

Operating Systems and Networking Software

Part of the inhumanity of the computer is that, once it is competently programmed and working smoothly, it is completely honest.

— Isaac Asimov

This quote by Isaac Asimov points out the basic difference between human intelligence and that which is attributed to computers. True computers can be designed and built to calculate, retain, and retrieve vast amounts of data in microseconds and display it in graphics and color beyond what human language is able to relate.¹ However, computers are programmed devices that are only able to operate on a set of rules designed by humans.

True, there are programs that attempt to give computers a form of artificial intelligence, but being only machines that work within a defined rule set, they can only respond in a completely honest manner. On the other hand, humans are capable of lying at any time and often do. We will not get into the philosophical or psychological reasons for why humans have a tendency toward lying. Whatever their reason may be, humans can be whimsical, whereas when a computer acts in that manner, it usually gets its guts torn out. So, now aren't you happy you are not a computer?


The essential piece of software each computer requires is an operating system. Without it, a computer would just sit and not do a meaningful piece of work, just like some humans we know. It is the basic process that operates on human requests and responds accordingly, if programmed to act in that manner. The network drivers embedded in the operating system communicate with the portions of a computer that interact with the network. The operating system assists other application programs to communicate with a server that

¹Try to tell the average human to produce a fancy graph on the fly!

is located remotely and can only be reached over the network. There are other programs involved in the network arena, but the purpose of this chapter is to cover the basic computer operating system and how it interacts with network components. There will also be discussion on network operating systems (NOS) and their place in the network.

4.1 Computer Operating System Basics

To understand computer operating systems and their place in the universe, it is essential to first discuss some computer design basics. Everyone by now has heard the acronym CPU (central processing unit). Some may say it means the computer itself, such as a personal computer, without any peripherals attached to it. In days gone by, a CPU could have taken up some serious floor space, filling a large room or many rooms with racks of equipment. Today, a desktop computer has roughly a footprint of one square foot. This represents a significant difference in floor space, but today's CPU also has major advantages in speed, storage, processing power, and energy consumption. Even though modern computers are far more capable than their early predecessors, they still operate pretty similarly when it comes to handling data.



4.1.1 CPU Basics

The CPU is the heart of any computer. Data and instructions flow into it so the data can be manipulated and acted upon in a controlled manner. Data and instructions are stored within the memory system of the computer. Figure 4-1 shows a block diagram of a basic CPU.

The memory storage area can be constructed of various storage devices ranging from semiconductor to magnetic media. For this section, all you need to know is this is where the instructions of a program and the data that program is to operate on reside. The memory interface contains circuitry that provides addressing information to the memory storage devices so that data may be retrieved. Once the data is received, it is passed to circuits that decode the retrieved data to determine if it is an instruction or data that needs to be operated on. If the latter, the appropriate input registers are loaded with the data. If it is determined that the retrieved data is an instruction, the arithmetic logic unit (ALU) is given the instruction. Depending on the instruction the ALU receives, it performs an operation on the data contained in the input

registers and places the result of that operation in the output registers so that data can be moved back to the memory system for storage.

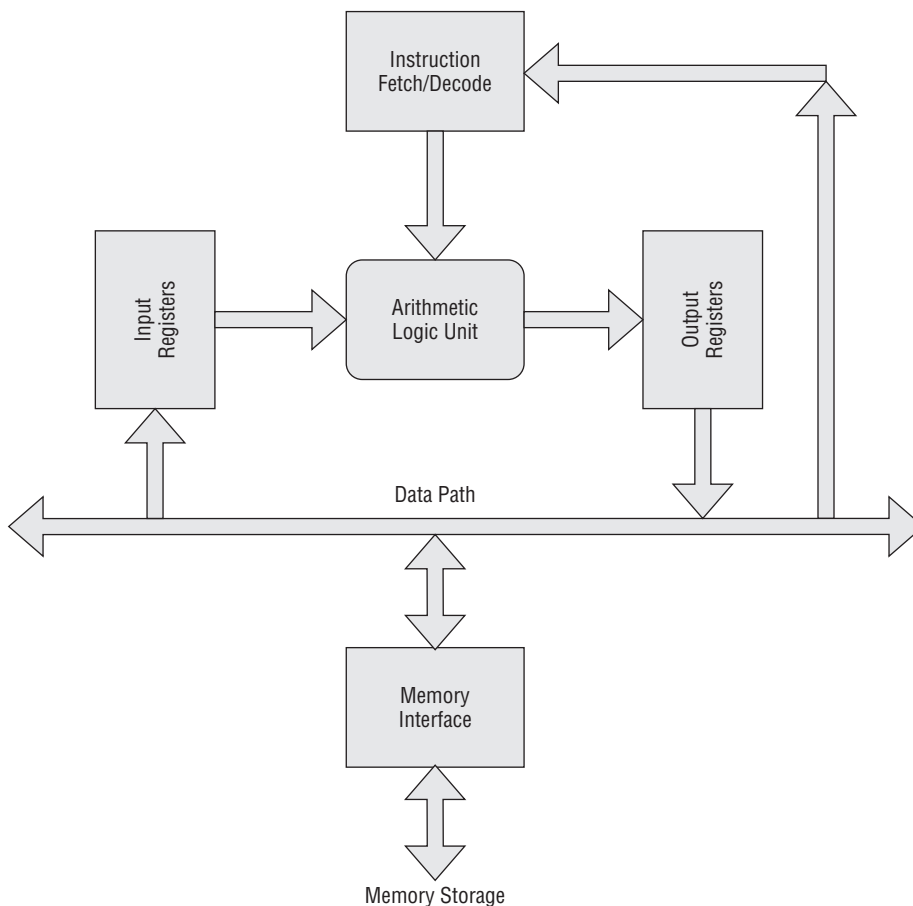


Figure 4-1 A block diagram of a basic CPU

The ALU is the device that performs mathematical operations on the data it is presented with. These are not only the basic functions of addition, subtraction, multiplication, and division, but also Boolean logic² such as `or`, `and`, and their negated logical functions. The ALU is solely responsible for actual mathematical manipulation of the data it is presented with. The remainder of the CPU functional blocks is solely for the purpose of retrieving data and seeing that it is returned to the memory system properly so it can be easily accessed if needed.

²A system of logical operations. The term *Boolean* comes from the name of the inventor of Boolean algebra, George Boole.

QUICK REVIEW

The Boolean algebra **or** function is usually indicated by a **+** sign between variables, such as $A+B=C$. A variable is usually true when its value is equal to 1 and false when its value is equal to 0. An **or** function result is true if any of the variables making up the function is true. A negated **or** function is usually referred to as a **nor** function and its value is false if any of the variables making up the function is true.

The Boolean algebra **and** function is indicated with a **·** sign between variables, such as $A·B=C$. An **and** function result is only true if all of the variables making up the function are true. A negated **and** function is usually referred to as a **nand** function and its value is only false if all the variables making up the function are true.

The following table shows two variables and the resultants of the **or**, **nor**, **and**, and **nand** functions.

A	B	OR	NOR	AND	NAND
0	0	0	1	0	1
0	1	1	0	0	1
1	0	1	0	0	1
1	1	1	0	1	0

This discussion is a simplification of what a CPU is. However, what once took racks of equipment is now contained on a single microprocessor chip. Current microprocessors are magnitudes more powerful than those early computers and use much

POP QUIZ

What function does an arithmetic logic unit provide?

more sophisticated designs that take advantage of bigger data paths, larger addressing capabilities, caching, look-ahead memory fetch,³ parallel and multiple processor technologies — to name a few.

The next section discusses the overall computer architecture and how the CPU interacts with those other computer subsystems.

³A memory fetch grabs the immediate contents of a memory location. Look ahead memory fetch is intuitively retrieving data from memory using the idea that memory fetching is mostly sequential and to save time memory contents would be retrieved in blocks of sequential memory addresses.

4.1.2 Computer Basics

A computer is a collection of subsystems under the control of the operating system, which is the driving intelligence behind the electrical circuits it runs over. Without an operating system, a computer is just a pile of chips, boards, wires, and circuits that would not do any useful thing. But, then again, an operating system is just a collection of ones and zeroes, which is just a bunch of useless information without a computer to execute those commands and instructions. So computers and their operating systems need each other to make a complete package.

In this section, we will be discussing a generic computer system. Most computers have the subsystems being discussed or at least some compatible variation of those subsystems. Figure 4-2 illustrates a block diagram of a basic computer system.

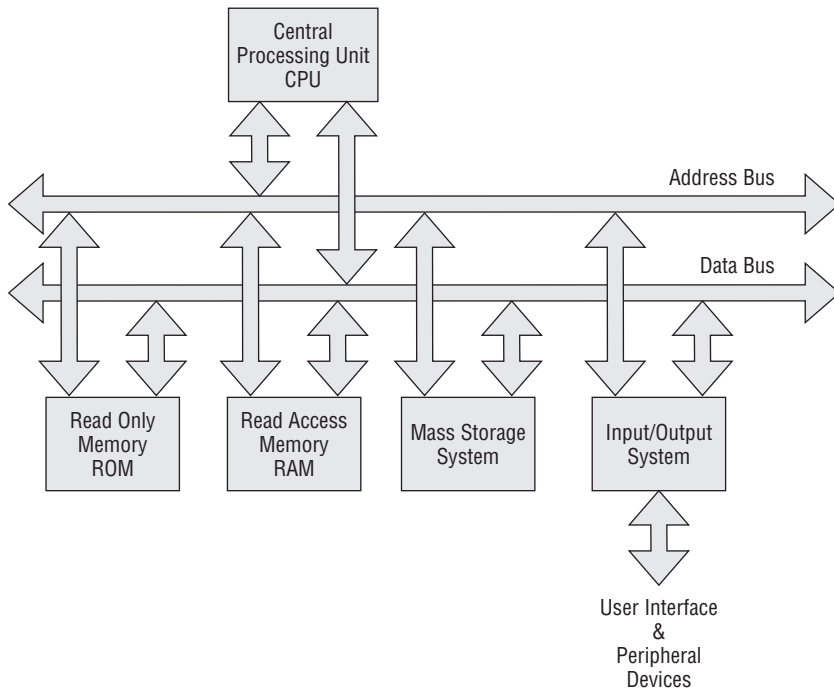


Figure 4-2 A block diagram of a basic computer


We already discussed the CPU portion of a computer. You know that it executes instructions and operates on data, but where is that data obtained? In Figure 4-2, the memory system is distributed across the ROM (read-only memory), RAM (random-access memory), and mass storage System. Why the need for different memory systems? Each has its own purpose within a computer system.

4.1.2.1 Read-Only Memory

When power is first applied⁴ to a computer, commands must be inputted into the CPU to initialize the computer system. A CPU is designed to output an initial address to the address bus to retrieve the first instructions from the ROM. The ROM is a fixed storage subsystem that has the initial boot-up instructions to initialize the system. Most boot-up programs perform both an initialization of the computer and a check of the subsystems to ensure they are functional. The ROM may consist of semiconductor devices that contain bits of the data making up the instructions to be executed that are not alterable by the user. However, current personal computer systems do allow for updates to the ROM software program for bug fixes or feature enhancements. ROM devices in this category are usually called *electrically alterable read-only memory devices*.⁵

This means the device can be written to if necessary using special sequences under control of the operating system. The boot-up code is critical for computer initialization. If this code becomes corrupted for any reason, the computer may not be usable and may require profes-

sional maintenance to restore it back to operation. For this reason, many computers flash warning messages and precautions when the ROM is being accessed under user control. ROMs can be upgraded safely, but do not attempt an upgrade without fully understanding the upgrade process. Typically, once the process has been initiated, it cannot be interrupted until it has completed and the computer has rebooted. If you ever have any questions about upgrading ROM, consult your computer documentation and, if necessary, contact the support staff of the computer's manufacturer.



POP QUIZ

Would it be advisable to cycle power to the computer while a ROM upgrade is in process?

4.1.2.2 Random-Access Memory

Random-access memory (RAM) consists of semiconductor devices that are used for temporary storage of program instructions and data. The usual design is

⁴Technically, you have power within the PC as soon as the battery is plugged in — in other words, when you press the “on” button on the node.

⁵The actual devices used in today's computers are called EEPROM (electrically erasable programmable read-only memory).

an array of these devices residing in the address space of the CPU. As their name implies, they can be accessed randomly no matter what address the information to be retrieved is residing at. This also means the CPU under program control may write data to locations within its address space and store the information for later retrieval. RAM space is usually controlled by the operating system, which designates locations for fixed buffer space for functions under its control and for use by the application programs that may be running at the time. Modern operating systems are capable of running multiple processes at the same time. Each of these processes require operational memory space, so it is critical that memory management be handled properly and as efficiently as possible.

All programs running under the control of the operating system must be well behaved and adhere to the memory space allocation given. When a program violates its memory space allotment, it may overwrite locations being used by other applications or the operating system. If a rogue application overwrites memory used by the operating system for control of the computer, there is a strong likelihood that machine control will be lost and the user will no longer be able to operate the computer under normal conditions. It is in these times that a computer may need to be rebooted to restore operation.

The amount of memory space a computer may contain is determined by how large an address a CPU is able to generate. In the early microprocessor-based PCs, the number of bits of address was only 16, which would allow for a maximum of 65,536 discrete memory addresses. You can determine the address space of a device by taking the number 2 and raising it to the power of the number of address bits that are generated by the CPU. For example:

- $2^{16} = 65,536$ for 16 address bits
- $2^{20} = 1,048,576$ for 20 address bits
- $2^{24} = 16,777,216$ for 24 address bits
- $2^{32} = 4,294,967,296$ for 32 address bits

RANDOM BONUS DEFINITION

active monitor — A node in a Token Ring LAN that is responsible for handling many boundary conditions and housekeeping functions, including generation of a common clock, elastic buffering, and removal of circulating high-priority tokens.

Earlier PCs were mostly character-based computers. Programs were smaller and not as memory-intensive as the visually oriented operating systems of today. As processor capabilities expanded with increased processing speeds and greater addressing ability, software became more sophisticated by taking advantage of these increased capabilities. In the early days, there was a constant battle between hardware designers and their software counterparts. The standing joke used to be that software is like a gas; it will occupy the space

that is provided. This is still pretty much true, but to the software developers' credit, they have done some totally marvelous things with the space they filled.

The real battle lines were drawn on the lines of cost. Hardware had fixed costs and increased rapidly as memory needed to be expanded. Those lines have been obliterated somewhat by the advances in chip design, with increased densities and lower power consumption of newer processor and memory chips.

Costs dropped dramatically and the capabilities of PCs expanded exponentially. This leads to the conclusion that there is a direct correlation between memory size and computer performance. A general rule of thumb is to buy as much memory as you can afford. However, it is really application-dependant. Applications such as gaming software require much more memory and processor speed, whereas someone who just wants to type a few reports can get by on a relatively smaller amount of memory and decreased processor speed. The marketplace puts PCs on the cutting edge of technology as consumers become more sophisticated. It can only keep pushing the demands on memory and processors to increase their abilities, and this is the driving force for today's technology.

POP QUIZ

True or false: The information contained within RAM is saved when the computer is powered off.

4.1.2.3 Mass Storage System

The mass storage system is comprised of a collection of multiple devices storing programs and information either in magnetic or optical media formats. The very earliest PCs used floppy disks to write and retrieve information in a somewhat nonvolatile manner when the computer was powered off. The "somewhat nonvolatile" comment is for anyone who had to suffer through the loss of information due to a flaw in the magnetic media or the electronics of the device controlling this media. If it can be easily written, it can be easily removed or erased.

Just as memory chips underwent improvement, so did magnetic media devices. Floppy disks went from single-sided to double-sided and higher densities. The last floppy disks were high-density 3.5-inch plastic-encased disks that were more reliable than their predecessors but still could suffer similar data losses. The highest density obtained with floppy disks was 1.44 MB, which is a lot for a typewritten document but far from having the capacity to store some of

ACRONYM ALERT

BOOTP — Bootstrap Protocol

today's programs. Programmers had to develop schemes to distribute their software using multiple floppy disks. A user had to sit by the computer during the installation of such a program and wait for the message to load the next disk. The process was tedious and time-consuming.

The development of optical storage devices, such as CD-ROM, increased storage capacities in a movable media format from just over 1 MB to the vicinity of 700 MB. This was a boon to both software developers and computer users. DVD devices, with their higher capacity for data storage, increased what CD-ROM could store by a factor of 10 — or roughly the ability to store 7 GB of information. Current day computers are shipped with optical drives that can read and write both CD-ROM and DVD media formats. Optical media now has read-write capability, but the process is slower than that of magnetic media. However, as a removable media storage system, it has many advantages over its magnetic predecessors. Even though optical disks are more robust as far as data retention, they still can be rendered unusable by physical damage. A severe scratch can make an optical disk unreadable.

Nonremovable disk storage systems are referred to as *hard disks*. They are “hard” because the magnetic media was originally sprayed on the surface of rigid aluminum disks, which were mounted within an enclosed airtight container to eliminate data corruption due to dust and other contaminants.

Magnetic media was bonded to a soft pliable Mylar surface, thus the name “floppy disk.” The advantages of hard disks are their ability to store vast amounts of information and its fast retrieval times. Initially, hard drives were commercially available only to users of large mainframe computers, but as development progressed on these devices, the pricing was such that it was commercially feasible to sell them to the PC market. The first PCs shipped with a whopping hard disk storage capacity of 5 MB. Many of today's graphics-intensive programs would not be able to load onto the drive, let alone the operating system or any other user data. It is not uncommon today to see laptops with 200 GB hard drives and desktops with 500 GB⁶ storage capacities. Hard drives are usually mounted within a computer's case, but many drives are sold as external drives communicating between the drive and computer over the USB port.

POP QUIZ

When a computer is first powered on, the first device it is most likely to read its initial set of instructions from is the _____

⁶This really is an amazing amount of data storage. Can you imagine what increases will be made within the next decade?

4.1.2.4 Input/Output System

A computer is not very useful if information cannot be entered into it or retrieved from it. The *input/output system* is a collection of circuits that allow for information to be entered by the user via a keyboard, pointing device, scanner, etc. It also provides a method for information to be displayed to the user. This can be in the form of video screens, teletype, printers, plotters, etc. These are the most common methods of input and output from a computer system. There are many specialized input/output devices for data entry and retrieval not mentioned in this section, but the idea is always the same: move information into the computer and retrieve it from the computer after it has operated on it.

Because input/output devices interacting with other physical devices and humans may experience timing differences with the CPU, there needs to be a way of storing the information and notifying the CPU when the data is present. Generally two schemes were devised to accomplish this. One is where the input/output devices are mapped to dedicated memory addresses and the CPU polls these locations to see if there is information that needs to be acted on. This is referred to as *memory-mapped I/O*. The other scheme is *interrupt-driven I/O*, where a device writes information into a dedicated register at a fixed port location and sets an interrupt requesting service from the CPU.

In a memory-mapped I/O system, the CPU determines which location it should poll under operating system control. In an interrupt-driven I/O system, the CPU responds to interrupts (and there may be many, depending on the number of I/O ports to be serviced). Interrupts adhere to a fixed interrupt priority scheme, which is hierarchal. The CPU can be processing an interrupt request and be preempted by a higher priority interrupt request.

Regardless of which I/O scheme is used in a computer, the operating system must be able to deal with input/output data requests. It must be able to determine when a device is acting unresponsive and either notify the user or take other action as determined by the program. Generally the operating system is responsible for data movement between the various systems within the computer. However, a user may be running an application, such as a word processor, which is running over the operating system. When a user depresses a key on the keyboard, the operating system reads the key and presents that

RANDOM BONUS DEFINITION

bit stuffing — A technique that provides a unique frame delimiter pattern yet maintains payload data transparency by inserting an extra 0 bit after every occurrence of five 1 bits in the payload data stream.

information to the word processor program, which may request that it also be displayed on the video screen.

On PCs, input/output connections are in the form of ports dedicated to either serial or parallel data communications. *Serial communications* refers to the information being passed one bit for each time interval, which is determined by the

POP QUIZ

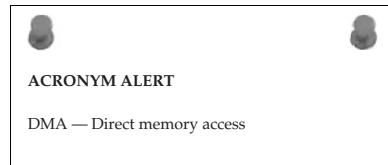
Name a device that you might find connected to a serial port.

speed of the port. Generally serial devices are slow data rate devices such as keyboards, modems, pointing devices, scanners, etc. However, with the development of Universal Serial Bus (USB), high-speed serial ports, devices such as hard disk drives and printers can be used due to the increased data rates on these ports. Parallel ports on older PCs were mostly relegated as printer ports. *Parallel data communications* means that data is sent a whole byte at a time for each cycle of the port. USB has become today's de facto standard for peripheral ports.

4.1.3 Operating System Basics

Operating systems in one form or another have been around since the inception of the first computer. Of course, the first computers were of the mainframe variety with character-oriented terminals.⁷ Users entered commands and data in the form of alphanumeric characters that could be found on any typewriter. Data retrieved from the computer could be displayed on the terminal screen for small queries, or, for larger reports, outputted to a printer.

The most basic form of an operating system is a file manager. It is able to create new files on the storage medium being used. It is also able to catalog the files for easy retrieval and has some sort of indexing ability similar to that of a filing cabinet. Computers and their operating systems were first designed to adopt systems that were similar to the



business practices of those days. The earlier computers were a high-speed filing system able to store, index, and retrieve data faster than a filing clerk.

Operating systems underwent some dramatic revisions with the introduction of the PC. Initially, these operating systems were similar to those found

⁷The first terminals were alpha-character-oriented. They were merely an electronic form of a typewriter. Graphic terminals that could display some sort of graphic (usually at low resolution by today's standards) were a later innovation in terminal design. Terminals connected to the computer via serial cable.

on the larger computers. They too were character-oriented. The major early PCs initially ran on proprietary operating systems such as Apple's DOS (Disk Operating System) and Tandy Radio Shack's TRS-DOS (usually phonetically pronounced *tris-dos*). The first cross-platform PC operating system to gain popularity was Digital Research's CP/M (Control Program for Microcomputers), originally designed to run on Intel 8080/8085 microprocessor-based computers. It migrated to the Zilog Z80 which was capable of executing the Intel 8080-based instruction set and was a mainstay of the Z80-based PCs for a number of years.

The major limitation of CP/M was that it was designed for 8-bit microprocessors and was only capable of addressing 64 KB of memory. As microprocessors moved up in capability, CP/M began to lose ground to other operating systems, mainly Microsoft's MS-DOS. Digital Research did finally release a 16-bit version as CP/M86, but it was not able to compete against the IBM/Microsoft juggernaut.

Initially, MS-DOS was locked up by IBM and was sold with the IBM PC as IBM DOS. Other PC manufacturers were on the outside looking in and attempted to adopt CP/M86, but the popularity of the IBM PC running MS-DOS left them far behind on the number of PCs being sold. The off-brand manufacturers eventually developed clone PCs that were able to run MS-DOS, thus boosting their PC sales. The developer of CP/M and CP/M86, Digital Research, also developed a clone to MS-DOS called DR-DOS to compete with Microsoft. The number of PCs now running MS-DOS caused IBM to lose their competitive edge and to eventually give up on the PC market.

Although CP/M was a cross-platform operating system, the hardware it was running over could have major differences. As a result, a CP/M program on one computer could not run on another computer from a different manufacturer. The portability of CP/M was the core operating system (sometimes referred to as the *kernel*). The CP/M kernel provided a common interface for user input and application programs that would run over different computer platforms. The computer manufacturers had their own software designer teams that would write the software code needed to allow the kernel to communicate with other hardware systems of the computer system. These pieces of code were referred to as *hardware drivers*.⁸ Each subsystem in a computer system could have its own driver if needed. An example of this is the mass storage subsystem. The kernel would call for a file and the driver would cause the floppy drive to seek the track and sector where the beginning of the file was located. The point is, although there was commonality as far as user interfaces and the applications able to run on CP/M, they could have been

⁸Hardware drivers are synonymous with device drivers. It is the code that is designed to allow the kernel of the operating system to properly communicate with the device/hardware no matter how different in design they may be. The device driver acts as a translator to allow for the correct operation of the device/hardware.

operating on computers whose hardware had substantial differences from one manufacturer to the next.

Soon after the IBM PC was introduced and its hardware specifications were published, clone PCs began to enter the marketplace. Since IBM opened its architecture, it was not able to legally protect its design, and

RANDOM BONUS DEFINITION

byte — A unit of data that is equal to 8 bits.

the PC marketplace ballooned overnight with clones from a number of hardware manufacturers. This phenomenon led to a PC base that not only was able to have the same operating system but also had hardware commonality, which was a boon to the peripheral manufacturers.

With the consolidation of today's PC marketplace, there are really only two variations of PCs. Today's PC users are either in the Apple Mac domain or the PC domain (PCs from various manufacturers able to run the various iterations of Microsoft DOS). Today, Apple

POP QUIZ

What is the acronym for a user interface that uses a point-and-click method of executing computer commands?

manufactures and markets laptops and desktop PCs based on its Macintosh family of computers. Macs were the first PCs that took advantage of a point-and-click-based operating system.⁹

Today's PC world is divided between the Mac operating system and Microsoft Windows operating system. Both are GUI (graphical user interface) based and use a graphical display screen and some sort of pointing device. However, even with the whiz-bang colorful interfaces, the operating system is basically performing the same functions as its predecessors. The only difference is that instead of parsing text instructions, the user input interpreter uses positional information, and if a mouse is used, a right, left, or double-click will cause the operating system to act on the object that is being pointed at on the video graphical display screen.

4.2 Network Operating System Basics

As the need grew for PCs to interconnect and share data and common resources, the opportunity arose for the design and marketing of network

⁹If this had caught on before Windows came out, it might have been a much different world today.

operating systems. The most common design of network operating systems was the client/server implementation. PCs were clustered for individual users (clients) to share files on the file server or print data files on printers under the control of a print server. Figure 4-3 illustrates an example of network running a network operating system (NOS).




Figure 4-3 A computer network under the control of a network operating system

Actually, “network operating system” is a bit of a misnomer in that the NOS really runs on computers that are servers placed in the network. Figure 4-3 shows a single¹⁰ file server and a single¹¹ print server. In reality, on large networks there could be multiple servers in use. Also, for a small office, the functionalities of both the file server and print server can be combined in a single server. Being a client/server application, the responsibility for authentication of clients with the authority to connect to the server depends on the server to verify that clients have the necessary valid security credentials. In larger networks with many clients, that function can be placed in entirely separate servers solely responsible for granting network access as well as the permission levels a user will have while logged into the network.

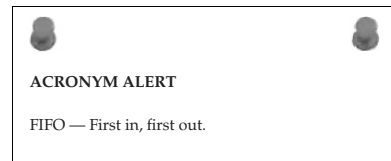
There are networks where the software that is being run on a local PC is actually an application located on the server. An example of this is a word processor program that has a fixed number of network licenses. The theory is that not all users would use the program simultaneously, so a company could save some costs by sharing applications over the network. Once all the licenses

¹⁰Just because they are single does not mean they are available.

¹¹See footnote 7.

are occupied, subsequent users would need to wait until another user logged out of the program, thus releasing the license. Users could be prevented from loading a program from a server if the network or the server is being heavily worked. Once the program is downloaded to the local PC, there is no further network interaction required until the application is released by the user. This interaction was called the *file services* portion of the NOS.

Print services were also an important piece of NOS. Printer requests would be queued to the print server servicing that portion of the network. A print server could have one or many printers under its control. As print job requests arrived at the print server, it would determine the printer the print job was to be outputted to. The



print server queued the print jobs on a first-come, first-served basis. Print jobs were stored on the print server and parceled out to the printer as fast as the printer was able to take the data. Today's network-ready printers are basically their own print server with the intelligence and storage capacity required to queue print jobs from a large user base.

There were many networking operating systems, but the most popular were Novell NetWare and Microsoft Networking. Novell utilized an IPX/SPX protocol stack to provide communications over its network. Both Novell and Microsoft have since migrated to supporting the TCP/IP protocol suite over their networks. TCP/IP is not a NOS; it is a protocol that controls communications between peers. A client/server application can be run over a network that uses TCP/IP protocol for communicating over the network, but the actual client/server application is independent on the protocol itself.

The majority of today's networks are TCP/IP-based networks that have a wide range of applications running over them. A workstation may have multiple sessions to various servers on the network simultaneously. Most people use e-mail and may be logged into a corporate mail server while running other applications to other servers over the same network. The need for a network server running a NOS is not required when running the TCP/IP protocol over a network.

4.2.1 Peer-to-Peer Networking

When discussing network operating systems, the context of the discussion is usually based around client/server networks. To perform peer-to-peer networking, where one computer can share data and resources with another computer, requires some sort of application program. The earlier versions of peer-to-peer networking were crude and cumbersome to configure and use. However, as Microsoft evolved its Windows operating system, they added

peer-to-peer as well as workgroup network capabilities. Windows was the first GUI-based operating system that was able to support this type of networking.

Windows users are able to share drive space and locally attached printers with other users on the same network using what is commonly referred to as *Windows networking*. Windows networking depends on the host names of each computer to be different if they reside within the same network. This was first accomplished with NetBIOS API (application programming interface) running on each Windows computer on the network. In today's networks, NetBIOS is usually run over the TCP/IP protocol. In this scenario, each computer has both a unique computer name and an IP address. The services NetBIOS provides are related to the Session layer of the OSI model.

On smaller networks, the computer broadcasts the name of the computer that it wants to establish a session with. On large networks, broadcasts can become intrusive and affect network throughput speeds. Large Windows networks will utilize a WINS (Windows Internet Name Service) server for computer name resolution. It maps computer host names to network addresses, thus eliminating multiple broadcasts on the network. WINS can be thought of as the name service for NetBIOS networks and is similar to a DNS (Domain Name Service) server in operation on a TCP/IP network.

Figure 4-4 shows a small peer-to-peer Windows-based network.

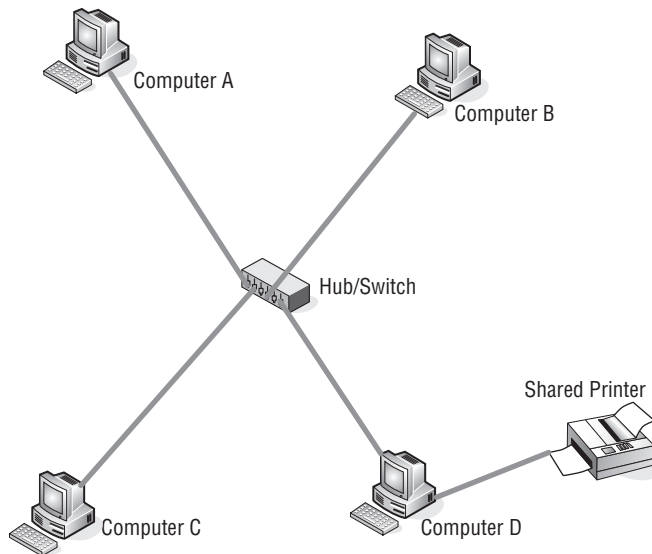


Figure 4-4 A small, Windows-based peer-to-peer network

In this figure, the PCs are labeled Computer A, B, C, and D. However, they may be named in any manner a user or network administrator chooses. It is

a good idea to select meaningful names such as `joes_pc`, `jims_pc`, and so on, to give a frame of reference for the PC. In larger companies, the computers may be named by department and function. Naming is purely arbitrary, but knowing what each PC is named can be helpful, especially when trying to troubleshoot network issues.

Within this network, NetBIOS provides computer name registration and resolution, a connection-oriented communication session service, and a connectionless communication for datagram distribution service.

RANDOM BONUS DEFINITION

`cheapernet` — Another name for 10BASE2.

Before a computer can either start a session or distribute datagrams on the network, it must use the NetBIOS name service to register its name. NetBIOS utilizes UDP port 137 for the name service. The NetBIOS name service functions are to add a name or group name, delete a name or group name, or find a name on the network.

Since in today's networks NetBIOS is run over TCP/IP, NBT (NetBIOS over TCP/IP) utilizes TCP port 139 for the session service. The session mode of NBT allows two computers to establish a connection to pass communications between them. The NetBIOS primitives¹² associated with the session service are as follows:

- **Call** — Opens a session to a remote computer using its NetBIOS name.
- **Listen** — Listens for session requests using NetBIOS name.
- **Hang Up** — Ends a session that had been previously established.
- **Send** — Sends a packet to the computer that a session has been established with.
- **Send No ACK** — Similar to Send but does not require a returned acknowledgement that the packet was received.
- **Receive** — Waits for the arrival of a packet from a computer a session has been established with.

The *datagram distribution service* is a connectionless service where messages are sent without regard to error detection or remediation. It is incumbent upon the application using this service to provide the necessary data error detection and recovery when needed. UDP port 138 is used by NBT for this datagram distribution service.

¹²This list is almost the same responses that one can expect from the family teenager. However, for a NetBIOS session these are the root terms used to describe a particular sequence within the session.

The primitives used for datagram distribution by NetBIOS are as follows:

- **Send Datagram** — Sends a datagram to a remote computer using its NetBIOS name.
- **Send Broadcast Datagram** — Sends a datagram to all the NetBIOS names that are registered on the network.
- **Receive Datagram** — Waits for the arrival of a packet from a Send Datagram process.
- **Receive Broadcast Datagram** — Waits for the arrival of a packet from a Send Broadcast Datagram process.

Fortunately, setting up a small Windows-based local network is easy to do. The previous discussion in this section gives you an appreciation of what is going on under that colorful GUI screen. The unfortunate part is that Windows, with all its various generations, had added twists and bends to the methods used to configure networking on a PC using the Windows operating system for its OS. It is the author's recommendation to review the documentation for your particular version of Windows before attempting to configure your PC for networking. The configuration overview as well as the screenshots in the remainder of this section are based on Windows XP.

Most of the PCs purchased within the last couple of years come pretty much network-ready. Many desktops come with an Ethernet NIC card¹³ installed, and many laptops not only have a hard-wired NIC for Ethernet connectivity but also have some sort of wireless connection interface. However, if you have an older PC that you would like to add to your network and it does not have a NIC installed, you have choices available to you to make your PC network-ready. Desktop computer models may either use an internal card, if there is an interface card slot available, or some sort of external solution. There are network interfaces available that will plug into the USB port. If you are not all that computer savvy, I recommend taking down as much information you have about your PC and visiting your local computer store. The sales clerk or computer support staff should be able to assist you in purchasing the appropriate solution to make your computer network-capable.

Older laptops can be easily made network-ready with the addition of a network PCMCIA card. The usual choice is either a card that supports a hard-wired Ethernet solution or a WLAN PCMCIA card, which enables you to connect to your local network wirelessly. The choice is solely dependent upon the current installed network. If this is an initial setup, I strongly suggest investigating a wireless solution. The beauty of a laptop is its mobility, and to have it tethered by an Ethernet cable may not be the ultimate network solution.

¹³Keep in mind, NIC = network interface controller.

NICs require drivers to be able to interoperate with the operating system. Windows has moved to the plug-and-play philosophy where the Windows operating system detects when new hardware has been installed. In most cases, with interface cards from larger manufacturers there is a high probability that Windows will have and load the appropriate driver. If your card is one that Windows is unable to auto-detect, the Windows wizard may request that you load a driver disk to complete the installation of the card. In most cases, there is usually a disk in the box with the card or documentation that will point you to a website or FTP server where the appropriate driver¹⁴ can be downloaded.¹⁵ You can use that downloaded file to complete the installation of the card.



With your wired Ethernet Interface installed, you can navigate to your local area connections properties. On Windows XP, click Start > Control Panel. On the Control Panel screen, select Network Connections for the classic view, or if using category view, select Network and Internet Connections. Select the Local Area Connection that is associated with the NIC card you have installed. With the icon for the interface selected, right-click and scroll to Properties. A window should appear labeled Local Area Connection Properties, similar to Figure 4-5.

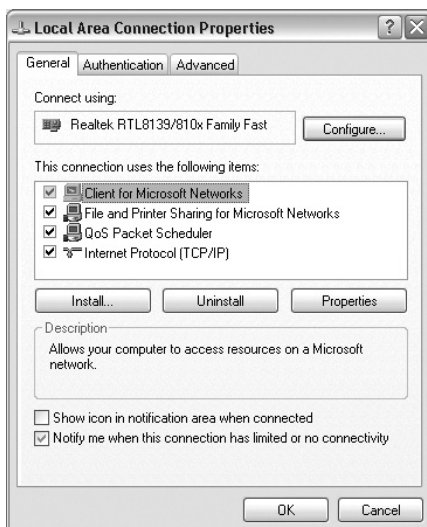


Figure 4-5 Windows XP Local Area Connection Properties

¹⁴Not to be confused with diver, one who deliberately jumps headfirst into water.

¹⁵Assuming that you have another computer that has network capability and is able to reach the Internet to get the file to download.

On this PC, Client for Microsoft Networks is already installed and enabled. If it is not yet installed on your PC, select the Install button and a new window will open labeled Select Network Component Type.

Select the Client component and click on the Add button. The Select Network Client window will open. Select Client for Microsoft Networks and click OK. If you want to share parts of your file system or locally attached printers, you must enable File and Print Sharing. In the Local Area Connection Properties window, click the Install button. When the Select Network Component type window appears, select Service and click on the Add button. The Select Network Service window will appear. Select File and Printer Sharing for Microsoft Networks, and then click OK. You now have Microsoft Networking enabled with file and printer services enabled. We will revisit both file and printer sharing in a bit. For now, it's on to how we get TCP/IP on this puppy.

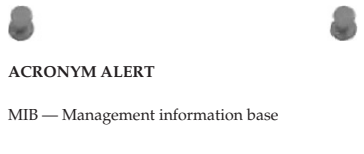
If you do not see Internet Protocol (TCP/IP) in the Local Area Connection Properties window, the protocol must be added. Click on the Install button in the Local Area Connection Properties window. When the Select Network Component window appears, select Protocol and click on the Add button. On the Select Network Protocol window, select Internet Protocol (TCP/IP) and click OK. The protocol has now been installed but must be configured.

Before getting into the configuration of TCP/IP on this Windows PC, a brief description is in order of the difference between a statically assigned IP address and an IP address that has been assigned by a server acting as a DHCP server. This topic will be covered and mentioned in other chapters, and by the time this book

is finished there will be no question that you will know the differences and how they come to be assigned. First, a statically assigned IP address is pretty obvious. It is an IP address that is assigned to the PC by a user or administrator and is the same IP address the computer will have assigned to it each time the PC is booted up.¹⁶ The only things that have to be known prior to assigning the static IP address is that the IP address is unique and not assigned to another computer on the same network segment, that the address to be assigned fits into the addressing scheme being used on that network segment, and, lastly, that the subnet mask assigned with the IP address is compatible with the IP

RANDOM BONUS DEFINITION

flooding — The process of sending a frame to all of a switch's ports, with the exception of the port the frame came in on.



¹⁶What it is *not* is an address that is applied via a static charge.

address and is the subnet mask assigned to that network segment. Static IP address assignment is not difficult in a small network, but it can become rather unwieldy in a large network. And if a network redesign is required with a change in IP address assignment for that network, it can become a support nightmare in very short order. If it can be avoided on the network you are setting up, it is recommended to do so and use a DHCP server for IP address assignment.

So, how does one come up with a DHCP server for their network? Of course, you could have an actual server running a DHCP service, but for a small network, such as that shown in Figure 4-5, it would be a waste of resources. There are many newer network devices that do run a DHCP service if configured to do so. Most routers, both wired and wireless, are capable of running a DHCP service. If the hub/switch shown in Figure 4-5 were replaced by a mini-router like those used for cable/DSL Internet access, you could have a DHCP service running on that network. The beauty of having a local DHCP server is that if there is ever a need to change a network's addressing scheme, default gateway, or the DNS servers being used, there is just a single point that requires configuration change. So there is a major support advantage of running a DHCP service on your network. It is easy to see the advantages of having such a service on large networks with many PCs. One reason to consider DHCP even for a small network is if there are laptops being used. The advantage of using a laptop for a PC is its portability and its mobility of moving from one network to another. Although it is doable, having to configure your TCP/IP setting each time you move from one network to another can grow old very quickly.

To set the IP properties of the installed NIC, click on Start > Settings > Control Panel. On the Control Panel, select Network Connections. Right-click on the Local Area Connection you are going to configure IP addressing on, and then select Properties. Select Internet Protocol (TCP/IP) and click on the Properties button. The window where properties can be configured will appear and look similar to that shown in Figure 4-6.

Notice that this interface is configured for obtaining an address dynamically from a DHCP server somewhere on the existing network. To do this, only the two radio buttons to automatically select these addresses need be selected. However, if you select to statically assign the IP address, each of the grayed fields needs to be filled in with the appropriate information.

- **IP address** — A unique IP address that is not currently used on the network segment where the computer is to be connected

RANDOM BONUS DEFINITION

host — Any node in an IP network.

- **Subnet mask** — The subnet address assigned to the network segment that the computer is to connected to.
- **Default gateway**¹⁷ — The IP address of the node that acts as the default gateway for the network segment the computer is connected to.

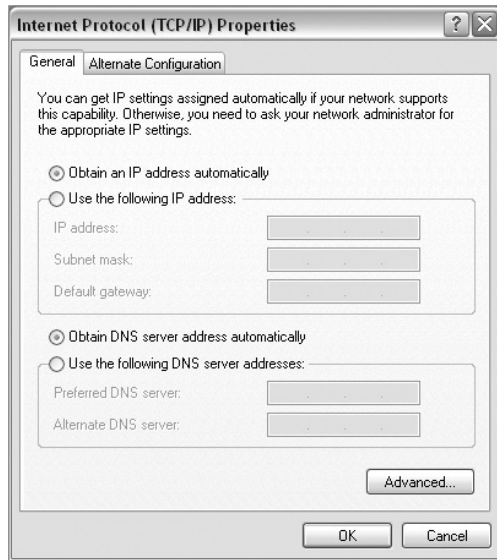
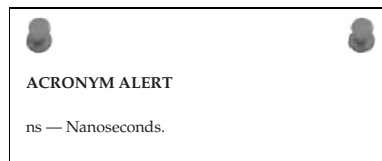


Figure 4-6 Windows XP Internet Protocol (TCP/IP) Properties screen

The DNS (Domain Name Service) server is required if the computer is going to attempt to connect to remote computers by using a domain name.¹⁸ In Figures 4-3 and 4-4 the networks are self-contained and it is assumed that someone is keeping track of IP addresses that have been assigned. In those situations, there is no need for a DNS server to reach the other PCs on the network. Each user will need a list of what those IP addresses are for all computers and other network resources, such as printers. However, in



¹⁷A quick definition of a default gateway is that it is the IP address of a node that is used when a computer needs to start a session with a computer that is not resident on the same network.

¹⁸A domain name server is a computer residing on the Internet providing requested services. For example, a web server may have a name like `www.mywebsite.com`. Since the IP protocol is dependent upon finding an address using numerical addresses, someone needs to resolve the name to a numeric address. This is the role of a DNS server and it gets its information from the authoritative service on the Internet where the name has been registered.

this current interconnected world the need for DNS is paramount. Figure 4-7 shows a small local network connected to the Internet using a router with a high-speed connection.

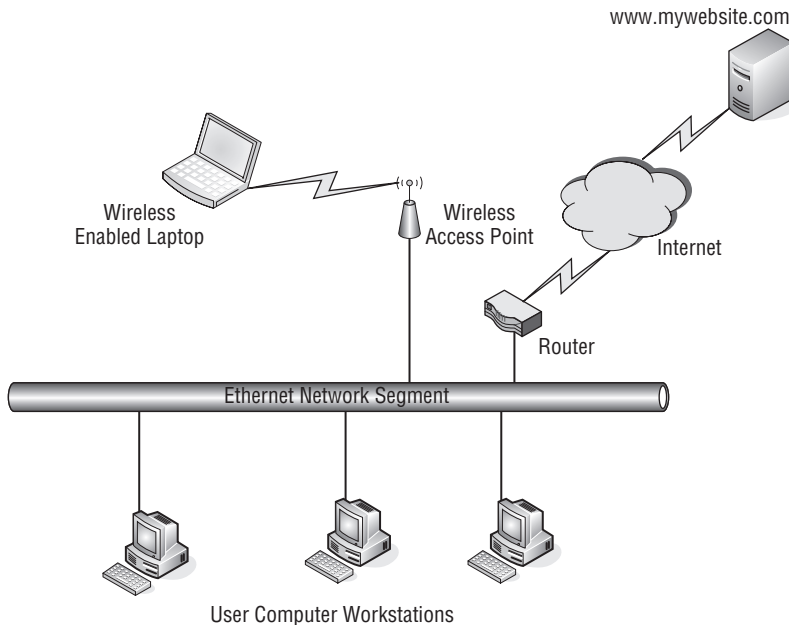


Figure 4-7 A small local network connected to the Internet

Usually when a user or company signs up with an Internet service provider (ISP), they are provided information such as the public IP address that is to be used on the router and its default gateway's IP address. The ISP also provides local DNS service located within the ISP's network, which can be pointed to for DNS name resolution. In a statically assigned IP scheme, these addresses would need to be entered in the appropriate fields of the Internet Protocols (TCP/IP) Properties window to enable the computer to query the provided DNS servers for name resolution when needed. This will need to be done for every computer on the network if they are to be able to connect to computers by IP host name. Most ISPs provide two DNS server addresses. Normally these would be called a *primary DNS address* and a *secondary DNS address*. The primary DNS address is entered in the Preferred DNS server box, whereas the secondary DNS address is entered in the Alternate DNS server box. The PC is now configured to communicate with other PCs on the local network and other computers that may be found on the Internet.

QUICK TIP

There are a couple of quick tests you may want to perform to verify the operation of the NIC card and the connectivity to the local network and the Internet.

1. Click on the Start button in the lower-left portion of your Windows screen.
2. Select Run.
3. In the Run window, enter `cmd` and click OK. A DOS window will open where DOS commands can be entered.
4. Type the command `ping 127.0.0.1`. You should receive back four messages stating “reply from 127.0.0.1.” This indicates that your NIC card is working properly with Ethernet and TCP/IP. If you receive “Request timed out” messages, your card has not been properly configured.
5. To verify your network connectivity, attempt to ping the local default gateway¹⁹ for your network. If you get “Request timed out” messages, verify your physical connection to the LAN.
6. If you get good responses back from the local default gateway, you may want to also check your connection to the Internet.
7. Ping the IP address of the router’s default gateway. If you get good responses, you are able to reach the Internet. If you receive “Request timed out messages” and you own the whole network, you will need to troubleshoot further. If you are on a company network, contact your network administrator.
8. DNS name resolution can be quickly checked if the Internet connectivity test passed successfully. Ping an Internet connected computer by its host name. For example, ping `www.mywebsite.com`. Receiving “Request timed out” messages may not be an indication of a problem with DNS. Some sites drop ping requests in order to combat denial-of-service attacks of their site. What you would want to see is that the name has been resolved to a numeric IP address. If so, then DNS appears to be working properly and you should be able to connect to the site using your web browser.
9. If DNS resolution does not appear to be working, verify the address you had entered on the Internet Protocol (TCP/IP) Properties. If there are no typos, you may want to attempt to ping the IP address of the DNS server. If there are no replies, you may want to attempt to ping the secondary DNS IP address. If you get a reply there, you may want to place the secondary DNS IP address in the preferred DNS server address field and test again, pinging by Internet host name. If problem persists, contact your ISP or your network administrator.

¹⁹This is the IP address inserted in the Internet Protocol (TCP/IP) Properties for the Default Gateway field. A default gateway is normally the IP address of a router located on your network that has access to the Internet.

This section configures a Windows-based PC not only for use on a Microsoft network but also for any TCP/IP-based network, which includes the Internet as we know it today. There will be changes coming such as IPv6,²⁰ but the basics will remain the basics. What is learned here is scalable to any new nuances that may be coming into the world of networking.

POP QUIZ

Name two network operating systems that are prominent in today's networking world.

4.2.1.1 File Sharing on a Peer-to-Peer Network

When we configured the NIC card on the PC to permit file sharing, we did not expound on how this is accomplished in a Microsoft Windows world. The strategy is to first determine what is needed to be shared between users. Whole drives, including hard drives, floppy drives, CD-ROM drives, and DVD drives, can be shared. However, any portion of the file system can be shared down to the lowest subdirectory within a directory structure. So this allows for drive, directory, and subdirectory file sharing, all of which can be accomplished over the local network.

From My Computer, right-click on the drive that you are willing to share. From the drop-down menu, select Sharing and Security. A new window will open showing the properties for the drive (see Figure 4-8).



Figure 4-8 Windows XP drive properties

²⁰We will cover this in Chapter 10, "The Network Layer."

Notice the message about the security risk that is involved in sharing a whole hard drive. You can proceed if you wish or you can back off to the directory you want to share. Multiple directories can be shared on a hard drive.

QUICK TIP

Proper planning can simplify sharing of directories over the network. Create a single folder that you want to share. Under that folder you can create other folders (subdirectories) that will be shared with the parent folder. The whole directory tree under the shared folder will be shared when you allow sharing on this folder.

One instance where it makes sense to share an entire drive is where removable media is concerned. Floppy drives, CD drives, and DVD drives can be both read and written to, as needed. The floppy drive is nowhere to be found on today's newer laptops, so if you need to generate a floppy disk with information from your laptop, share the drive on the desktop to accomplish that task. Granted, it may not be as fast as a directly connected floppy drive, but it can get you by in a pinch.

Enabling file sharing is only half of the task. You may want to create user accounts on the PC. This can be accomplished under the User Accounts section of the Control Panel. For other computers to use the shared folder, they will need to map a network drive. This can be done from My Computer by selecting the Tools drop-down menu and then Map Network Drive. This window is illustrated in Figure 4-9.

The format shown on this window is `\\server`, which would be the NetBIOS computer name of the computer where the shared directory is located. An example would be `\\joe_pc`. However, with TCP/IP enabled on the network connection, this also may be an IP address of the computer where the shared directory is located. The command format would be similar but with the IP address of the computer is placed where the computer name had been. An example would be `\\192.168.5.154`. The `\share` is the name assigned to the shared entity,

RANDOM BONUS DEFINITION

router — A network node that operates at the Network layer.

POP QUIZ

What can be shared using Windows file sharing?

whether it is a drive or directory on the hard drive. The naming is fairly arbitrary and the owner of the computer can use any name he or she pleases. However, the owner must play nice and give the name to the user who would be sharing the data contained in that directory. Without the proper shared name, the share cannot be established. If a guest account or user account has been created for that user, they will be prompted for the account prior to gaining access to the shared data. However, for file sharing to work properly, the computer with the shared directory must be powered on and connected to the network before its shared resources can be accessed.





Figure 4-9 Windows XP Map Network Drive screen

4.2.1.2 Printer Sharing on a Peer-to-Peer Network

In today's networking world there are network-ready printers that act as their own print server. They can obtain a network IP address, be given a name, and will allow themselves to be mapped to from other computers connected to the network. This section does not deal with those printers but with the printers that are locally connected to computer on the network.

These printers may be locally connected to a network PC with a parallel port, serial port, or USB port.²¹ To share a locally connected printer, select Printers and Faxes from the Control Panel. Select the printer to be shared by pointing to it and clicking the right-mouse button. In the drop-down menu, select Sharing. A new window similar to the window in Figure 4-10 will appear on the screen.



²¹Extra credit: What is the benefit and the disadvantage for each of the port types? (This is a question that you will have to research — unless you already know).

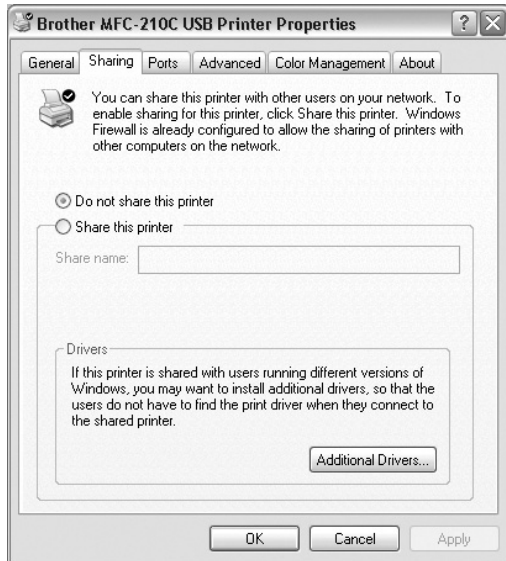


Figure 4-10 Windows XP Printer Sharing screen

Select the radio button to share this printer and enter a share name in the box provided. Windows will attempt to enter a name that is being used locally, but this can be changed as needed. For this example, it a high-speed laser printer connected to Flo the secretary's computer, and other users in the department would like access to that print resource, so a share name may be something like `flo_printer`. Other computer users on the network can then go to Control Panel and select Printers and Faxes and then Add a Printer. They may either browse the network for Flo's printer or enter the name directly, as discussed in the file sharing section. For the sake of this example, the name may appear as `\\flo_pc\flo_printer`, where `flo_pc` is the server name of the computer and `flo_printer` is the share name for the laser printer sitting by her computer. If needed, the IP address assigned to Flo's computer can be used in place of a server name.

QUICK TIP

The use of IP addresses in place of server names is indicative of static IP address assignment. If the network is designed to use dynamic IP address assignment, this could cause problems for users on the network since a computer's assigned IP address could theoretically change each time it is booted up.

Remember that a shared resource in a peer-to-peer network environment assumes that the resource is available on the network. The computer providing the source must be powered on and connected to the network for the resource to be shared.

POP QUIZ

Which printers connected to a network-connected computer can be shared with other users on the network?

4.3 Other Operating Systems

So far in this chapter, we have concentrated on the client aspect of networks and the Internet. However, many computers on the Internet and within the corporate environment are large computers running a wide range of applications. Although there are many similar applications that can run on a PC and offer the same type of service, they may not be equally able to handle many users at the same time. Large computers were initially designed and used to service multiuser environments, whereas the small computer or PC was initially designed with the single user in mind. As a result, the operating systems that control these large machines are much more robust when it comes to handling a large number of simultaneous users.

This section will concentrate on the network aspects of these operating systems and how they are used within both the corporate network environment and the Internet.

4.3.1 Unix

Unix was first developed by AT&T Bell Labs as a multiuser operating system. It was initially designed to handle many users connected simultaneously and all sitting in front of character-based terminals. These terminals were connected to terminal concentrators that were able to aggregate a number of users for ease of communications with the computer the Unix operating system was running on. TCP/IP had not been implemented and the Internet was in its earliest planning stages.

Since its inception, Unix, because of its kernel design was able to be ported to a number of different computer platforms from a variety of computer manufacturers. Later, the operating system program was emulated and offered by other software vendors and computer manufacturers. The discussion in this section will cover the basics to get a Unix-based computer onto a

TCP/IP network. Since these are usually specialized computers from many manufacturers, it would be difficult to get into specifics for all the variations and iterations, so consider this a familiarization with the requirements to make a Unix-based computer network-able.

Unix is a flat file operating system, which basically means that most of the configuration files are in readable text. Configuration is accomplished using one of the resident text processor programs that are part of the utilities that come with the operating system. The appropriate files can be edited as needed to configure the TCP/IP settings on the computer. Usually, systems of this vintage have system administrators who maintain and update the `/etc/hosts`²² file. The information that needs to be modified includes the following:

RANDOM BONUS DEFINITION

trap — A message that originates from a network management client to a network management server to notify the server of a notable event.

- The host name
- The IP address assigned to the interface
- The subnet mask being used for the network segment the computer resides on
- The IP address of the DNS server that is going to be used
- The default gateway that is residing on the same network segment as the computer

The version of Unix you are working with will determine which files and syntax of commands will need to be used. Luckily, most iterations of Unix have resident help in the form of the *man pages*. These pages are an online manual and the common syntax is `man <command>`, where `<command>` is the command you need help with. You will be informed if the command does not exist. When in doubt, issue the `man` command and you will get a complete description of the command along with the various switches that are used by the command.

Newer versions of Unix come with configuration utility programs that assist with the network settings and configuration. Edits of the related network files are automated for ease of use, but essentially it performs the same edits that an administrator can do with a text editor.

²²The Unix `/etc` directory contains configuration files for devices connected to the computer. The `hosts` file aids in host name to IP address resolution. For further information on the Unix directory structure, including the full contents of the `/etc` directory, consult the operating manual supplied with your Unix system.

The following are a few useful commands for troubleshooting network issues on a Unix computer:

- `arp` — Displays a table that shows the IP address to physical MAC address relation for nodes on the same subnet with the Unix computer. This is useful when there are connectivity issues between the Unix computer and that host. If there is an `arp` entry for the problem node, there is a possible Physical layer issue.

```
arp -a
```

- `ping` — An important troubleshooting command that helps to determine that the TCP/IP stack is configured properly on the Unix computer, that the network interface is configured properly, that the default gateway is reachable, and that domain name services are configured properly.

```
ping 127.0.0.1
```

If no response is received, you need to verify that TCP/IP services have been loaded and are running on the Unix computer.

```
ping <address of default gateway>
```

If no response is received, verify that TCP/IP has been bound to the NIC. Check that the operating system has been configured properly as far as the NIC's hardware address and the proper interrupt request number. If the operating system is configured properly, check for a Physical layer issue.

```
ping <address on another subnet>
```

This verifies that the subnet mask has been properly set in the TCP/IP configuration and that the request is sent to the default gateway correctly. If no response is received, check settings to verify that the default gateway is set correctly in the TCP/IP parameters after you were successful in pinging the default gateway.

```
ping <Internet hostname>23
```

This will verify that the DNS service is correctly configured on the TCP/IP stack. If no response is received, attempt to ping the configured DNS server using its IP address. If no reply is received, there may be a connectivity issue. Repeat the ping test to the default gateway. If that passes, verify the settings in the TCP/IP configuration.



²³Internet host name is the fully qualified domain name (FQDN) of the host server you are attempting to reach. An example of a FQDN for a host name would be `www.google.com`.

- `netstat` — A network status command that will display status and information on the network interfaces²⁴ configured on the Unix computer.

The following are some switches that can be used with the `netstat` command:

- `-a` — Displays information on all interfaces.
- `-i` — Displays configuration information.
- `-n` — Displays IP addresses.
- `-r` — Displays routing table information.
- `ifconfig` — Used to display information on the interfaces that are found on the Unix computer. These interfaces can be Ethernet or other types of interfaces.
- `route` — Used to add static routes to the Unix computer's routing table.
- `traceroute` — A useful tool to show the nodes an echo request²⁵ needs to pass through to reach its intended target. The target address may be either a numeric address or an alphanumeric Internet host name.

```
traceroute <address>
```

POP QUIZ

Which command can be used to verify the TCP/IP stack has been properly configured on a Unix computer?

4.3.2 Linux

Linux²⁶ has many similarities and commonalities to Unix. However, it was designed more for the desktop environment even though it will run on larger computers. The number of Linux variations is too many to mention, and each has its own piece of window dressing when it comes to configuration. Similar to Unix, Linux can be configured with a text editor, if necessary.

The variables that are configured are part of a script that is loaded each time a Linux computer is booted. Therefore, changes in network configuration would require a reboot so that these scripts can be executed with the new variables

²⁴Network interfaces on a computer can be of the LAN variety (NICs) or interfaces for WANs, such as a WAN card for a T1 line.

²⁵Echo request is part of the ICMP protocol primarily utilized by the `ping` command. The ICMP components of a `ping` command are `echo request` (the ping to a target IP address) and `echo reply` (a successful response from that target). `traceroute` uses these components to verify the path by receiving and logging the network nodes that the echo request passed through on its way to the target IP address.

²⁶One of the Unix-like operating systems.

in place. The Linux distribution being used will determine the name of the script. In some distributions, the script responsible for initializing the kernel for networking may have the name `rc.inet1`, whereas the script that starts the networking services may be named `rc.inet2`. Again, the distribution and vintage of Linux being used may cause these file names to be totally different. You should consult the documentation for your Linux version prior to configuring or making network changes on the Linux computer.

The networking information for the kernel runtime can be accessed and displayed through the `/proc` file system. The `/proc` file system is usually mounted when the computer is first booted. If it is not mounted, there will be a message stating that `procfs` is not supported by the kernel. If this is the case, the kernel will need to be recompiled with `procfs` support enabled.

Most Linux distributions come with a set of binaries²⁷ containing all the applications and utilities needed for networking support. These applications and utilities may change from time to time with updates to the kernel and the networking utilities. These updates and applications need to be recompiled in order to be used as part of the Linux operating system.

The following are a few of the basic networking configuration and monitoring commands:

- `hostname` — Sets the name of the computer entered in the `/etc/hosts` file.

```
hostname <name of the computer>
```

- `ifconfig` — Allows the interface to be available to the kernel networking layer. This command is normally a portion of the network initialization script that is executed at system boot-up.

```
ifconfig <interface> <assigned IP address>
```

- The first interface required to be activated is the loopback interface.²⁸ The following `ifconfig` command configures this interface:

```
ifconfig 10 127.0.0.1
```

²⁷Binary files are programs that have already been compiled for the system the program is to be executed on. Since Linux can run over many various platforms, application programs need to be compiled on the computer to execute properly. To save users time, many Linux OS providers have already compiled these programs for the platform they are and are considered to be included binaries with the operating system. An example of different platforms would be those that are built around the Intel family of microprocessors versus those computers that have been designed and built using the Motorola 68000 microprocessor family.

²⁸The loopback interface on a computer is a logical network interface which will allow for testing of applications requiring network connectivity. Using this adapter permits the testing of those applications even though the computer is not connected to a network. An example of this would be a computer that is running as a web server testing itself by launching a web browser and navigating to the loopback IP address of 127.0.0.1. The web browser will bring up the server's own home page. A less sophisticated use is in checking the IP stack of the computer by pinging the IP address 127.0.0.1. If no response is returned, there is a problem with the IP stack of that computer.

The following entry in the host table is inserted upon execution of this command:

```
localhost 127.0.0.1
```

RANDOM BONUS DEFINITION

wire speed — The maximum frame and data rate that is supported on a given interface.

Configuration of an Ethernet interface is accomplished using the following command:

```
ifconfig eth0 <interface address> netmask <interface subnet mask>
```

Status of an Ethernet interface can be obtained by executing the following command:

```
ifconfig eth0
```

- **route** — Used to add or delete routes from the kernel's routing table.

```
route [add|del] [-net | -host] target [if]
```

- **add** — Adds a route.
- **del** — Removes a route.
- **-net** — Specifies it is a network route.
- **-host** — Specifies a host address.
- **target** — Specifies the address of either the network or host.
- **if** — Specifies the network interface the route should be directed to (optional).

To add a default gateway, execute the following command:

```
route add default gw <address of gateway node>
```

- **netstat** — As in Unix, a useful command to verify the operation and status of the Linux network components.

```
netstat [-nr, -i, -ta]
```

- **-nr** — Displays the kernel's route table with IP addresses displayed in dotted numerical notation.
- **-i** — Displays interface statistics for currently configured network interfaces.
- **-ta** — Displays a list of both active and passive TCP sockets. This command option can also be modified to also show UDP (-u), RAW (-w), and Unix sockets (-x).
- **arp** — Displays the kernel's ARP table.

```
arp -a
```

Linux is a very robust and feature-rich operating system that is under constant development and improvement. The commands in this section are just a beginning when it comes to Linux. Much more investigation is required, and the information that is available from a wide range of sources is beyond the scope of this section and book.

POP QUIZ

True or false: The name Linux is a derivative of the words Unix lite.

4.3.3 Sun Solaris

Sun Microsystems initially developed the Solaris operating system for their Sun SPARC workstations. It has been ported to X86 Intel-based computers and is distributed and supported by Sun Microsystems. Like Linux, it has similarities and commonalities with the Unix operating system. The latest release of Sun's operating system is Solaris 10.

Although Solaris-based workstations are capable of operating in a standalone (not networked) environment, the operating system provides strong networking tools to allow it to be interconnected not only to the local LAN but the Internet.

Solaris does provide a number of installation programs that will configure the built-in installations.

Enabling a network interface on a Solaris computer requires the following actions:

1. Install device drivers.
2. Reboot to reconfigure the system.
3. Assign an IP address on the interface.
4. Create a hosts file entry to map the IP address to the host name.
5. Configure the interface to pass traffic.

The IP address is assigned to an interface when the IP address is entered into the hostname file located in the `/etc` directory. As with Unix and Linux, this can be accomplished with the use of a text editor.

An interface is configured to allow IP traffic with the use of the `ifconfig` command. The command can also be issued to verify the operation of an interface and to monitor its health. Issuing the `ifconfig -a` command displays all active interfaces on the computer. Incorrect configuration of an interface will result in an error message being returned stating "no such interface." To enable an interface, issue the following command:

```
/usr/bin/ifconfig eri0 up
```

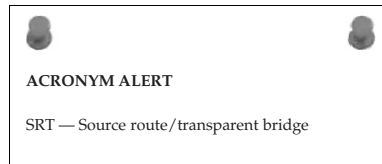
To verify connectivity over TCP/IP with other hosts on the network, issue the following command, which will display the kernel's ARP table:

```
arp -a
```

The flags that can be returned in the ARP table are as follows:

- P — Indicates a published address
- S — Indicates a static address
- U — Indicates an unresolved address
- M — Indicates a mapped address for multicast

Solaris allows for manual tuning of protocol transmission parameters for increased performance. This can be accomplished with the use of the `ndd` command. Using `ndd` parameter options for TCP, UDP, IP, and ARP will display a list of parameter values related to that particular protocol. An example of this would be the issuing of the command `ndd /dev/tcp \?` to display a list of all the parameters that are currently related to TCP.



Like Unix and Linux, Solaris uses the `netstat` command to display network statistics and to verify the operational status of network interfaces.

`netstat` is capable of displaying the following statistics:

- Data collection by protocol type
- Statistics grouped by node address, which may be IPv4, IPv6, or Unix-based
- Data related to DHCP
- Multicast grouped interface data
- Details of the routing table
- Data associated to STREAMS²⁹
- State and status of all IP interfaces
- State of all active logical and physical interfaces, routes, and sockets

`netstat` can display protocol statistics for packets of the following types: TCP, UDP, RAWIP, IPv4, IPv6, ICMPv4, ICMPv6, and IGMP. Each of these

²⁹STREAMS is a flexible programming model used for Unix communications services. It allows for the definition of standard interfaces for character input and output both within the kernel and between the kernel and the rest of the Unix system. It is a collection of system calls, kernel resources, and kernel routines.

packet types has specific parameters associated with it. Generally they display the total number of packets in and out and those that are in error. When monitored, these counters can be used to point out possible problem areas.

Issuing a `netstat -m` command will display the system calls, standard libraries, and kernel associated with writing network applications that use the STREAMS package. Additional details on this function can be obtained by reading the man page for the `streamio` command.

Sun Solaris version 10 can be obtained free of charge from the download section of the Sun Microsystems web page. A version with documentation can be ordered directly from Sun for a nominal charge. If you are interested in learning more about the configuration and maintenance of a Sun system, the X86 version can be loaded on any i86 Intel microprocessor-based computer.

POP QUIZ

List some of the Solaris network commands that are similar to those found in Unix and Linux.

4.4 Chapter Exercises

1. If you have a network-capable PC, try using a few of the network utilities discussed in this chapter.
2. Open a DOS window by running `cmd` from Start, Run. Enter the command `ipconfig` and note what is displayed.
3. Issue the command `ipconfig /all` and note what is displayed.
4. If your network allows your PC to access the Internet, execute this command `tracert <insert your favorite website URL>` and hit the Return key. Note the results. You may want to repeat this with other Internet addresses.
5. To display information about all the interfaces on a Unix computer, which command would need to be issued?
6. What is used on the Internet to find the numeric address of a computer host that resides on the Internet?
7. True or false: Floppy disks are the fastest form of magnetic media.
8. True or false: AT&T is the sole provider for the Unix operating system.
9. Can you name at least one Linux distribution?
10. If a microprocessor designer wanted to allow his newest chip design to access a greater amount of memory space, what might he do to accomplish this?

4.5 Pop Quiz Answers

1. What function does an arithmetic logic unit (ALU) provide?

The ALU performs mathematical operations on the data it is presented with.

2. Would it be advisable to cycle power to the computer while a ROM upgrade is in process?

No.

3. True or false: The information contained within RAM is saved when the computer is powered off.

False.

4. When a computer is first powered on, the first device it is most likely to read its initial instructions from is the ROM.

5. Name a device that you may find connected to a serial port.

Generally serial devices are slow data rate devices such as keyboards, modems, pointing devices, scanners, etc. However, with the development of Universal Serial Bus (USB) high-speed serial ports, devices such as hard disk drives and printers can be used due to the increased data rates on these ports.

6. What is the acronym for a user interface that uses a point-and-click method of executing computer commands?

Graphical user interface (GUI)

7. Name two network operating systems that are prominent in today's networking world.

- Novell Netware
- Microsoft Windows networking

8. What can be shared using Windows file sharing?

- Drives
- Directories
- Subdirectories

9. Which printers connected to a network-connected computer can be shared with other users on the network?

All of the ones designated for sharing.

10. Which command can be used to verify the TCP/IP stack has been properly configured on a Unix computer?

`ping 127.0.0.1`

11. True or false: The name Linux is a derivative of the words Unix lite.
False — The correct answer is Unix-like.
12. List some of the Solaris network commands that are similar to those found in Unix and Linux.
 - netstat
 - ping
 - traceroute

