

Part I

NETWORKING FUNDAMENTALS


COPYRIGHTED MATERIAL

Chapter 1

AN INTRODUCTION TO NETWORKS

Imagine 20 years ago working in an office with little or no computer equipment. It's hard to imagine now, isn't it? One could say that we take for granted a lot of what we have gained in technology the past few decades. Now, imagine having to send a memo to everyone in the company. Back then we used interoffice mail; today we use e-mail. This is one form of communication that only became available due to the introduction and growth of networks.

This chapter focuses on the basic concepts surrounding how a network works, including the way it sends information and what it uses to send information. This information is covered only to a minor degree by the A+ certification exam. However, if you have interest in becoming a service technician, this information will prove to be very useful, as you will in all likelihood find



Adapted from *A+ Complete Study Guide, Deluxe Edition*
by David Groth and Dan Newland
ISBN 0-7821-4052-1 976 pages \$59.99

yourself asked to troubleshoot both hardware and software problems on existing networks. Included in this chapter is information on:

- ▶ What is a network?
- ▶ Network types
- ▶ Media types
- ▶ Connectivity devices

**NOTE**

If you find that the material in this chapter interests you, you might consider studying for, and eventually taking, CompTIA's Network+ exam. It is a generic networking certification (similar to A+, only it is for network-related topics). You can study for it using Sybex's Network+ Study Guide materials available at www.sybex.com.

WHAT IS A NETWORK?

Stand-alone personal computers, first introduced in the late 1970s, gave users the ability to create documents, spreadsheets, and other types of data and save them for future use. For the small business user or home computer enthusiast, this was great. For larger companies, however, it was not enough. The larger the company, the greater the need to share information between offices, and sometimes over great distances. The stand-alone computer was not enough for the following reasons:

- ▶ Their small hard drive capacities were inefficient.
- ▶ To print, each computer required a printer attached locally.
- ▶ Sharing documents was cumbersome. People grew tired of having to save to a diskette, then taking that diskette to the recipient. (This procedure was called “sneakernet.”)
- ▶ There was no e-mail. Instead, there was interoffice mail, which was not reliable and frequently was not delivered in a timely manner.

To address these problems, *networks* were born. A network links two or more computers together to communicate and share resources. Their success was a revelation to the computer industry as well as businesses.

Now, departments could be linked internally to offer better performance and increase efficiency.

You have heard the term “networking” in the business context, where people come together to exchange names for future contact and to gain access to more resources. The same is true with a computer network. A computer network allows computers to link to each other’s resources. For example, in a network every computer does not need a printer connected locally to print. Instead, one computer has a printer connected to it and allows the other computers to access this resource. Because they allow users to share resources, networks offer an increase in performance as well as a decrease in the outlay for new hardware and software.

LANs vs. WANs

Local area networks (LANs) were introduced to connect computers in a single office. *Wide area networks (WANs)* came to expand the LANs to include networks outside of the local environment and also to distribute resources across distances. Today, LANs can be seen in many businesses, from small to large. WANs are becoming more widely accepted as businesses are becoming more mobile and as more of them are spanning greater and greater distances. It is important to have an understanding of LANs and WANs as a service professional, because when you’re repairing computers, you are likely to come in contact with problems that are associated with the computer being connected to a network.

Local Area Networks (LANs)

The 1970s brought us the minicomputer, which was a smaller version of the mainframe. Whereas the mainframe used *centralized processing* (all programs ran on the same computer), the minicomputer used *distributed processing* to access programs across other computers. As depicted in Figure 1.1, distributed processing allows a user at one computer to use a program on another computer as a “back end” to process and store the information. The user’s computer is the “front end,” performing the data entry. These allowed programs to be distributed across computers rather than centralized. This was also the first time computers used cable to connect rather than phone lines.




FIGURE 1.1: Distributed processing

By the 1980s, offices were beginning to buy PCs in large numbers. Also, portables were introduced, allowing computing to become mobile. Neither PCs nor portables, however, were efficient in sharing information. As timeliness and security became more important, diskettes were just not cutting it. Offices needed to find a way to implement a better means to share and access resources. This led to the introduction of the first type of PC LAN: ShareNet by Novell. LANs are simply the linking of computers to share resources within a closed environment. The first simple LANs were constructed a lot like Figure 1.2.




FIGURE 1.2: A simple LAN

After the introduction of ShareNet, more LANs sprouted. The earliest LANs could not cover a great distance. Most of them could only stretch across a single floor of the office and could support no more than 30 users. Further, they were still simple, and only a few software programs supported them. The first software programs that ran on a LAN were not capable of permitting more than one user at a time to use a program (this constraint was known as *file locking*). Nowadays, we can see multiple users accessing a program at one time, limited only by restrictions at the record level.

Wide Area Networks (WANs)

By the late 1980s, networks were expanding to cover ranges considered geographical in size and were supporting thousands of users. Wide area networks, first implemented with mainframes at massive government expense, started attracting PC users as networks went to this whole new level. Businesses with offices across the country communicated as if they were only desks apart. Soon the whole world would see a change in its way of doing business, across not only a few miles but across countries. Whereas LANs are limited to single buildings, WANs are able to span buildings, states, countries, and even continental boundaries. Figure 1.3 gives an example of a simple WAN.




FIGURE 1.3: A simple WAN

Networks of today and tomorrow are not limited anymore by the inability of LANs to cover distance and handle mobility. WANs play an important role in the future development of corporate networks worldwide. Although the primary focus of this chapter is LANs, we will feature a section on WAN connectivity. This section will briefly explain the current technologies and what you should expect to see in the future. If you are interested in more information on LANs or WANs, or if you plan on becoming a networking technician, check your local library resources or the Internet.

Primary Network Components

Putting together a network is not as simple as it was with the first PC network. You can no longer consider two computers cabled together a fully functional network. Today, networks consist of three primary components:

- ▶ Servers
- ▶ Clients or workstations
- ▶ Resources

No network would be complete without these three components working together.

Servers

Servers come in many shapes and sizes. They are a core component of the network, providing a link to the resources necessary to perform any task. The link it provides could be to a resource existing on the server itself or a resource on a client computer. The server is the “leader of the pack,” offering directions to the client computers regarding where to go to get what they need.

Servers offer networks the capability of centralizing the control of resources and can thus reduce administrative difficulties. They can be used to distribute processes for balancing the load on the computers and can thus increase speed and performance. They can also offer the departmentalizing of files for improved reliability. That way, if one server goes down, then not all of the files are lost.

Servers perform several tasks. For example, servers that provide files to the users on the network are called file servers. Likewise, servers that host printing services for users are called print servers. (There are other tasks as well, such as remote access services, administration, mail, etc.) Servers can be *multi-purpose* or *single-purpose*. If they are multi-purpose, they can be, for example, both a file server and a print server at the same time. If the server is a single-purpose server, it is a file server only or print server only.

Another distinction we use in categorizing servers is whether they are dedicated or nondedicated:

Dedicated Servers These are assigned to provide specific applications or services for the network, and nothing else.

Because a *dedicated server* is specializing in only a few tasks, it requires fewer resources from the computer that is hosting it than a nondedicated server might require. This savings in overhead may translate to a certain efficiency and can thus be considered as having a beneficial impact on network performance.

Nondedicated Servers These are assigned to provide one or more network services and local access. A *nondedicated server* is expected to be slightly more flexible in its day-to-day use than a dedicated server. Nondedicated servers can be used not only to direct network traffic and perform administrative actions, but often to serve as a front end for the administrator to work with other applications or services. The nondedicated server is not really what some would consider a true server, because it can act as a workstation as well as a server.

Many networks use both dedicated and nondedicated servers in order to incorporate the best of both worlds, offering improved network performance with the dedicated servers and flexibility with the nondedicated servers.

Workstations or Client Computers

Workstations are the computers that the users on a network do their work on, performing activities such as word processing, database design, graphic design, e-mail, and other office or personal tasks. Workstations are basically nothing more than an everyday computer, except for the fact that they are connected to a network that offers additional resources. Workstations can range from a diskless computer system to a desktop system. In network terms, workstations are also known as *client computers*. As clients, they are allowed to communicate with the servers in the network in order to use the network's resources.

It takes several items to make a workstation into a client. You must install a *network interface card (NIC)*, a special expansion card that allows the PC to talk on a network. You must connect it to a cabling system that connects to another computer (or several other computers). And you must install some special software, called *client software*, which allows the computer to talk to the servers. Once all this has been accomplished, the computer will be "on the network."

To the client, the server may be nothing more than just another drive letter. However, because it is in a network environment, the client is able

to use the server as a doorway to more storage or more applications, or through which it may communicate with other computers or other networks. To a user, being on a network changes a few things:

- ▶ They can store more information, because they can now store data on other computers on the network.
- ▶ They can now share and receive information from other users, perhaps even collaborating on the same document.
- ▶ They can use programs that would be too large for their computer to use by itself.

Network Resources

We now have the server to share the resources and the workstation to use them, but what about the resources themselves? A *resource* (as far as the network is concerned) is any item that can be used on a network. Resources can include a broad range of items, but the most important ones include:

- ▶ Printers and other peripherals
- ▶ Files
- ▶ Applications
- ▶ Disk storage

When an office can purchase paper, ribbons, toner, or other consumables for only one, two, or maybe three printers for the entire office, the costs are dramatically lower than the costs for supplying printers at every workstation. Networks also give more storage space to files. Client computers are not always able to handle the overhead involved in storing large files (for example, database files) because they are already heavily involved in the day-to-day work activities of the users. Because servers in a network can be dedicated to only certain functions, a server can be allocated to store all the larger files that are worked with every day, freeing up disk space on client computers. Similarly, applications (programs) no longer need to be on every computer in the office. If the server is capable of handling the overhead an application requires, the application can reside on the server and be used by workstations through a network connection.

**NOTE**

The sharing of applications over a network requires a special arrangement with the application vendor, who may wish to set the price of the application according to the number of users who will be using it. The arrangement allowing multiple users to use a single installation of an application is called a *site license*.

**BEING ON A NETWORK BRINGS RESPONSIBILITIES**

You are part of a community when you are on a network, which means that you need to take responsibility for your actions. First of all, a network is only as secure as the users who use it. You cannot just randomly delete files or move documents from server to server. You do not own your e-mail, so anyone in your company's management can choose to read it. Additionally, printing does not mean that if you send something to print now it will print immediately—yours may not be the first in line to be printed at the shared printer. Plus, if your workstation has also been set up to be a nondedicated server, you cannot turn it off.

Network Operating Systems (NOSs)

PCs use a disk operating system that controls the file system and how the applications communicate with the hard disk. Networks use a network operating system (NOS) to control the communication with resources and the flow of data across the network. The NOS runs on the server. Many companies offer software to start a network. Some of the more popular network operating systems at this time include Unix, Novell's NetWare, and Microsoft's Windows NT Server (or Windows 2000). Although several other NOSs exist, these three are the most popular.

Back in the early days of mainframes, it took a full staff of people working around the clock to keep the machines going. With today's NOSs, servers are able to monitor memory, CPU time, disk space, and peripherals, without a baby-sitter. Each of these operating systems allows processes to respond in a certain way with the processor.

With the new functionality of LANs and WANs, you can be sitting in your office in Milwaukee and carry on a real-time electronic "chat" with a coworker in France, or maybe print an invoice at the home office in

California, or manage someone else's computer from your own while they are on vacation. Gone are the days of disk-passing, phone messages left but not received, or having to wait a month to receive a letter from someone in Hong Kong. NOSs provide this functionality on a network.

Network Resource Access

Now that we have discussed the makeup of a typical network, let's discuss the way resources are accessed on a network. There are generally two resource access models: peer-to-peer and server-based. It is important to choose the appropriate model. How do you decide what type of resource model is needed? You must first think about the following questions:

- ▶ What is the size of the organization?
- ▶ How much security does the company require?
- ▶ What software or hardware does the resource require?
- ▶ How much administration does it need?
- ▶ How much will it cost?
- ▶ Will this resource meet the needs of the organization today and in the future?
- ▶ Will additional training be needed?

Networks today cannot just be put together at the drop of a hat. A lot of planning is required before implementation of a network to ensure that whatever design is chosen will be effective and efficient, and not just for today but for the future as well. It is the forethought of the designer that will create the best network with the least amount of administrative overhead. In each network, it is important that a plan be developed to answer the previous questions. The answers will help decide the type of resource model to be used.

Peer-to-Peer Networks

A peer-to-peer network is a network where the computers act as both workstations and servers. An example of a peer-to-peer resource model is shown in Figure 1.4.

Peer-to-peer networks are great for small, simple, and inexpensive networks. In fact, this model can be set up almost immediately, with little extra hardware required. Windows 3.11, Windows 95, and Windows NT

are popular operating system environments that support a peer-to-peer resource model.




FIGURE 1.4: The peer-to-peer resource model

There is no centralized administration or control in the peer-to-peer resource model. However, this very lack of centralized control can make it difficult to “administer” the network; for the same reason, it’s not very secure. Moreover, because each computer is acting as both a workstation and server, it may not be easy to locate the resources. The person who is in charge of the file may have moved it without anyone’s knowledge. Also, the users who work under this arrangement need more training, because they are not only users but also administrators.

Will this type of network meet the needs of the organization today and in the future? Peer-to-peer resource models are generally considered the right choice for companies where there is no expected future growth. For example, the business might be small, possibly an independent subsidiary of a specialty company, and has no plans on increasing its market size or number of employees. Companies that are expecting growth, on the other hand, should not choose this type of model. Although it could very well meet the needs of the company today, the growth of the company will necessitate making major changes over time. If a company chooses to set up a peer-to-peer resource model simply because it is cheap and easy to install, it could be making a costly mistake. The company’s management may find that it will cost them more in the long run than if they had chosen a server-based resource model.

Server-Based Resource Model

The server-based model is better than the peer-to-peer model for large networks (say 25 users or more) that need a more secure environment and centralized control. Server-based networks use a dedicated, centralized server. All administrative functions and resource sharing are performed from this point. This makes it easier to share resources, perform backups,

and support an almost unlimited number of users. It also offers better security. However, it does need more hardware than that used by the typical workstation/server computer in a peer-to-peer resource model. Additionally, it requires specialized software (the NOS) to manage the server's role in the environment. With the addition of a server and the NOS, server-based networks can easily cost more than peer-to-peer resource models. However, for large networks, it's the only choice. An example of a server-based resource model is shown in Figure 1.5.




FIGURE 1.5: The server-based resource model

Will this type of network meet the needs of the organization today and in the future? Server-based resource models are the desired models for companies that are continually growing or that need to initially support a large environment. Server-based networks offer the flexibility to add more resources and clients almost indefinitely into the future. Hardware costs may be more, but, with the centralized administration, managing resources becomes less time-consuming. Also, only a few administrators need to be trained, and users are only responsible for their own work environment.



TIP

If you are looking for an inexpensive, simple network with very little setup required, and there is really no need for the company to grow in the future, then the peer-to-peer network is the way to go. If you are looking for a network to support many users (more than 25), strong security, and centralized administration, consider the server-based network your only choice.

Whatever you decide, be sure to take the time to plan. A network is not something you can just “throw together.” You don’t want to find out a few months down the road that the type of network you chose does not meet the needs of the company. This could be a timely and costly mistake.

Network Topologies

A *topology* is a way of “laying out” the network. Topologies can be either physical or logical. *Physical topologies* describe how the cables are run. *Logical topologies* describe how the network messages travel. Deciding which type of topology to use is the next step when designing your network.

You must choose the appropriate topology in which to arrange your network. Each type differs by its cost, ease of installation, fault tolerance (how the topology handles problems like cable breaks), and ease of reconfiguration (like adding a new workstation to the existing network).

There are five primary topologies (some of which can be both logical and physical topologies):

- ▶ Bus (can be both logical and physical)
- ▶ Star (physical only)
- ▶ Ring (can be both logical and physical)
- ▶ Mesh (can be both logical and physical)
- ▶ Hybrid (usually physical)

Each topology has its advantages and disadvantages. Chapter 3, “Network Topologies and Types,” covers each of these topologies and compares and contrasts each one.

Network Communications

You have chosen the type of network and arrangement (topology). Now the computers need to understand how to communicate. Network communications use protocols. A *protocol* is a set of rules that govern communications. Protocols detail what “language” the computers are speaking when they talk over a network. If two computers are going to communicate, they both must be using the same protocol.

There are different methods used to describe the different protocols. We will discuss two of the most common types in Chapter 3, “Network Topologies and Types,” and Chapter 4, “The OSI Model.”

Network Architectures

Network architectures define the structure of the network, including hardware, software, and layout. We differentiate each architecture by the hardware and software required to maintain optimum performance levels. The major architectures in use today are Ethernet, Token Ring, ARCNet, and AppleTalk. Chapter 3, “Network Topologies and Types,” covers this information in greater detail.

NETWORK MEDIA

We have taken a look at the types of networks, network architectures, and the way a network communicates. To bring networks together, we use several types of media. A medium is the material on which data is transferred from one point to another. There are two parts to the medium, the network interface card and the cabling. The type of network card you use depends on the type of cable you are using, so let’s discuss cabling first.

Cabling

When the data is passed from one computer to another, it must find its way onto the medium that is used to physically transfer data from computer to computer. This medium is cable. It is the network interface card’s role to prepare the data for transmission, but it is the cable’s role to properly move the data to its intended destination. It is not as simple as just plugging it into the computer. The cabling you choose must support both the network architecture and topology. There are four main types of cabling methods: twisted-pair cable, coaxial cable, fiber-optic cable, and wireless. We’ll summarize all four cabling methods following the brief descriptions below.

The Network Interface Card (NIC)

The network interface card (NIC) provides the physical interface between computer and cabling. It prepares data, sends data, and controls the flow of data. It can also receive and translate data into bytes for the CPU to

understand. It communicates at the Physical layer of the OSI model and comes in many shapes and sizes.

Different NICs are distinguished by the PC bus type and the network for which they are used. This section describes the role of the NIC and how to choose the appropriate one. The following factors should be taken into consideration when choosing a NIC:

- ▶ Preparing data
- ▶ Sending and controlling data
- ▶ Configuration
- ▶ Drivers
- ▶ Compatibility
- ▶ Performance

Preparing Data

In the computer, data moves along buses in parallel, as on a four-lane interstate highway. But on a network cable, data travels in a single stream, as on a one-lane highway. This difference can cause problems transmitting and receiving data, because the paths traveled are not the same. It is the NIC's job to translate the data from the computer into signals that can flow easily along the cable. It does this by translating digital signals into electrical signals (and in the case of fiber-optic NICs, to optical signals).

Sending and Controlling Data

For two computers to send and receive data, the cards must agree on several things. These include the following:

- ▶ The maximum size of the data frames
- ▶ The amount of data sent before giving confirmation
- ▶ The time needed between transmissions
- ▶ The amount of time needed to wait before sending confirmation
- ▶ The amount of data a card can hold
- ▶ The speed at which data transmits

If the cards can agree, then the sending of the data is successful. If the cards cannot agree, then the sending of data does not occur.

In order to successfully send data on the network, you need to make sure the network cards are of the same type (i.e., all Ethernet, all Token Ring, all ARCNet, etc.) and they are connected to the same piece of cable. If you use cards of different types (for example, one Ethernet and one Token Ring), neither of them will be able to communicate with the other (unless you use some kind of gateway device, such as a router).

Additionally, network cards can send data in either full-duplex or half-duplex modes. *Half-duplex communication* means that between the sender and receiver, only one of them can transmit at any one time. In *full-duplex communication*, a computer can send and receive data simultaneously. The main advantage to full-duplex over half-duplex communication is performance. Network cards (specifically Fast Ethernet network cards) can operate twice as fast (200Mbps) in full-duplex mode than they do normally in half-duplex mode (100Mbps).

Configuration

The NIC's configuration includes things like a manufacturer's hardware address, IRQ address, base I/O port address, and base memory address. Some may also use DMA channels to offer better performance.

Each card must have a unique hardware address. If two cards have the same hardware addresses, neither one of them will be able to communicate. For this reason, the IEEE committee has established a standard for hardware addresses and assigns blocks of these addresses to NIC manufacturers, who then hard-wire the addresses into the cards.

Configuring a NIC is similar to configuring any other type of expansion card. The NIC usually needs a unique IRQ channel and I/O address, and possibly a DMA channel. Token Ring cards often have two memory addresses that must be excluded in reserved memory to work properly.

Drivers

For the computer to use the network interface card, it is very important to install the proper device drivers. These drivers communicate directly with the network redirector and adapter. They operate in the Media Access Control sublayer of the Data Link layer of the OSI model.

PC Bus Type

When choosing a NIC, use one that fits the bus type of your PC. If you have more than one type of bus in your PC (for example, a combination ISA/PCI), use a NIC that fits into the fastest type (the PCI, in this case).

This is especially important in servers, as the NIC can very quickly become a bottleneck if this guideline isn't followed.

Performance

The most important goal of the network adapter card is to optimize network performance and minimize the amount of time needed to transfer data packets across the network. There are several ways of doing this, including assigning a DMA channel, using a shared memory adapter, and deciding to allow bus mastering.

If the network card can use DMA channels, then data can move directly from the card's buffer to the computer's memory, bypassing the CPU. A shared memory adapter is a NIC that has its own RAM. This feature allows transfers to and from the computer to happen much more quickly, increasing the performance of the NIC. Shared system memory allows the NIC to use a section of the computer's RAM to process data. Bus mastering lets the card take temporary control of the computer's bus to bypass the CPU and move directly to RAM. This is more expensive, but can improve performance by 20 to 70 percent. However, EISA and MCA cards are the only ones that support bus mastering.

Each of these features can enhance the performance of a network interface card. Most cards today have at least one, if not several, of these features.

Media Access Methods

You have put the network together in a topology. You have told the network how to communicate and send the data, and you have told it how to send the data to another computer. You also have the communications medium in place. The next problem you need to solve is how do you put the data on the cable? What you need now are the *cable access methods*, which define a set of rules for how computers put data on and retrieve it from a network cable. The four methods of data access are:

- ▶ Carrier Sense Multiple Access with Collision Detection (CSMA/CD)
- ▶ Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)
- ▶ Token Passing
- ▶ Polling

For more information on these methods, see Chapter 3, “Network Topologies and Types.”

CONNECTIVITY DEVICES

It’s the cabling that links computer to computer. Most cabling allows networks to be hundreds of feet long. But what if your network needs to be bigger than that? What if you need to connect your LANs to other LANs to make a WAN? What if the architecture you’ve picked for your network is limiting the growth of your network along with the growth of your company? The answer to these questions is found in a special class of networking devices known as *connectivity devices*. These devices allow communications to break the boundaries of local networks and let your computers talk to other computers in the next building, the next city, or the next country.

There are several categories of connectivity devices, but we are going to discuss the six most important and frequently used. They are:

- ▶ Repeaters
- ▶ Hubs
- ▶ Bridges
- ▶ Routers
- ▶ Brouters
- ▶ Gateways

These connectivity devices have made it possible to lengthen the distance of the network to almost unlimited distances. For a complete discussion of these devices, check out Chapter 7, “Network Connectivity Devices.”

WHAT’S NEXT

In the next chapter, you’ll be introduced to computer networks and how they evolved. You’ll learn how to decide whether setting up a network will work for your company, and, if you decide that it will work, this chapter will help you determine the qualities to look for when hiring staff to maintain your network.