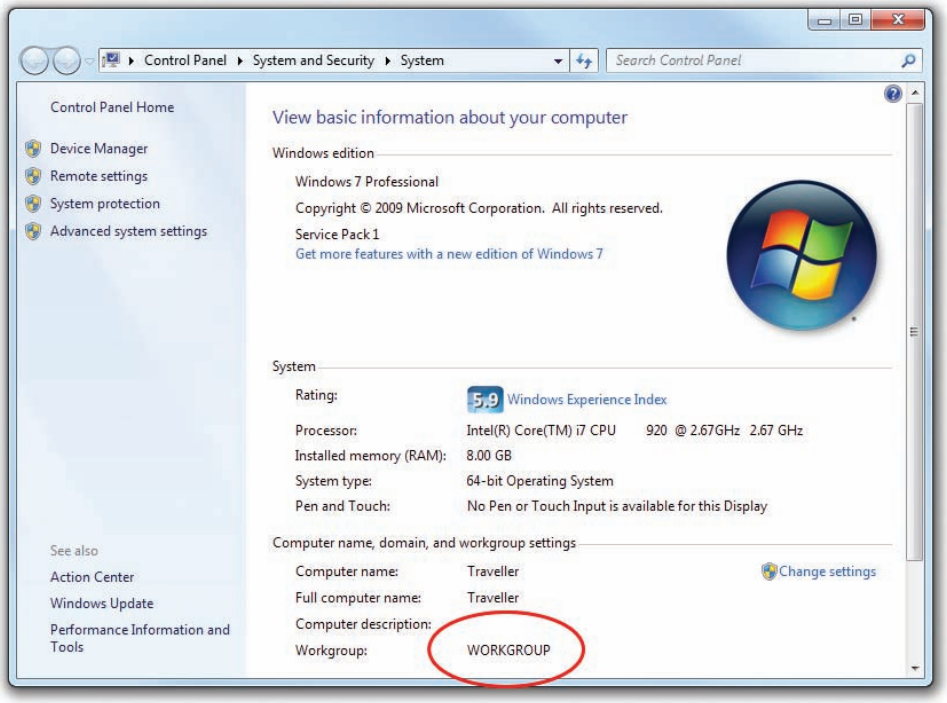

• **Figure 5.21** Default workgroup

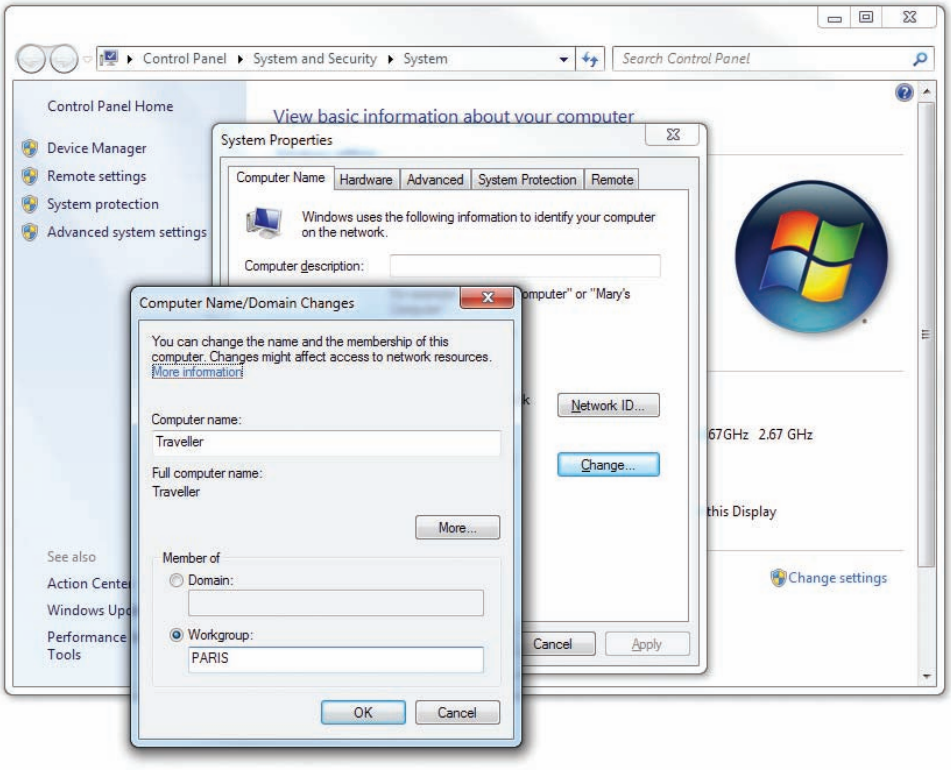

• **Figure 5.22** Changing the workgroup in advanced settings

• **Figure 5.23** Windows 7 logon screen

To appreciate this problem, let's watch a typical folder share take place on a network of Windows 7 systems.

Sharing a Folder

All Windows computers can share folders and printers out of the box. Sharing a folder in Windows 7 is easy—just right-click on the folder and select Share with | Specific people to get to the File Sharing dialog box (see Figure 5-24). On most Windows 7 systems, you'll see options called Homegroup in the context menu—ignore these for now as all will be explained in the next section.

By default, you'll see every user account that's currently on this system. You may give an account Read or Read/Write permission, while the person who created the folder is assigned as Owner. The following list describes these permissions:

- **Read** You can see what's in the folder. You may open files in the folder, but you can't save anything back into the folder.
- **Read/Write** Same as Read but you can save files into the folder.
- **Owner** Same as Read/Write plus you can set the permissions for other users on the folder.

So all this sharing seems to work quite nicely, except for one big issue: When you log on to your computer, you are accessing a user name and database on that computer. The accounts you are giving access to are stored on your computer, so how do you give someone from another computer access to that shared folder? You have to give that other person a valid user

All versions of Windows come with a far more powerful and much more complex form of permissions based on the NTFS file system. We'll save the big discussion of this more advanced form of permissions for Chapter 16.

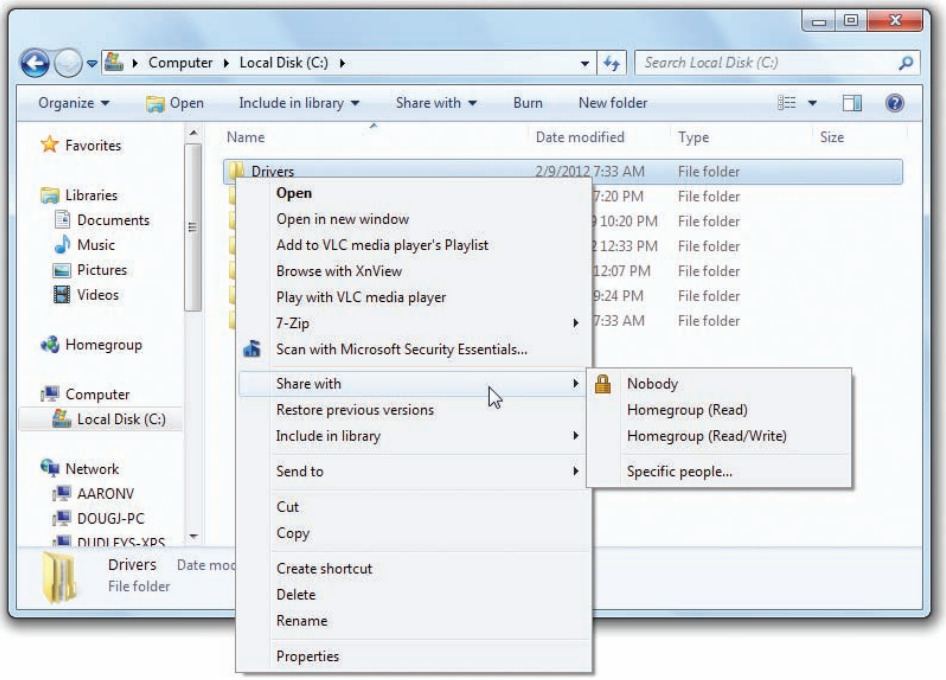

• **Figure 5.24** Folder Sharing dialog box

name and password. We use the nomenclature <computer name>\<user name> to track our logons. If you log on to Computer A as Mike, we say you are logged on to ComputerA\Mike. This nomenclature comes in very handy when networked computers become part of the process.

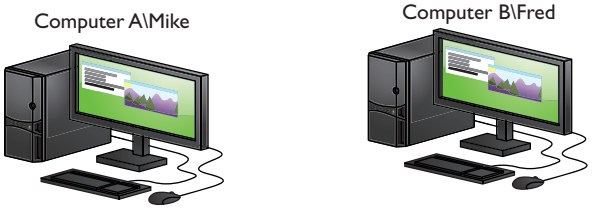

• **Figure 5.25** Computers A and B

Figure 5.25 shows an account called Mike on Computer A. Assume there is a shared folder called Timmy on Computer A and Mike has read/write permission.

A person fires up Computer B, logging in as Fred. He opens his Network menu option and sees Computer A, but when he clicks on it he sees a network password prompt (see Figure 5.26).

The reason is that the person is logged on as ComputerB\Fred and he needs to be logged on as ComputerA\Mike to access this folder successfully. So the user needs to know the password for ComputerA\Mike. This isn't a very pretty way to protect user names and passwords. So what can you do? You have three choices:

1. You can make people log on to shares as just shown.
2. You can create the same accounts (same user name and same password) on all the computers and give sharing permissions to all the users for all the shares.
3. You can use one account on all computers. Everyone logs on with the same account, and then all shares are by default assigned to the same account.

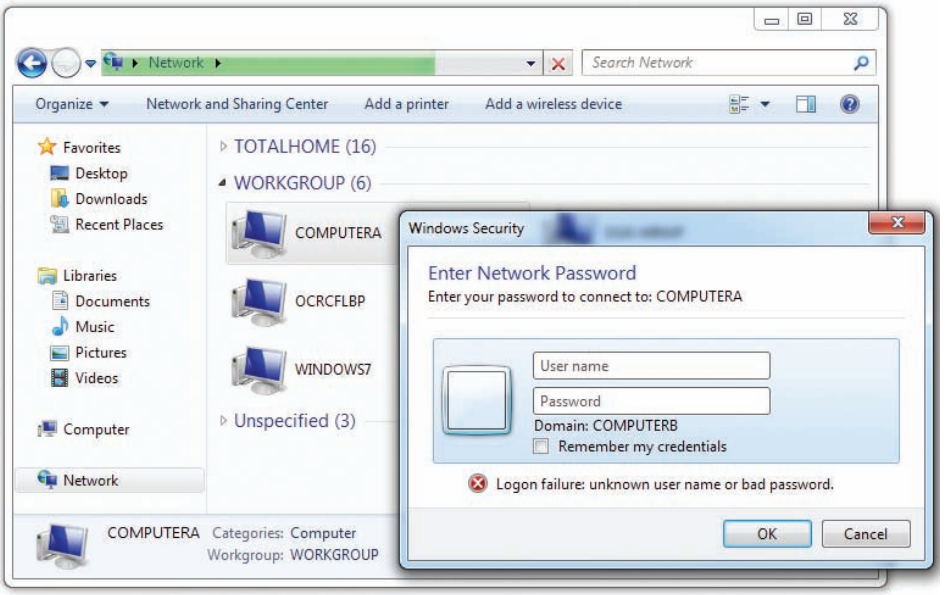

• **Figure 5.26** Prompt for entering user name and password

Domains

Larger networks that need more control use **domains**. Opposite the decentralized nature of workgroups, domains require a specific server to control access to the network's resources. This means tracking each user, each resource, and what each user can do to each resource.

To use a domain on a network of Windows machines, you must have a computer running a version of Windows Server (see Figure 5.27). Windows Server is a completely different, much more powerful, and much more expensive version of Windows. Current editions of this specialized OS include Windows Server 2008 and Windows Server 2008 R2.

An administrator creates a domain on the Windows Server system, which makes that system the *domain controller* (*DC*). The administrator also creates new user accounts on the domain controller. These accounts are called *domain accounts*. Once a network is set up as a domain, each PC on the network needs to join the domain (which kicks you off the workgroup). When you log on to a computer that's a member of a domain, Windows will prompt you for a user name instead of showing you icons for all the users on the network (see Figure 5.28).

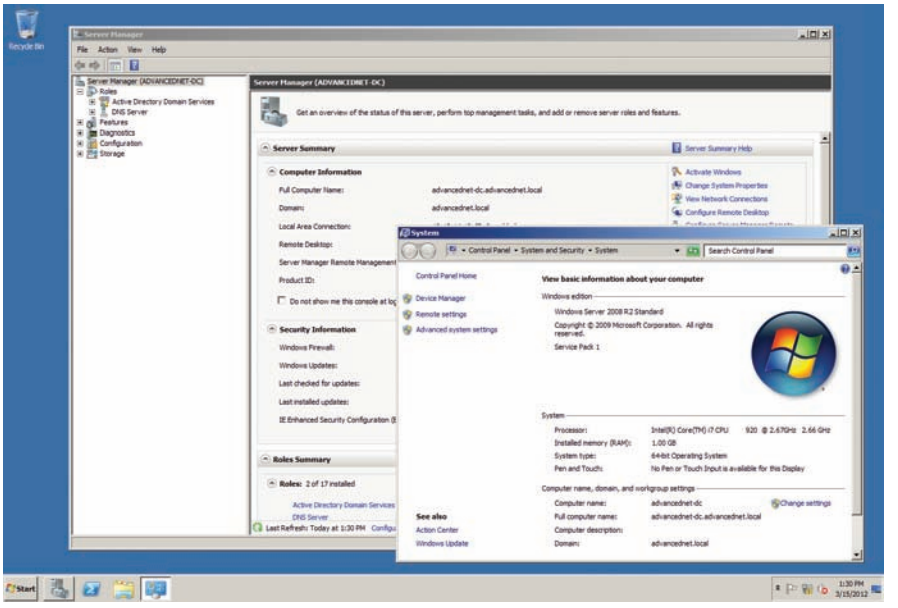

• **Figure 5.27** Windows Server

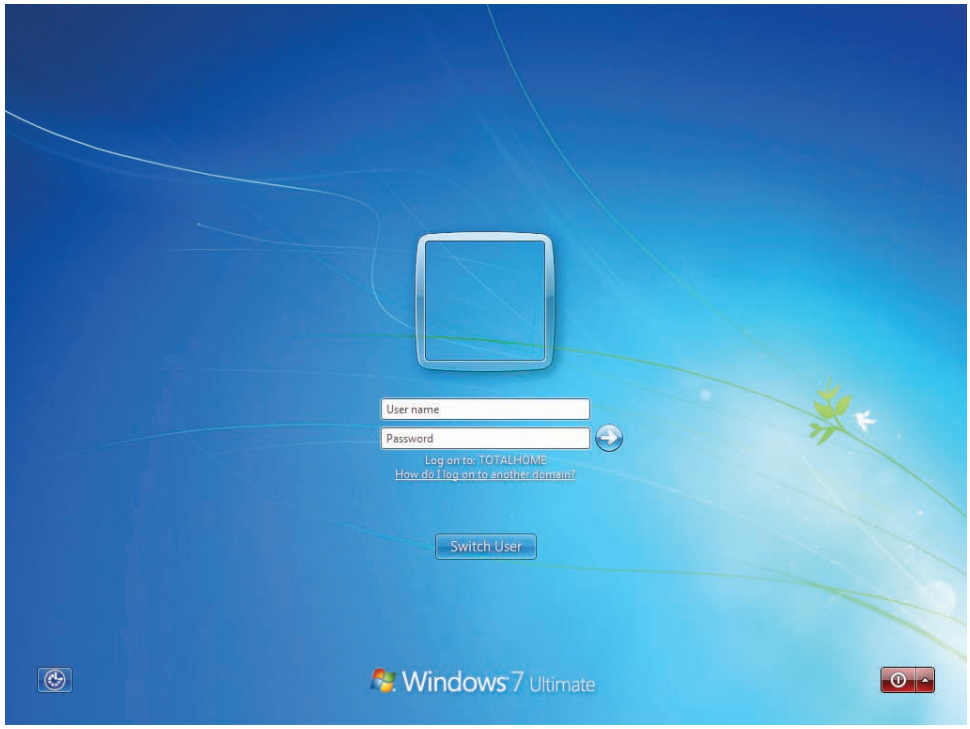

• **Figure 5.28** Domain logon screen

There is much more to a Windows domain than single sign-on. For the CompTIA A+ certification, however, that's the big selling point. If you want to delve deeper into Windows domains, consider pursuing the CompTIA Network+ certification or one of the Microsoft certifications.

When using a domain, you don't log on to your computer. Instead, you log on directly to the domain. All user accounts are stored on the domain controller, as shown in Figure 5.29. A lot of domains have names that look like Web addresses, like totalhome.com, totalhome.local, or even just totalhome. Using the previous nomenclature, you can log on to a domain using <domain>\<domain user name>. If the domain totalhome.local has a user

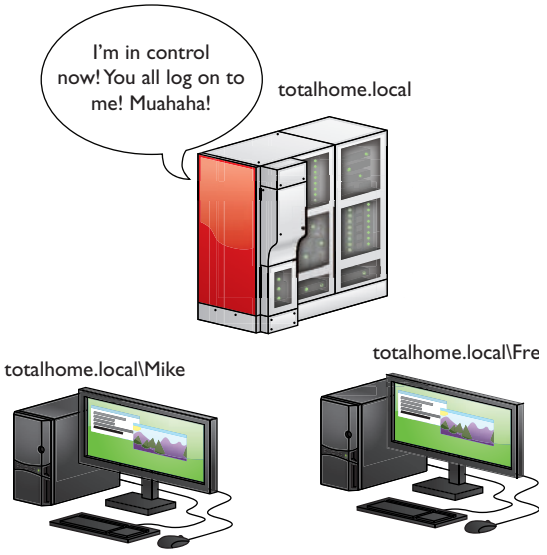

• **Figure 5.29** Domain network

account called Mike, for example, you would use totalhome.local\Mike to log on.

One of the best features of domains is that you can log on to any computer on the domain using the same domain account. You don't have to set up local accounts on each computer. We call this feature *single sign-on*, and for most users, this is the biggest benefit to using a Windows domain.

Homegroups

The problem with workgroups is that they provide almost no security and require lots of signing on to access resources. Domains provide single sign-on and lots of security, but require special servers and lots of administration. To address this, Microsoft introduced a feature in Windows 7 called **HomeGroup**.

HomeGroup uses the idea that people want to connect data, not folders. Most people want to share their music, not their My Music folder. So homegroups skip folders completely and share Windows 7 libraries.

A homegroup connects a group of computers using a common password—no special user names required. Each computer can be a member of only one homegroup at a time. Let's make a homegroup and see how this works.

To make a homegroup, open the HomeGroup Control Panel applet. Assuming you currently connect to a workgroup and haven't already created a homegroup, you'll see a dialog box like the one shown in Figure 5.30.

Click the *Create a homegroup* button to create a homegroup. You'll then see the Create a Homegroup dialog box shown in Figure 5.31.

Notice the five options: Pictures, Music, Videos, Documents, and Printers. Remember that those first four are the libraries you learned about in Chapter 4? The Documents checkbox is probably not checked, but go ahead and check it to share all five things. Click Next to see the homegroup's password (see Figure 5.32)

• **Figure 5.30** Default HomeGroup dialog box

Homegroups are not available in Windows XP or Windows Vista.

You may want to review what you learned about libraries in the previous chapter.

Microsoft refers to the technology as HomeGroup, but drops the capitalization to homegroup when talking about the groups themselves. Look for it to appear either way on the CompTIA A+ exams.

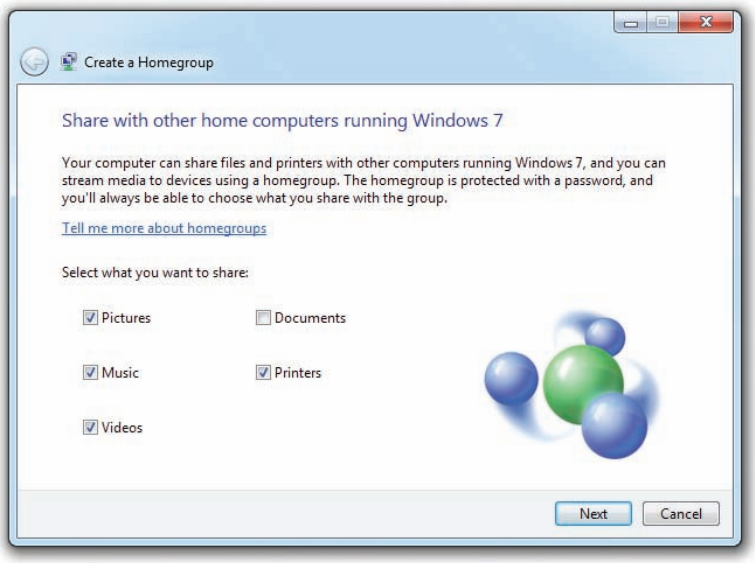

Figure 5.31 Create a Homegroup dialog box

Interestingly, all homegroup data is encrypted between systems.

Perhaps you've heard that you shouldn't write down passwords? Well, this password is so long that you might *need* to write it down. The dialog box even gives you a way to print it out! Click Next one more time to see the dialog box shown in Figure 5.33. This is the dialog box you will now see every time you click the HomeGroup applet in the Control Panel.

Let's look at this carefully. Notice where it says *Share libraries and printers* and, a bit lower, *How do I share additional libraries?* By default, homegroups share libraries, not individual folders. The Music, Pictures, Videos, and Documents libraries are shared by default. Although printers get their own checkbox, this setting remains the same as a normal printer share. It's

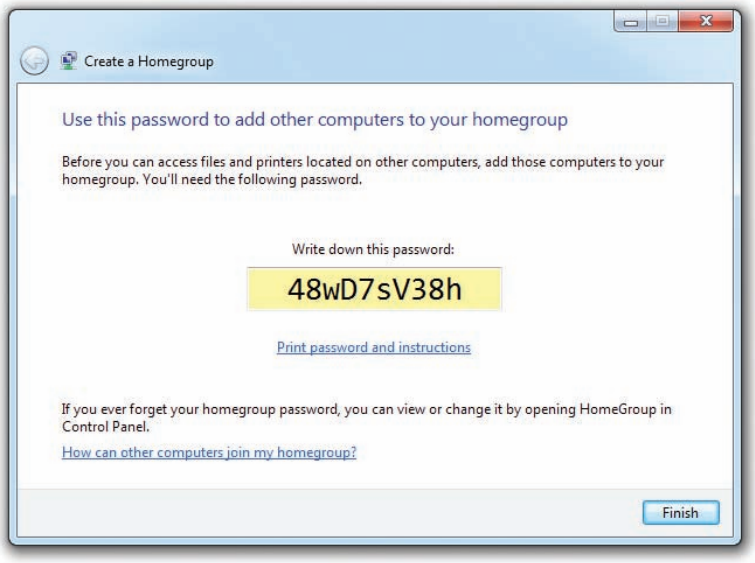

Figure 5.32 The homegroup's password

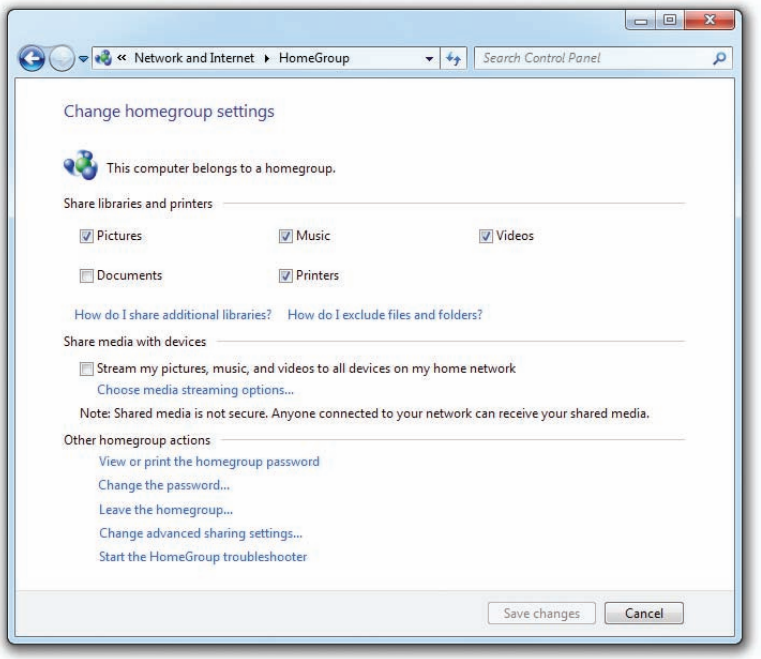

• **Figure 5.33** Homegroup configured

just a handy place to add printer sharing, as even the most basic users like to share printers.

Once you've created a homegroup, go to another computer on the network and open the HomeGroup Control Panel applet. Assuming all the factors stated earlier, you will see a dialog box like Figure 5.34.

Click the *Join now* button, enter the password, choose which libraries you want to share with everyone else, and the new computer is in the homegroup!

Remember that homegroups share libraries, not folders, by default.

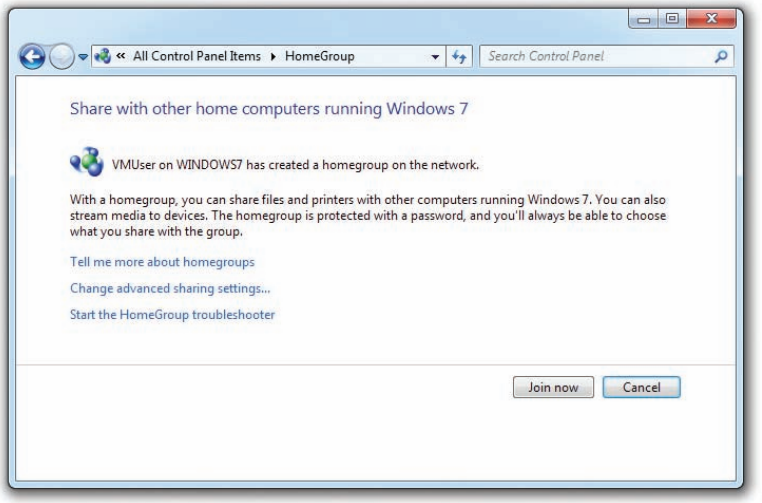

• **Figure 5.34** HomeGroup showing an existing homegroup

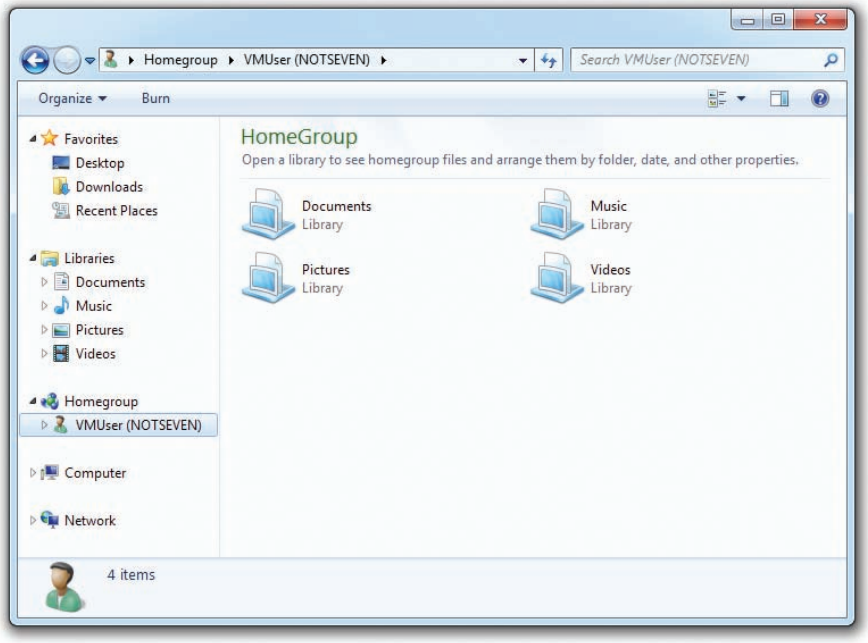

• **Figure 5.35** Using homegroups

Access the files shared through a homegroup by opening Windows Explorer, as shown in Figure 5.35. To see what others are sharing, select the corresponding computer name. You can then open those libraries to see the shared folders.

Sharing more libraries is easy, and, if you'd like, you can even share individual folders. Just right-click on the library or folder and select Share with, as shown in Figure 5.36.

Notice you have four options: Nobody, Homegroup (Read), Homegroup (Read/Write), and Specific people. The Nobody option means the item is not shared.

By sharing libraries with homegroups, Microsoft hides folders for most users, helping users share their stuff (documents, pictures, music, and videos) instead of folders. Homegroups fit a very specific world: smaller, non-domain home networks, but within that realm, they work wonderfully.

Once you create a homegroup, you can access it from Windows Explorer.

Windows Explorer also adds a *Share with* toolbar button that works exactly like the menu shown in Figure 5.36.

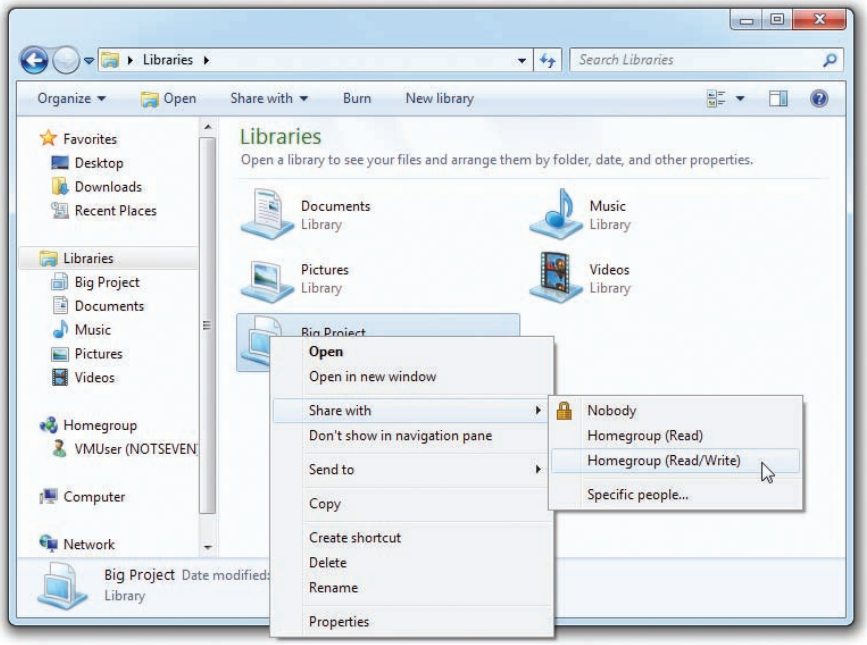

• **Figure 5.36** The Share with menu

Chapter 5 Review

Chapter Summary

After reading this chapter and completing the exercises, you should understand the following facts about networks and resource sharing.

Describe the basic functions of a network, including identifying common devices and connectors

- Topologies define how data is moved between computers in a network. The most important types of topology are bus, ring, mesh, and star. Almost all networks today use the hybrid star bus topology.
- Networks have one or more client machines with NICs for both defining the nodes and creating frames for sending data to other nodes. Some kind of network cabling or wireless signal connects the nodes together.
- Each NIC has a unique MAC address that identifies a node. Frames include the MAC address for both the sender and recipient, some kind of data payload, and a data check.
- Most networks today employ some form of Ethernet to define things like frame type and cabling. The vast majority of 10/100/1000BaseT networks run on unshielded twisted-pair cabling. Switches provide the central connection point in the Ethernet star bus topology. Most network cables connect with RJ-45 connectors.
- Varieties of UTP include CAT 1, CAT 3, CAT 5, CAT 5e, and CAT 6. Each can be plenum-grade or use a PVC jacket.

Discuss the differences between a LAN and a WAN and the importance of TCP/IP

- Network protocols enable you to name and address computers within a network, removing the need to rely solely on the MAC address to get data from one node to another.
- NetBIOS/NetBEUI protocols were used on early Windows networks for naming purposes. Because they relied on broadcasts, the move to scale up in size required a new protocol for the larger network.
- You can connect LANs together to form a WAN using routers to make the connections. Routers don't pass broadcast frames.
- All modern operating systems use TCP/IP as the default protocol. TCP is the part that handles getting the data between computers, while IP handles the addressing scheme that gives us something more powerful and flexible than MAC addresses. Other protocols in the protocol stack handle other duties.
- The IP address is the unique identification number for your system on the network. Part of the address identifies the network, and part identifies the local computer (host) address on the network. The subnet mask tells the computer which part of its IP address is the network ID.
- DNS servers enable users to type in names for remote nodes rather than having to remember and type the numerical IP addresses. Many systems are set up with DHCP enabled, where a DHCP server will provide all the TCP/IP settings needed by the systems at boot.

Perform basic resource sharing

- Microsoft designed Windows networks to work in one of three categories: workgroups, domains, or homegroups. These three organizations differ in control, number of machines needed, compatibility, and security.
- Every computer on a network needs the same workgroup name to be able to share resources. By default, all computers on the network are assigned to a workgroup called WORKGROUP. Workgroups require a valid user name and password to log on to each node in the workgroup.
- When using a domain, you don't log on to your computer. Instead, you log on directly to the domain controller, a node running Windows Server that stores all user accounts. This enables single sign-on for all resources available to your user account on the domain computers.
- Microsoft added HomeGroup to Windows 7 to simplify sharing within a small LAN. Once you add connected computers into a homegroup, sharing is simple, with no further logon required for accessing shared resources such as libraries.

Key Terms

10BaseT *(131)*
100BaseT *(131)*
1000BaseT *(131)*
bandwidth *(132)*
broadcast *(135)*
bus *(128)*
client *(125)*
cyclic redundancy check *(131)*
default gateway *(138)*
domain *(145)*
domain name service (DNS) *(138)*
Ethernet *(131)*
frame *(127)*
HomeGroup *(147)*
hub *(132)*
hybrid *(128)*
IP address *(137)*
local area network (LAN) *(136)*
media access control (MAC) address *(130)*
mesh *(128)*
NetBIOS Extended User Interface (NetBEUI) *(135)*
network ID *(137)*
network interface controller (NIC) *(127)*
network protocol *(135)*
network technology *(129)*
resource *(127)*
RJ-11 *(133)*
RJ-45 *(133)*
ring *(128)*
router *(136)*
segment *(132)*
server *(125)*
shielded twisted pair (STP) *(133)*
star *(128)*
star bus *(131)*
subnet mask *(138)*
switch *(132)*
T568A *(134)*
T568B *(134)*
topology *(128)*
Transmission Control Protocol/Internet Protocol (TCP/IP) *(137)*
unshielded twisted pair (UTP) *(132)*
Web browser *(125)*
Web server *(125)*
wide area network (WAN) *(136)*
workgroup *(141)*

Key Term Quiz

Use the Key Terms list to complete the sentences that follow. Not all terms will be used.

1. The _____ version of Ethernet runs at 100 Mbps.
2. Ethernet cables typically use a(n) _____ connector.
3. The most common Ethernet topology is _____.
4. The _____ tells the computer which part of its IP address is the network ID.
5. Most Ethernet networks have a central box called a(n) _____.
6. All NICs have a unique identifier called a(n) _____.
7. Routers enable multiple LANs to connect into a(n) _____.
8. Windows 7 enables sharing of libraries using _____.
9. A Windows _____ enables single sign-on.
10. The numbers 192.168.5.36 are an example of a(n) _____.

Multiple-Choice Quiz

1. How many bits are in a MAC address?
   - **A.** 24
   - **B.** 36
   - **C.** 48
   - **D.** 64
2. Which protocol enables the use of names such as www.totalsem.com rather than IP addresses?
   - **A.** DNS
   - **B.** MAC
   - **C.** IP
   - **D.** TCP
3. What is the minimum CAT level cable required for a 100BaseT network?
   - **A.** CAT 1
   - **B.** CAT 5
   - **C.** CAT 5e
   - **D.** CAT 6
4. Which OS enables you to implement HomeGroup?
   - **A.** Mac OS X
   - **B.** Windows XP
   - **C.** Windows 7
   - **D.** Windows Server
5. Which of the following is an example of a hybrid topology?
   - **A.** Bus
   - **B.** Ring
   - **C.** Star
   - **D.** Star bus
6. Which of the following is an advantage of a domain-based network over a workgroup-based network?
   - **A.** Ease of administration
   - **B.** Cheaper
   - **C.** Single sign-on
   - **D.** Faster to implement
7. A typical CAT 6 cable uses which connector?
   - **A.** RJ-11
   - **B.** RJ-45
   - **C.** Plenum
   - **D.** PVC
8. Why would you use STP over UTP cabling?
   - **A.** Cheaper
   - **B.** Easier to install
   - **C.** Better to avoid interference
   - **D.** They're interchangeable terms
9. What kind of frame gets received by all NICs in a LAN?
   - **A.** Ethernet
   - **B.** Broadcast
   - **C.** WAN
   - **D.** DNS
10. Internet Explorer, Mozilla Firefox, and Google Chrome are all examples of what?
    - **A.** Web servers
    - **B.** DNS
    - **C.** Web browsers
    - **D.** IP addresses
11. Which of the following does 255.255.255.0 most probably represent?
    - **A.** An IP address
    - **B.** A subnet mask
    - **C.** A default gateway
    - **D.** A DNS address
12. Which of the following is the default network organization for a new Windows installation?
    - **A.** LAN
    - **B.** Workgroup
    - **C.** Domain
    - **D.** HomeGroup

13. What is a remote computer providing resources called?
    - **A.** Web browser
    - **B.** DNS
    - **C.** Server
    - **D.** Router
14. What is the total amount of data that a network can handle called?
    - **A.** Protocol
    - **B.** Broadcast
    - **C.** Bandwidth
    - **D.** Client
15. What do you call the discrete "chunks" of data that networks transmit?
    - **A.** Frames
    - **B.** Stacks
    - **C.** LANs
    - **D.** Workgroups

Essay Quiz

1. Most computers come with onboard NICs. Check out your NIC and determine its fastest speed. Can your NIC support lower speeds as well? Show how you came to this determination.
2. Leah's home network has a large number of shared folders across four different systems and she's getting tired of continually logging on to these shares. She doesn't want to install a domain. What can she do to reduce the logon hassle?
3. Mike's office is composed of roughly ten PCs all running 100BaseT NICs and a 100BaseT switch on CAT 5 cable. He wants to "speed up the network" by going up to 1000BaseT. Explain to him the process of doing this upgrade and give him your opinion of whether this is a good idea.
4. Your company is running a Windows 7 network as a workgroup. There are 15 computers and 15 users on the network. Your boss wants to upgrade to a Windows domain. Do some research to determine the costs involved with upgrading your network. In your report to your boss, include the cost of a new server as well as a copy of Windows Server that will support 15 computers.
5. Even though NetBIOS/NetBEUI is long gone, Windows systems still use computer names. Do some research on the term SMB and explain how modern Windows systems still support network names.

Lab Projects

• Lab Project 5.1

To access the Internet, you must have a DNS server address. Typically this is automatically configured when your system starts, but you can add custom DNS servers to increase your Internet speed.

1. Go to www.grc.com and download the Domain Name Speed Benchmark utility. This utility compares your DNS server to alternatives that may be faster.
2. Run the utility. Did the utility find a faster DNS server? (Google has some very popular, very fast DNS servers.)
3. Figure out how to replace your DNS server for your system with one of these alternatives. Be sure to remember how to return to the original DNS settings in case this doesn't work!
4. Try browsing with your new settings. Do you notice a difference?

• Lab Project 5.2

A good tech should know their network. Take a tour of your network to answer the following questions:

1. What is the network ID for your network?
2. What is the workgroup or domain name?
3. What is the default DNS server?
4. What is the default gateway?

• Lab Project 5.3

It's very common for smaller Windows networks to all be members of a workgroup called WORKGROUP.

1. Put every computer into a different workgroup. How does this change the ability to see computers in the network folder? How does this change the way you share folders?
2. Put half the computers into one workgroup and the other half into another workgroup to see the changes as discussed in the previous question.
3. Assuming you have Windows 7 systems, put some of the computers into a homegroup and leave some out of the homegroup. Can you still share folders? What happens to non–Windows 7 computers in this situation?