

In this chapter we introduce you to the Internet and lay bare our intentions for this book.

We start by asking the seminal question, "What is the Internet?" From there, we advance from the philosophical to the practical by offering our suggestions on how to best use and enjoy this book.

Next, we tackle an oft-asked query: "Do I need to learn Unix?" (as well as the only slightly less-asked query, "What is this Unix thing anyway?").

Introduction

What is the Internet?

The Internet is the name for a group of worldwide information resources. These resources are so vast as to be well beyond the comprehension of a single human being. Not only is there no one who understands all of the Internet, there is no one who even understands most of the Internet.

The roots of the Internet lie in a collection of computer networks that were developed in the 1970s. They started with a network called the Arpanet that was sponsored by the United States Department of Defense. The original Arpanet has long since been expanded and replaced, and today its descendants form the global backbone of what we call the Internet.



FUN TIP: Technology marches on: The first experimental network using Internet-like technology involved four computers and was built in 1969. This was 56 years after the invention of the zipper, 37 years after the introduction of the first parking meter, and 13 years prior to the development of the first IBM personal computer.

It would be a mistake, though, to think of the Internet as a computer network, or even a group of computer networks connected to one another. From our point of view, the computer networks are simply the medium that carries the information. The beauty and utility of the Internet lie in the information itself.

Take a moment and skim through the catalog that comprises the last part of this book. Notice the enormous variety. As we start to work together, this is how we want you to think of the Internet: not as a computer network, but as a huge source of practical and enjoyable information.

But this is only the beginning. We would also like you to develop an appreciation of the Internet as a people-oriented society. Put simply, the Internet allows millions of people all over the world to communicate and to share. You communicate by either sending and receiving electronic mail, or by establishing a connection to someone else's computer and typing messages back and forth. You share by participating in discussion groups and by using the many programs and information sources that are available for free.

Does this mean that we are saying that the Internet resources will become as important to you as your telephone and your post office? Yes, that is exactly what we are saying.

In learning how to use the Internet, you are embarking upon a great adventure. You are about to enter a world in which well-mannered people from many different countries and cultures cooperate willingly and share generously. They share their time, their efforts, and their products. (And you will, too.)



For one more moment, take another look at the catalog. Each of those items is there because some person or some group volunteered their time. They had an idea, developed it, created something worthwhile, and then made it available to anyone in the world.

Thus, the Internet is much more than a computer network or an information service. The Internet is living proof that human beings who are able to communicate freely and conveniently will choose to be social and selfless.

The computers are important because they do the grunt work of moving all the data from place to place, and executing the programs that let us access the information. The information itself is important because it offers utility, recreation, and amusement.

But, overall, what is most important is the people. The Internet is the first global forum and the first global library. Anyone can participate, at any time: the Internet never closes. Moreover, no matter who you are, you are always welcome. You will never be excluded for wearing the wrong clothes, having the wrong colored skin, being the wrong religion, or not having enough money.

A cynic might say that the reason the Internet works so well is that there are no leaders. Actually, there is some truth to this. As unbelievable as it sounds, nobody actually "runs" the Internet. Nobody is "in charge" and no single organization pays the cost. The Internet has no laws, no police, and no army. There are no real ways to hurt another person, but there are many ways to be kind. Perhaps, under the circumstances, it is only natural for people to learn how to get along. (Although this does not stop people from arguing.)

What we choose to believe is that, for the first time in history, unlimited numbers of people are able to communicate with ease, and we are finding it is in our nature to be communicative, helpful, curious, and considerate.

That is the Internet.

Using the Internet

Using the Internet means sitting at your computer screen and accessing information. You might be at work, at school, or at home, using virtually any type of computer (including a PC or a Macintosh).

A typical session might begin with you checking your electronic mail. You can read your messages, reply to those that require a response and, perhaps, send a message of your own to a friend in another city.

You might then read a few articles in some of the worldwide discussion groups: jokes from one of the humor groups, or perhaps recipes for a dinner that you are planning for the weekend. Maybe you are following a discussion about Star Trek or philosophy or literature or aviation.

After leaving the discussion groups, you might play a game, or read an electronic magazine, or search for some information on a computer in another country.

This is what using the Internet means and this is what we will show you how to do in this book.

How to Use This Book

The title of this book contains the word "Complete", so let's take a minute to discuss exactly what we mean by that.

We promise to teach you everything you need to know to be able to use the Internet and its basic resources. With the catalog that forms the latter part of the book, we offer a large, representative list of all that is available. Our intention is that for most people, most of the time, this book will be all that you need.

Since the purpose of this book is to show you how to understand and use the Internet, we need to start by providing the proper background and by teaching you some technical details. This is what we will do in Chapters 2, 3, and 4. These chapters explain the basics: how the Internet is organized, how you connect to the Internet and, most important, how to understand Internet addressing.

As you know, when you use a telephone you need a number to dial, and when you mail a letter you need a postal address. Similarly, the Internet has its own official "addresses". Each person and each computer has its own address, and Chapter 4 explains how the system works. Almost everything is built around these addresses, so it is crucial that you read this chapter at the beginning of your Internet career.

Once you have read Chapters 2, 3, and 4, you can go wherever your interests take you. If you want to send electronic mail, start with Chapters 5 and 6. If you want to participate in a Usenet discussion group, read Chapters 9, 10, and 11, and then whichever of chapters 12, 13, 14, or 15 is appropriate. Once you understand the basics, there is no special order for learning.

If you are not sure where to start, take a look at the catalog. Read through until you find something that interests you, and then turn to the chapter that shows you how to access that service.

If you are using this book to teach a course, you can follow the chapters in order. We have carefully designed the material to cover the most important topics first. You can start from the beginning and teach your students (or yourself) as much as you want, leaving the rest for another day.

Do You Need to Know Unix?

The answer is no, but read on.

Unix is a family of operating systems (master control programs) that are used to control computers. Virtually all types of computers can run Unix. Conversely, there are many variations of Unix that run on all sizes of computers.

When we use the word "Unix" as a general term, we mean more than an operating system. We are actually referring to an entire culture with its own language, technical terms, conventions, traditions, and a wide variety of computer-oriented facilities. To many people, the Unix culture is intimately connected to the Internet. Some people consider the Internet to be part of the Unix culture. Other people consider Unix to be part of the Internet culture.

The truth of it is that the Internet very much has a life of its own. Most of the Internet computers use Unix, but the details are hidden from you. Thus, you do not need to learn Unix *per se* in order to use the Internet.

However, you do need to understand the rudiments of using your own computer system. You must know how to start work, enter commands, use the keyboard (and mouse if you have one), and stop work when you are finished. It is also helpful—almost indispensable—to be able to manipulate data files, so you can save and retrieve information, and so you can create and edit your own information. For example, when you send someone a message using electronic mail, it is convenient to be able to use a text editing program to compose the message ahead of time.

When you use a Unix system to access the Internet, it is often difficult to say where Unix ends and the Internet begins. The programs that you will be using will all be Unix programs, perhaps not built into the operating system itself, but very much a part of the Unix culture. For example, many people read their electronic mail using the Unix mail program.

However, you are perfectly justified in keeping your knowledge of Unix to a practical minimum and spending your time, instead, exploring and using the Internet. There is no need for you to become a Unix expert.

Should You Learn Unix?

Having just told you that you do not need to become a Unix expert in order to use the Internet, we are now going to tell you that, yes, we do recommend that you learn Unix. You do not need to, but you ought to, and here is why.

As we mentioned earlier, Unix is a lot more than an operating system. It is actually a large, worldwide culture that is intimately connected to the Internet. On its own, Unix has a lot to offer and, hence, is worth learning.

Moreover, if you are using Unix to access the Internet, there are a great many advantages to having some technical knowledge of the underlying system. For example, you will probably want to send messages, either to another person or to a discussion group. To compose such messages, you should be able to use one of the Unix text editing programs.

As you work with the Internet, you may also want to collect information that you find on computers around the world. This means that you will need to know how to create and manipulate data files.

There are many such examples and they all illustrate the fact of life that, regardless of what anyone tells you, it is always a good idea to learn something about the technical details of your computer system. People have gone to a great deal of trouble to make computer systems easy to use and to make certain parts of the Internet easy to use, but—and this is an important but—you will be far better off if you learn more than the bare minimum.

Wait a minute, we hear some of you saying, that may be fine for the people who are using Unix-based systems, but I use a PC (or a Macintosh or some other fine non-Unix system) to access the Internet, and the people who set everything up promised it would be user friendly. They told me that I could select anything I wanted from a menu and that I wouldn't have to spend any time learning anything as evil and horrible as Unix. Besides, computers scare me.

Well, far be it from us to deny you the comfort and security of your own fears. Moreover, the last thing that we would do in a family-oriented book like this would be to disagree contentiously with those misguided souls who firmly believe that the future of computing lies in the populace-at-large selecting items from menus.

Suffice it to say that, as you use the Internet, you will from time to time connect to remote computers which will be running Unix. At such times, a basic knowledge of the Unix culture and its tools can provide practical benefits.

In order to help you, we recommend the book *A Student's Guide to Unix* by Harley Hahn (McGraw-Hill College Division, 1993). That book and this Internet book complement one another nicely.

Hint

Learning the fundamentals of Unix will make your work with the Internet more comfortable and will expand your possibilities.



CHAPTER 22



In order to use the Internet well, you need to understand something of what it is and how it works. In this chapter, we will start with the idea of computer networks, introduce a few basic ideas and terms, and then move quickly to the Internet itself.

After explaining the basic services upon which the Internet depends, we will survey the large variety of resources you will be able to use. If you have been wondering, "What can the Internet do for me?", this is the chapter that will explain it all.

Understanding the Internet

If you hang around the Internet at all, you will from time to time hear people talk about something mysterious called "TCP/IP". At the end of this chapter we will discuss TCP/IP and show you what it really is. (What? You have never heard of TCP/IP? Obviously, you have been hanging around the wrong sorts of people. Well, we can fix that.)

Our Friend the Network

The term *network* refers to two or more computers connected together. There are a number of reasons to connect computers into networks but the two most important (from our point of view) are:

- to allow human beings to communicate
- to share resources

Once you start using the Internet, you can send messages to anyone else on the Internet. You can even send messages to people who use other networks that are connected to the Internet.

As for sharing, computer managers arrange networks so resources that are expensive or difficult to maintain can be used by anyone on the network. For instance, a manager might attach a costly printer to a network so that everybody who needs it can use the same printer. On the Internet, we share information resources rather than pieces of hardware. For example, the items in the catalog (later in this book) are some of the many resources that can be used by anyone on the Internet.

A *local area network*, or LAN, is a network in which the computers are connected directly, usually by some type of cable. When we connect LANs together, we call it a *wide area network* or WAN. Most wide area networks are connected via leased telephone lines, although a variety of other technologies, such as satellite links, are used as well. The wide area connections for most of the Internet travel over some telephone system or another. Indeed, the bottleneck in establishing Internet service within developing countries is usually due to the lack of a reliable telephone system.

Here is a typical example of a network. Imagine yourself sitting in a room full of computers in the Social Sciences Computing Facility at a major university. (The weather outside happens to be cold and rainy, but since it has nothing to do with the example, we won't mention it.) Your computer is connected in a LAN to all the other computers in the room and to the computers within people's private offices throughout the building. This arrangement is shown in Figure 2-1.

There are a number of other LANs on the campus. For example, the Psychology department has its own network of computers, as does the Math

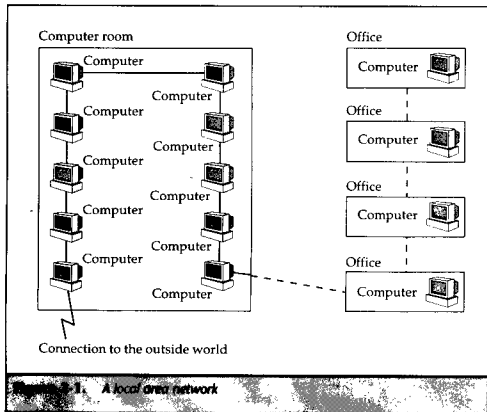


Figure 2-1. A local area network

department, the Computer Science department, and so on. Each of these LANs is connected to a high-speed link, called a *backbone*, to form a campus-wide WAN. This is shown in Figure 2-2.

Although we have used a university as an example, many types of organizations use similar arrangements: companies, governments, research facilities, other types of schools, and so on. If an organization is small, it may have only a single LAN. Large organizations may have multiple LANs connected into one or more complex WANs. Such organizations usually have a full-time staff to care for and feed the networks.

How are the LANs connected? By special-purpose computers called *routers*. The job of a router is to provide a link from one network to another. We use routers to connect LANs (to form WANs) and to connect WANs (to form even larger WANs). In other words, you can consider the computers within the Internet to be connected into LANs and WANs by a large number of routers. However, there is more to the picture that we need to explain. But first, let's take a moment to talk about hosts and terminals.

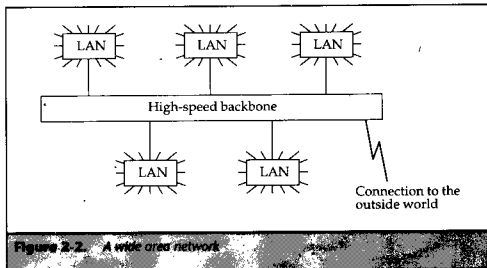


Figure 2-2. A wide area network

Hosts and Terminals

There are two meanings for the word *host* that you should know about.

First, within the Internet, each separate computer is called a host. For example, you might tell someone that he can find the information he wants by connecting to a host in Switzerland. If your computer is connected to the Internet, then it too is a host even though you may not share any resources with the rest of the Internet.

You may also see a computer referred to as a *node*. Here is why. If you draw a diagram of points and lines to represent the connections within a network, each computer will be a point and each connection will be a line. In the part of mathematics that deals with such diagrams, each such point is called a "node". Network specialists have borrowed this term to refer to any computer that is connected to a network. Thus, "node" is a more technical synonym for "host".

Hint

If you are trying to see if someone is a nerd, ask him "What is the technical term for a computer that is attached to the Internet?" If he says "node", chances are, he is a nerd.

The second meaning of the word "host" has to do with how certain computer systems are set up.

In general, there are two ways you might use a computer. You might have it all to yourself, or you might share it. For example, when you use a PC or a Macintosh, you are the only person using the computer. Single-user computers, especially the more powerful ones, are often called *workstations*.

Some computers, however, are made to support more than one user at the same time. These multi-user systems are often referred to as host computers.

A large mainframe computer, for example, can act as a host for hundreds of users at the same time. More commonly, a smaller computer, perhaps one that looks no larger than a PC, will act as a host for a small group of users. The Unix operating system, which we discussed in Chapter 1, is a multi-user system. Although some people use a Unix computer as a personal workstation, many Unix computers are used as hosts to support multiple users.


When you have your own computer, you interact by using the keyboard, screen, and (possibly) a mouse. These devices are part of the computer. With a multi-user computer, each person has his or her own *terminal* to use. A terminal has a keyboard, screen, perhaps a mouse, and not much more. All the terminals are connected to the host, which provides the computing power for everybody. This arrangement, called a *time-sharing system*, is shown in Figure 2-3.

Thus, there are two meanings for the word "host". Within the Internet, each computer is called a host. Within a time-sharing system, the main computer that supports each user on a separate terminal is also called a host. Of course, if such a computer were connected to the Internet, it would be both a time-sharing host and an Internet host.

Client/Server Systems

As you know, one of the principle uses of a network is to allow the sharing of resources. Much of the time, this sharing is implemented by two separate programs, each running on different computers. One program, called the *server*, provides a particular resource. The other program, called the *client*, makes use of that resource.

For instance, say that you are working with a word processing program that is running on your own PC. You tell the program that you want to edit a particular file that is stored on another computer on your network. Your program will pass a message to that computer asking it to send the file. In this case, your word processing program is the client while the program that accepts the request and sends the file is the server. More precisely, it is a file server.



On local area networks, where the hardware is close by and visible, it is common for people to use the word "server" to refer to the actual computer that runs the server program. For example, one can imagine a network manager showing around the President of the United States and, while pointing to an otherwise nondescript beige machine, proudly announcing, "This is our file server." (To which proper etiquette requires that the President respond by nodding his head and saying, "How nice for you.")

On the Internet, hardware is normally not visible, and the terms "client" and "server" usually refer to the programs that ask for and provide services.

Here is an important example. Many Internet sites provide a service called a "Gopher" (which we will discuss in more detail later). Briefly, a Gopher allows you to select items from menus. Each time you select an item, the Gopher performs the required task. For instance, if the item describes a particular piece of information (such as "News of the Day"), the Gopher will retrieve this information and display it for you.

When you use a Gopher, two different programs are involved. First, there is the program that provides your interface. This is the program that interprets your keystrokes, displays the menus, and generally makes sure your requests are carried out. This program is called the Gopher client.

The other program is the one that supplies whatever it is the Gopher client has requested on your behalf. This program is called the Gopher server.

The beauty of this system is that the client and server programs do not necessarily run on the same computer. Indeed, more often than not, the client and server programs reside on different computers. For instance, you may be sitting in front of a PC in Rio Linda, California, using a Gopher to read the "News of the Day" at the National Security Agency in Virginia, three thousand miles away.

In this case, the Gopher client is a program running on your PC, while the Gopher server is a program running on a supercomputer on the other side of the country.

All of the Internet services make use of this *client/server relationship*. Learning how to use the Internet actually means learning how to use each of the client programs. Thus, in order to use an Internet service, you must understand:

1. How to start the client program for that service
2. How to tell the client program which server to use
3. Which commands you must use with that type of client

Your job is to start the client and tell it what to do. The client's job is to connect to the appropriate server and to make sure that your commands are carried out correctly.

Each type of Internet client has its own commands and conventions. For example, the commands you use with a Gopher client are different from the commands you use with, say, an Archie client (another Internet service). Fortunately, there are only a handful of basic services you need to learn.

As you learn about Internet services, you will find that the client programs—such as Gopher or Archie—have many different commands. However, you will also find that you actually need to learn only a few of the most important commands.

The Internet client programs all come with their own built-in help functions. Thus, if a situation arises in which you require an esoteric command that you have not yet learned, you can use the built-in help to find what you need.

And, of course, you always have this book.

X Window and X Clients

There is one special type of client/server system, called X Window, that you may need to know about. If your computer uses X Window, this section will explain

the basic concepts you should understand. As you will see, X Window gives you certain advantages when you use the Internet.

If you are not an X Window user, feel free to move on to the next topic. Still, X Window is used widely, and you may want to skim this section to see what all the fuss is about.

X Window is used with Unix systems that support graphical user interfaces. A *graphical user interface*, or *GUI*, allows you to operate a computer by using not only your keyboard, but also a mouse or some other type of pointing device. With the help of your mouse, you select items from menus and manipulate objects on the screen. You can run more than one program at the same time, each of which can reside in its own rectangular area called a *window*.

Of course, the idea of a GUI (pronounced "goo-ee") is nothing new. If you have ever used a Macintosh, or a PC running Microsoft Windows or OS/2, then you have used a GUI.

As we mentioned in Chapter 1, there are many different types of Unix. X Window was developed in order to provide a standard set of tools for programmers who develop graphical applications and a standard interface for users to interact with those applications. For convenience, X Window is usually referred to as X. For example, a friend tells you that he has a program you might want to use. You might ask him, "Does it run under X?"

Although there is nothing in the design of X that says it must be used with Unix, in practice, this is almost always the case.

Hint

At one time, an operating system named V (the letter V) was developed at Stanford University. To work with V, the programmers created a windowing system which they named W. Later, the W system was sent to someone at MIT who used it for a new windowing system which he named X. The modern X Window system, first developed in 1984, grew out of this initial effort. Today, the current version of X Window is version 11 release 5, usually written as "X11 R5". Note that the official name of this system is the singular "X Window", not the plural "X Windows".

In X terminology, the three devices you use to interact with your computer—the keyboard, the screen, and the mouse—are referred to as a *display*. X lets you run more than one program at the same time on a single display. As part of the graphical user interface, each program resides in its own window on the screen. When you want to switch from one program to another, you use the mouse to move from one window to another.

When you use X, the details of maintaining the GUI for all the programs that are running are handled by a single program called a *display server* or an *X server*.

For example, say that you have four programs running at the same time, each of which resides in its own window. As you work, you can move the windows or even change their size. Now, suppose one of the programs needs to draw a circle on the screen. Rather than doing the work itself, it sends a message to the X server—the program that is controlling the screen—telling it to draw a circle of a particular size at a particular location. The X server actually does the work.

This division of labor has several important advantages. First, it means that the entire GUI is controlled by a single program that will ensure that everything works as it is supposed to. For example, the window in which a program is running may be partially obscured by another window. The program itself does not need to know this, nor does it care. The X server will handle the details.

Second, when a programmer designs a new program, he or she does not have to worry about the user interface. All that is necessary is for the program to call on X in the standard way whenever such work needs to be done. This makes for smaller, more reliable programs that are portable from one X system to the next.

Since all X servers provide the same functions, a program that is written to depend on an X server will run under any X system. You can, for instance, find an X program anywhere on the Internet, copy it to your computer, and run it under your own graphical user interface. Once you learn how to use the Internet file transfer service, you will be able to acquire many such graphical programs for free.

The third advantage of X is that the graphical user interfaces are more or less standard.

The part of the system that provides the look and feel of your interface is called the *window manager*. (Technically, the window manager is itself a program that runs on top of X.) There are two widely used window managers, named *Motif* and *Open Look*. There are differences between the two, but they are not profound.

If you saw identical programs running on two computer screens, one using Motif and one using Open Look, the main thing you would notice is that the appearance of the windows and other graphical elements would be somewhat different. In addition, you would use the mouse and the menu system slightly differently.

However, the basic concepts are the same no matter which window manager you use. Certainly, the programs inside the windows would not change.

Basically, if you know how to use one X system, you know how to use them all.

Figure 2-4 shows the screen of an X display in which you are running five different programs at the same time, each one in its own window. Notice that, as these programs execute, they depend on a single X server program to maintain the user interface.

By now, this should all remind you of the client/server relationship that we discussed in the previous section. Indeed, this is the case. The programs you run are clients and, as such, are referred to as *X clients*. They request