

# Chapter 1: Introducing Linux

---

## *In This Chapter*

- ✓ Explaining what Linux is
- ✓ Going over what Linux distributions typically include
- ✓ Discovering what Linux helps you manage
- ✓ Getting started

**I** bet you've heard about Linux. If you're wondering what exactly Linux is and what it can help you do, this chapter is all about answering those questions. Here I provide a broad-brushstroke picture of Linux and tell you how you can start using it right away.

By the way, this book covers Linux for Intel 80x86 and Pentium processors (basically any PC that can run any flavor of Windows).

## *What Is Linux?*

You know that your PC is a bunch of *hardware* — things you can touch, like the system box, monitor, keyboard, and mouse. The system box contains the most important hardware of all — the *central processing unit* (CPU), the microchip that runs the *software* (any program that tells the computer how to do your bidding), which you actually *can't* touch. In a typical Pentium 4 PC, the Pentium 4 microprocessor is the CPU. Other important hardware in the system box includes the memory (RAM chips) and the hard drive — and one program has to run all this stuff and get it to play nice: the operating system.

The *operating system* is software that manages all the hardware and runs other software at your command. You, the user, provide those commands by clicking menus and icons or by typing some cryptic text. Linux is an operating system — as are UNIX, Windows 98, Windows 2000, and Windows XP. The Linux operating system is modeled after UNIX; in its most basic, no-frills form, the Linux operating system also goes by the name *Linux kernel*.

The operating system is what gives a computer — any computer — its personality. For example, you can run Windows 98 or Windows XP on a PC — and on that same PC, you can *also* install and run Linux. That means, depending on which operating system is installed and running at any particular time, *the same PC* can be a Windows 98, Windows XP, or Linux system.

## Does Linux really run on any computer?

Linux runs on many different types of computer systems — and it does seem able to run on nearly any type of computer. Linus Torvalds and other programmers originally developed Linux for the Intel 80x86 (and compatible) line of processors. Nowadays, Linux is also available for systems based on other processors — such as those with AMD’s 64-bit AMD64 processors, the Motorola 68000 family; Alpha AXP’s; Sun SPARC’s

and UltraSPARC’s; Hewlett-Packard’s HP PA-RISC; the PowerPC and PowerPC64 processors; and the MIPS R4x00 and R5x00. More recently, IBM has released its own version of Linux for its S/390 mainframe. This book covers Linux for Intel 80x86 and Pentium processors (these have in common a basic physical structure known as *IA-32 architecture*).

The primary job of an operating system is to load software (computer programs) from the hard drive (or other permanent storage) into the memory and get the CPU to run those programs. Everything you do with your computer is possible because of the operating system — so if the operating system somehow messes up, the whole system freezes up. You know how infuriating it is when your favorite operating system — maybe even the one that came with your PC — suddenly calls it quits just as you were about to click the Send button after composing that long e-mail to your friend. You try the three-finger salute (pressing Ctrl+Alt+Del), but nothing happens. Then it’s time for the Reset button (provided your computer’s builders were wise enough to include one). Luckily, that sort of thing almost never happens with Linux — it has a reputation for being a very reliable operating system.



In technical mumbo jumbo, Linux is a *multiuser, multitasking operating system*. All this means is that Linux enables multiple users to log in, and Linux can run more than one program at the same time. Nearly all operating systems are multiuser and multitasking these days, but when Linux first started in 1994, *multiuser* and *multitasking* were big selling points.

## Linux distributions

A *Linux distribution* consists of the Linux *kernel* (the operating system) and a collection of applications, together with an easy-to-use installation program. By the way, most people just say *Linux* to refer to a specific Linux distribution.

You find many Linux distributions, and each includes the standard Linux operating system and the following major packages:

- ♦ **The X Window System:** The graphical user interface.
- ♦ **One or more graphical desktops:** Among the most popular are GNOME and KDE.

- ◆ **A selection of applications:** Linux programs come in the form of ready-to-run software, but the *source code* (the commands we humans use to tell the computer what to do) is included (or easily available), as is its documentation.

Current Linux distributions include a huge selection of software — so much that it usually requires multiple CD-ROMs or a single DVD-ROM (which this book includes).



The development and maintenance of the Linux kernel, software packages in a Linux distribution, and the Linux distributions themselves are organized as open source projects. In a nutshell, *open source* means access to the source code and the right to freely redistribute the software without any restrictions. There's a lot more to the definition than my succinct note. To find out more about the details of what open source means and the acceptable open source licenses, you can visit the Open Source Initiative Web site at [www.opensource.org](http://www.opensource.org).

Table 1-1 lists a few major Linux distributions along with a brief description for each. There are many more Linux distributions besides the ones I show in Table 1-1.



To find out more about Linux distributions, visit DistroWatch.com at [www.distrowatch.com](http://www.distrowatch.com). At that Web site, you can read up on specific distributions as well as find links for ordering CDs for specific distributions.

Table 1-1

Major Linux Distributions

<i>Distribution</i>	<i>Description</i>
Debian GNU/Linux	This non-commercial distribution started in 1993 and continues to be a popular distribution with many volunteer developers around the world contributing to the project. The installation is harder because in addition to the basic steps you'd expect, you also need to provide information about your PC's hardware. However, after you have installed the basic Debian system, you can install and upgrade Debian packages easily with a package installer called <code>apt-get</code> where <code>apt</code> stands for the Advanced Package Tool. Debian is available free of charge from <a href="http://www.debian.org">www.debian.org</a> .
Fedora Core	This distribution is the successor to Red Hat Linux — the Linux distribution from Red Hat. Fedora Core 1, released in November 2003, was the successor to the Red Hat Linux 9. Fedora Core is freely available. Fedora Core uses the Red Hat Package Manager (RPM) format for its software packages. You can download Fedora Core from <a href="http://fedora.redhat.com">fedora.redhat.com</a> .

(continued)

**Table 1-1 (continued)**

<i>Distribution</i>	<i>Description</i>
Gentoo Linux	This is a non-commercial, source-based (meaning that all software is provided in source code form) distribution that first appeared in 2002. The installer provides some binary packages to get the Linux going, but the idea is to compile all source packages on the user's computer. This makes it time-consuming to build a full-fledged Gentoo system with the latest graphical desktops, multimedia, and development tools because all the packages have to be downloaded and compiled. Gentoo Linux is freely available from <a href="http://www.gentoo.org">www.gentoo.org</a> .
Knoppix	This Live CD distribution is based on Debian and named after its developer Klaus Knopper of Germany. Knoppix can be used as a recovery tool (to fix problems with an already-installed Linux system) because you can run Knoppix directly from the CD without having to first install it on the hard drive. The Knoppix CD stores software in compressed format and Knoppix decompresses the programs on the fly. Using this approach, Knoppix can pack up to 2GB of software on a CD. Knoppix uses the Debian package management. For information on downloading Knoppix free of charge, visit the Knoppix Web site at <a href="http://www.knopper.net/knoppix/index-en.html">www.knopper.net/knoppix/index-en.html</a> .
Linspire	This commercial distribution was first released in 2002 under the name LindowsOS. Linspire uses the Debian package format and offers software download, for a fee, through what it calls the Click-N-Run Web-based interface. You can download a Live CD version called LindowsLive! via BitTorrent, a peer-to-peer file sharing system. For more information about Linspire, visit <a href="http://www.linspire.com">www.linspire.com</a> .
Mandrakelinux	This popular distribution began life as a 1998 release of Red Hat Linux with an easy-to-use installer and with KDE as the default desktop. Mandrake Linux is freely available. Mandrake software packages use the Red Hat Package Manager (RPM) format. You can download Mandrake Linux from <a href="http://www.mandrakelinux.com">www.mandrakelinux.com</a> .
MEPIS Linux	This Debian-based Live CD distribution was first released in July 2003. It also includes a graphical installer that can be launched from the CD to install MEPIS on the hard drive. MEPIS has good hardware detection and it comes with Java and multimedia software, which makes it popular. MEPIS uses the Debian package format. You can download slightly older versions of MEPIS free of charge from <a href="http://www.mepis.org">www.mepis.org</a> .
Slackware Linux	This is one of the oldest distributions, having been first released in 1992. Slackware uses compressed tar files for its packages and provides a text-based installer with limited automatic detection of hardware. All software configurations are done by editing text files. Slackware is freely available from <a href="http://www.slackware.com">www.slackware.com</a> .

<i>Distribution</i>	<i>Description</i>
SUSE Linux	This commercial distribution focuses on the desktop and includes some proprietary components that prevent its redistribution. SUSE comes with the YaST installation and configuration tool. SUSE Linux uses Red Hat Package Manager (RPM) packages. Although SUSE does not provide ISO image files for SUSE Linux, you can install it via FTP over the Internet. Also, a Live CD version of SUSE called SUSE Live Eval is available free of charge from the SUSE Web site. Visit <a href="http://www.suse.com">www.suse.com</a> for more information about SUSE Linux.
Xandros Desktop	This distribution is the successor to Corel Linux and is based on Debian. Xandros is aimed at first-time Linux users with an installer that can repartition the hard drive. The versatile Xandros File Manager is a key selling point of Xandros. However, Xandros includes some proprietary components that prevent redistribution. A limited version of Xandros Desktop called the Open Circulation Edition is available for free through BitTorrent. Visit <a href="http://www.xandros.com">www.xandros.com</a> for more information about Xandros.

As you can see from the brief descriptions in Table 1-1, some of the Linux distributions such as Knoppix and MEPIS are in the form of Live CDs. A *Live CD* includes a Linux kernel that you can boot and run directly from the CD, without having to first install it on your hard drive. Such Live CD distributions can be handy if you want to try out a distribution before you decide whether to install it or not.

Many Linux distributions, such as SUSE Linux and Xandros Desktop, are commercial products that you can buy online or in computer stores and bookstores. If you have heard about *open source* and the *GNU (GNU's Not UNIX)* license, you may think that no one can sell Linux for profit. Luckily for companies that sell Linux distributions, the GNU license — also called the GNU General Public License (GPL) — does allow commercial, for-profit distribution, but requires that the software be distributed in source-code form, and stipulates that anyone may copy and distribute the software in source-code form to anyone else. Several Linux distributions are available free of charge under the GPL, which means that the publisher may include these distributions on a DVD-ROM with this book and that you may make as many copies of the DVD as you like.

## ***Making sense of version numbers***

Both the Linux kernel and a Linux distribution have their own version numbers, not to mention the many other software programs (such as GNOME and KDE) that come with the Linux distribution. The version numbers for the Linux kernel and the Linux distribution are unrelated, but each has particular significance.

### ***Linux-kernel version numbers***

After Linux kernel version 1.0 was released on March 14, 1994, the loosely knit Linux development community adopted a version-numbering scheme. Version numbers such as 1.X.Y and 2.X.Y, where X is an even number, are considered the stable versions. The last number, Y, is the patch level, which is incremented as problems are fixed. For example, 2.6.7 is a typical, stable version of the Linux kernel. Notice that these version numbers are in the form of three integers separated by periods — *Major.Minor.Patch* — where *Major* and *Minor* are numbers denoting the major and minor version numbers, and *Patch* is another number representing the patch level.

Version numbers of the form 2.X.Y with an odd X number are beta releases for developers only; they may be unstable, so you should not adopt such versions for day-to-day use. For example, when you look at version 2.5.75 of the Linux kernel, notice the 5 — that tells you it's a beta release. Developers add new features to these odd-numbered versions of Linux.

You can find out about the latest version of the Linux kernel online at [www.kernel.org](http://www.kernel.org).

### ***Distribution-specific version numbers***

Each Linux distribution has a version number as well. These version numbers are usually of the form X.Y, where X is the major version and Y the minor version. Nowadays, if the minor version number is zero, it's simply dropped — as in Fedora Core 1 and Fedora Core 2. Unlike with the Linux-kernel version numbers, no special meaning is associated with odd and even minor versions. Each version of a Linux distribution includes specific versions of the Linux kernel and other major components, such as GNOME, KDE, and various applications.

The developers of active Linux distributions usually release new versions of their distribution on a regular basis — every six months or so. For example, SUSE Linux 9.0 was released in October 2003 and 9.1 was released in April 2004. Typically, each new major version of a Linux distribution provides significant new features.



Debian always has at least three releases at any time — *unstable*, *testing*, and *stable*. The *stable* release is the latest officially released distribution that most users would use. The *unstable* version is the distribution being worked on by developers. The *testing* distribution contains packages that have gone through some testing, but are not ready for inclusion in the stable release yet.

### ***Linux Standard Base (LSB)***

Linux has become important enough that there is a standard for Linux called the Linux Standard Base (or LSB, for short). LSB is a set of binary standards

that should help reduce variations among the Linux distributions and promote portability of applications. The idea behind LSB is to provide application binary interface (ABI) so that software applications can run on any Linux (or other UNIX) systems that conform to the LSB standard. The LSB specification references the POSIX standards as well as many other standards such as the C and C++ programming language standards, the X Window System version 11 release 6 (X11R6), and the Filesystem Hierarchy Standard (FHS). LSB version 1.2 (commonly referred to as LSB 1.2) was released on June 28, 2002. LSB 1.3 came out in January 2003, and LSB 2.0 was released on August 30, 2004.

The LSB specification is organized into two parts — a common specification that remains the same across all types of processors and a set of hardware-specific specifications, one for each type of processor architecture. For example, LSB 1.2 has architecture-specific specifications for Intel 32-bit (IA32) and Power PC 32-bit (PPC32) processors. LSB 1.3 adds a specification for the Intel 64-bit (IA64) architecture and IBM S390 processors, in addition to the ones for IA32 and PPC32. LSB 2.0 includes specification for the AMD 64-bit (AMD64) processors.

An LSB certification program exists, and by now, several Linux distributions (such as Red Hat Linux 9, Red Hat Enterprise Linux 3 for x86, SUSE Linux 9.1, and Sun Wah Linux Desktop 3.0) are certified to be LSB 1.3 compliant IA32 runtime environments. You can expect more distributions to be LSB 2.0 certified in the near future.

To discover more about LSB, visit [www.linuxbase.org](http://www.linuxbase.org). The latest list of LSB-certified systems is available at [www.opengroup.org/lsb/cert/cert\\_prodlist.tpl](http://www.opengroup.org/lsb/cert/cert_prodlist.tpl).

## What a Linux Distribution Includes

A Linux distribution comes with the Linux kernel and a whole lot more software. These software packages include everything from the graphical desktops to Internet servers to programming tools to create new software. In this section, I briefly describe some major software packages that come bundled with typical Linux distributions. Without this bundled software, Linux wouldn't be as popular as it is today.

### GNU software

At the heart of a Linux distribution is a collection of software that came from the GNU Project. You get to know these GNU utilities only if you use your Linux system through a text terminal (or a graphical window that mimics one) — a basic *command-line interface* that puts nothing much on-screen but a prompt at which you type in your commands. The GNU software is one of the basic parts of any Linux distribution.

## What is the GNU Project?

GNU is a recursive acronym that stands for *GNU's Not UNIX*. The GNU Project was launched in 1984 by Richard Stallman to develop a complete UNIX-like operating system. The GNU Project developed nearly everything needed for a complete operating system except for the operating system kernel. All GNU software was distributed under the GNU General Public License (GPL). GPL essentially requires that the software is distributed in source-code form and stipulates that any user may copy, modify, and

distribute the software to anyone else in source-code form. Users may, however, have to pay for their individual copies of GNU software.

The Free Software Foundation (FSF) is a tax-exempt charity that raises funds for work on the GNU Project. To find out more about the GNU Project, visit its home page at [www.gnu.org](http://www.gnu.org). You can find information about how to contact the Free Software Foundation and how to help the GNU Project.

As a Linux user, you may not realize the extent to which all Linux distributions rely on GNU software. Nearly all the tasks you perform in a Linux system involve one or more GNU software packages. For example, the GNOME graphical user interface (GUI) and the command interpreter (that is, the Bash shell) are both GNU software programs. By the way, the *shell* is the command-interpreter application that accepts the commands you type and then runs programs in response to those commands. If you rebuild the kernel or develop software, you do so with the GNU C and C++ compiler (which is part of the GNU software that accompanies Linux). If you edit text files with the `ed` or `emacs` editor, you're again using a GNU software package. The list goes on and on.



Table 1-2 lists some of the well-known GNU software packages that come with most Linux distributions. I show this table only to give you a feel for all the different kinds of things you can do with GNU software. Depending on your interests, you may never need to use many of these packages, but knowing they are there in case you ever need them is good.

**Table 1-2 Well-Known GNU Software Packages**

<i>Software Package</i>	<i>Description</i>
Autoconf	Generates shell scripts that automatically configure source-code packages
Automake	Generates <code>Makefile.in</code> files for use with Autoconf
Bash	The default shell — command interpreter — in Linux
Bc	An interactive calculator with arbitrary precision numbers



<i>Software Package</i>	<i>Description</i>
Binutils	A package that includes several utilities for working with binary files: <code>ar</code> , <code>as</code> , <code>gasp</code> , <code>gprof</code> , <code>ld</code> , <code>nm</code> , <code>objcopy</code> , <code>objdump</code> , <code>ranlib</code> , <code>readelf</code> , <code>size</code> , <code>strings</code> , and <code>strip</code>
Coreutils	A package that combines three individual packages called <code>Fileutils</code> , <code>Shellutils</code> , and <code>Textutils</code> and implements utilities such as <code>chgrp</code> , <code>chmod</code> , <code>chown</code> , <code>cp</code> , <code>dd</code> , <code>df</code> , <code>dir</code> , <code>dircolors</code> , <code>du</code> , <code>install</code> , <code>ln</code> , <code>ls</code> , <code>mkdir</code> , <code>mkfifo</code> , <code>mknod</code> , <code>mv</code> , <code>rm</code> , <code>rmdir</code> , <code>sync</code> , <code>touch</code> , <code>vdir</code> , <code>basename</code> , <code>chroot</code> , <code>date</code> , <code>dirname</code> , <code>echo</code> , <code>env</code> , <code>expr</code> , <code>factor</code> , <code>false</code> , <code>groups</code> , <code>hostname</code> , <code>id</code> , <code>logname</code> , <code>nice</code> , <code>nohup</code> , <code>pathchk</code> , <code>printenv</code> , <code>printf</code> , <code>pwd</code> , <code>seq</code> , <code>sleep</code> , <code>stty</code> , <code>su</code> , <code>tee</code> , <code>test</code> , <code>true</code> , <code>tty</code> , <code>uname</code> , <code>uptime</code> , <code>users</code> , <code>who</code> , <code>whoami</code> , <code>yes</code> , <code>cut</code> , <code>join</code> , <code>nl</code> , <code>split</code> , <code>tail</code> , <code>wc</code> , and so on
Gnuchess	A chess game
GNU C Library	For use with all Linux programs
Cpio	Copies file archives to and from disk or to another part of the file system
Diff	Compares files, showing line-by-line changes in several different formats
Ed	A line-oriented text editor
emacs	An extensible, customizable full-screen text editor and computing environment
Findutils	A package that includes the <code>find</code> , <code>locate</code> , and <code>xargs</code> utilities
Finger	A utility program designed to enable users on the Internet to get information about one another
Gawk	The GNU Project's implementation of the AWK programming language
GCC	Compilers for C, C++, Objective C, and other languages
Gdb	Source-level debugger for C, C++, and Fortran
Gdbm	A replacement for the traditional <code>dbm</code> and <code>ndbm</code> database libraries
Gettext	A set of utilities that enables software maintainers to internationalize (that means make the software work with different languages such as English, French, Spanish, and so on) a software package's user messages
Ghostscript	An interpreter for the Postscript and Portable Document Format (PDF) languages
Ghostview	An X Window System application that makes Ghostscript accessible from the GUI, enabling users to view Postscript or PDF files in a window
The GIMP	The GNU Image Manipulation Program is an Adobe Photoshop-like image-processing program

(continued)

**Table 1-2 (continued)**

<i>Software Package</i>	<i>Description</i>
GNOME	Provides a graphical user interface (GUI) for a wide variety of tasks that a Linux user may perform
Gnumeric	A graphical spreadsheet (similar to Microsoft Excel) that works in GNOME
grep package	Includes the <code>grep</code> , <code>egrep</code> , and <code>fgrep</code> commands that are used to find lines that match a specified text pattern
Groff	A document-formatting system similar to <code>troff</code>
GTK+	A GUI toolkit for the X Window System (used to develop GNOME applications)
Gzip	A GNU utility for compressing and decompressing files
Indent	Formats C source code by indenting it in one of several different styles
Less	A page-by-page display program similar to <code>more</code> , but with additional capabilities
Libpng	A library for image files in the Portable Network Graphics (PNG) format
m4	An implementation of the traditional UNIX macro processor
Make	A utility that determines which files of a large software package need to be recompiled, and issues the commands to recompile them
Mtools	A set of programs that enables users to read, write, and manipulate files on a DOS file system (typically a floppy disk)
Ncurses	A package for displaying and updating text on text-only terminals
Patch	A GNU version of Larry Wall's program to take the output of <code>diff</code> and apply those differences to an original file to generate the modified version
RCS	The Revision Control System is used for version control and management of source files in software projects
Sed	A stream-oriented version of the <code>ed</code> text editor
Sharutils	A package that includes <code>shar</code> (used to make shell archives out of many files) and <code>unshar</code> (to unpack these shell archives)
Tar	A tape archiving program that includes multivolume support; the capability to archive sparse files (files with big chunks of data that are all zeros), handle compression and decompression, and create remote archives; and other special features for incremental and full backups
Texinfo	A set of utilities that generates printed manuals, plain ASCII text, and online hypertext documentation (called Info), and enables users to view and read online Info documents
Time	A utility that reports the user, system, and actual time that a process uses

## GUIs and applications

Face it — typing cryptic Linux commands on a terminal is boring. For average users, using the system through a *graphical user interface* (GUI, pronounced “gooey”) — one that gives you pictures to click and windows (small *w*) to open — is much easier. This is where the X Window System, or X, comes to the rescue.

X is kind of like Microsoft Windows, but the underlying details of how X works is completely different from Windows. Unlike Windows, X provides the basic features of displaying windows on-screen, but it does not come with any specific look or feel for graphical applications. That look and feel comes from GUIs, such as GNOME and KDE, which make use of the X Window System.

Most Linux distributions come with the X Window System in the form of XFree86 or X.Org X11 — implementations of the X Window System for x86 systems. XFree86 and X.Org X11 work with a wide variety of video cards available for today’s PCs.



Until recently, XFree86 from the XFree86 Project ([www.xfree86.org](http://www.xfree86.org)) was the most commonly used X Window System implementation for x86 systems. However, around version 4.4, some changes to the XFree86 licensing terms caused concerns to many Linux and UNIX vendors — they felt that the licensing terms were no longer compatible with the GNU General Public License (GPL). In January 2004, several vendors formed the X.Org Foundation ([www.x.org](http://www.x.org)) to promote continued development of an open source X Window System and graphical desktop. The first release of X.Org X11 uses the same code that was used by XFree86 4.4, up until the time when the XFree86 license changes precipitated the creation of X.Org Foundation.

As for the GUI, Linux distributions include one or both of two powerful GUI desktops: KDE (K Desktop Environment) and GNOME (GNU Object Model Environment). If both GNOME and KDE are installed on a PC, you can choose which desktop you want as the default — or switch between the two. KDE and GNOME provide desktops similar to those of Microsoft Windows and the Mac OS. GNOME also comes with the Nautilus graphical shell that makes finding files, running applications, and configuring your Linux system easy. With GNOME or KDE, you can begin using your Linux workstation without having to know cryptic Linux commands. However, if you ever need to use those commands directly, all you have to do is open a terminal window and type them at the prompt.

Linux also comes with many graphical applications. The most noteworthy program is the GIMP (GNU Image Manipulation Program), a program for working with photos and other images. The GIMP’s capabilities are on a par with Adobe Photoshop.

Providing common productivity software — such as word-processing, spreadsheet, and database applications — is an area in which Linux used to be lacking. This situation has changed, though. Linux comes with the OpenOffice.org office productivity applications. In addition, you may want to check out these prominent, commercially available office productivity applications for Linux that are not included on this book's companion DVD-ROM:

- ◆ **Applicxware Office:** Now called Anyware Desktop for Linux, this office package is a good example of productivity software for Linux. You can find it at [www.vistasource.com](http://www.vistasource.com).
- ◆ **StarOffice:** From Sun Microsystems ([www.sun.com/staroffice](http://www.sun.com/staroffice)), StarOffice is another well-known productivity software package.
- ◆ **CrossOver Office:** From CodeWeavers ([www.codeweavers.com/site/products](http://www.codeweavers.com/site/products)), you can use CrossOver Office to install your Microsoft Office applications (Office 97, Office 2000, and Office XP) in Linux.

As you can see, there's no shortage of Linux office applications that are compatible with Microsoft Office.

## Networks

Linux comes with everything needed to use the system in networks so that the system can exchange data with other systems. On networks, computers that exchange data have to follow well-defined rules or protocols. A *network protocol* is a method that the sender and receiver agree upon for exchanging data across a network. Such a protocol is similar to the rules you might follow when you're having a polite conversation with someone at a party. You typically start by saying hello, exchanging names, and then taking turns talking. That's about the same way network protocols work. The two computers use the protocol to send bits and bytes back and forth across the network.

One of the most well known and popular network protocols is Transmission Control Protocol/Internet Protocol (TCP/IP). TCP/IP is the protocol of choice on the Internet — the “network of networks” that now spans the globe. Linux supports the TCP/IP protocol and any network applications that make use of TCP/IP.

## Internet servers

Some popular network applications are specifically designed to deliver information from one system to another. When you send electronic mail (e-mail) or visit Web sites using a Web browser, you use these network applications (also called Internet services). Here are some common Internet services:

- ◆ Electronic mail (e-mail) that you use to send messages to any other person on the Internet using addresses like [joe@someplace.com](mailto:joe@someplace.com).
- ◆ World Wide Web (or simply, Web) that you browse using a Web browser.

- ◆ News services, where you can read newsgroups and post news items to newsgroups with names such as `comp.os.linux.networking` or `comp.os.linux.setup`.
- ◆ File-transfer utilities that you can use to upload and download files.
- ◆ Remote login that you can use to connect to and work with another computer (the remote computer) on the Internet — assuming you have the required username and password to access that remote computer.

Any Linux PC can offer these Internet services. To do so, the PC must be connected to the Internet and it must run special server software called *Internet servers*. Each of the servers uses a specific protocol for transferring information. For example, here are some common Internet servers that you find in Linux:

- ◆ `sendmail` is the mail server for exchanging e-mail messages between systems using SMTP (Simple Mail Transfer Protocol).
- ◆ Apache `httpd` is the Web server for sending documents from one system to another using HTTP (Hypertext Transfer Protocol).
- ◆ `vsftpd` is the server for transferring files between computers on the Internet using FTP (File Transfer Protocol).
- ◆ `innd` is the news server for distribution of news articles in a store-and-forward fashion across the Internet using NNTP (Network News Transfer Protocol).
- ◆ `in.telnetd` allows a user on one system to log in to another system on the Internet using the TELNET protocol.
- ◆ `sshd` allows a user on one system to securely log in to another system on the Internet using the SSH (Secure Shell) protocol.

## Software development

Linux is particularly well suited to software development. Straight out of the box, it's chock-full of software-development tools such as the compiler and libraries of code needed to build programs. If you happen to know UNIX and the C programming language, you will feel right at home programming in Linux.

As far as the development environment goes, Linux has the same basic tools (such as an editor, a compiler, and a debugger) that you might use on other UNIX workstations, such as those from IBM, Sun Microsystems, and Hewlett-Packard (HP). What this means is that if you work by day on one of these UNIX workstations, you can use a Linux PC in the evening at home to duplicate that development environment at a fraction of the cost. Then you can either complete work projects at home or devote your time to software you write for fun and to share on the Internet.



Just to give you a sense of Linux's software-development support, here's a list of various features that make Linux a productive software-development environment:

- ◆ GNU C compiler, `gcc`, can compile ANSI-standard C programs.
- ◆ GNU C++ compiler (`g++`) supports ANSI-standard C++ features.
- ◆ The GNU compiler for Java (`gcj`) can compile programs written in the Java programming language.
- ◆ The GNU `make` utility enables you to compile and link large programs.
- ◆ The GNU debugger, `gdb`, enables you to step through your program to find problems and to determine where and how a program failed. (The failed program's memory image is saved in a file named `core`; `gdb` can examine this file.)
- ◆ The GNU profiling utility, `gprof`, enables you to determine the degree to which a piece of software uses your computer's processor time.
- ◆ Subversion, Concurrent Versions System (CVS), and Revision Control System (RCS) maintain version information and control access to the source files so that two programmers don't inadvertently modify the same source file at the same time.
- ◆ The GNU `emacs` editor prepares source files and even launches a compile-link process to build the program.
- ◆ Perl is a scripting language that you can use to write scripts to accomplish a specific task, tying together many smaller programs with Linux commands.
- ◆ The Tool Command Language and its graphical toolkit (`Tcl/Tk`) enable you to build graphical applications rapidly.
- ◆ Python is an interpreted programming language comparable to Perl and Tcl (for example, the Fedora Core installation program, called `anaconda`, is written in Python).
- ◆ Dynamically linked shared libraries allow your actual program files to be much smaller because all the library code that several programs may need is shared — with only one copy loaded in the system's memory.

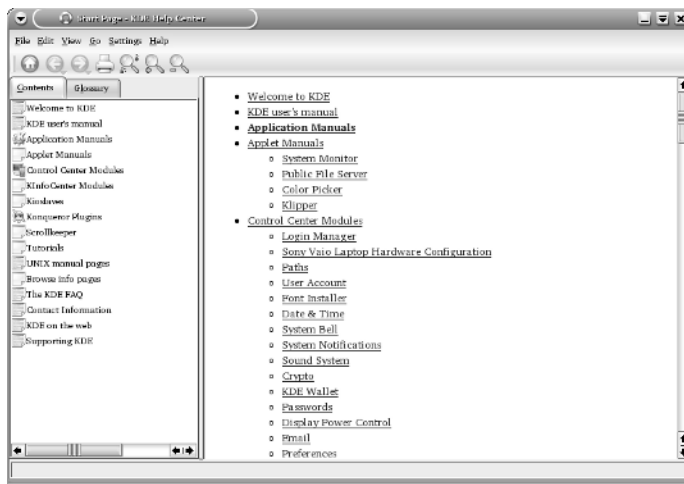
## ***Online documentation***

As you become more adept at using Linux, you may want to look up information quickly — without having to turn the pages of (ahem) this great book, for example. Luckily, Linux comes with enough online information to jog your memory in those situations when you vaguely recall a command's name, but can't remember the exact syntax of what you're supposed to type.

If you use Linux commands, you can view the manual page — commonly referred to as the *man page* — for a command by using the `man` command. (You do have to remember that command in order to access online help.)

You can also get help from the GUI desktops. Both GNOME and KDE desktops come with help viewers to view online help information. In both GNOME and KDE, select Main Menu→Help. You can then browse the help information by clicking the links on the initial help window. Figure 1-1 shows a typical help window — this one from a KDE desktop.

**Figure 1-1:**  
Online help  
is available  
from the GUI  
desktops.



## What Linux Helps You Manage

As an operating system, Linux acts as the intermediary through which you, as the “lord of the system,” manage all the hardware. The hardware includes the system box, the monitor, the keyboard, the mouse, and anything else connected to the system box. The catchall term *peripheral* refers to any equipment attached to the system. If you use a laptop computer, all your hardware is packaged into the laptop.

Inside that system box is the system’s brain — the microprocessor (Intel Pentium 4, for example) or the central processing unit (CPU) — that performs the instructions contained in a computer program. When the microprocessor is running a computer program, that program’s instructions are stored in the memory or RAM. RAM stands for *Random Access Memory* (that means any part of the memory can be accessed randomly — in arbitrary order).

The system box has another crucial component — the hard drive (or hard disk, as it is sometimes called). The hard drive is the permanent storage space for computer programs and data. It's permanent in the sense that the contents don't disappear when you power off the PC. The hard drive is organized into files, which are in turn organized in a hierarchical fashion into directories and subdirectories (somewhat like organizing papers in folders inside the drawers of a file cabinet).

To keep a Linux system running properly, you or someone else has to make sure that the hardware is working properly and that the files are backed up regularly. There is also the matter of security — making sure that only legitimate people can access and use the system. These tasks are called *system administration*.

If you are using Linux at a big facility with many computers, a full-time system administrator probably takes care of all system administration tasks. On the other hand, if you are running Linux on a home PC, you are the system administrator. Don't let the thought frighten you. You don't have to know any magic incantations or prepare cryptic configuration files to be a system administrator. Most Linux distributions include many graphical tools that makes system administration a "point-and-click" job, just like running any other application.

## ***Disks, CD-ROMs, and DVD-ROMs***

Some Linux distributions come on a single DVD-ROM. After installation, the Linux kernel and all the applications are stored on your hard drive — which is where your PC looks first when you tell it to do something.

Typically, the hard drive is prepared to use Linux during the installation process. After that, you usually leave the hard drive alone except to back up the data stored there or (occasionally) to install new applications.

Using CD-ROMs or DVD-ROMs in Linux is easy. While you are logged in at the GNOME or KDE desktop, just pop in a CD or DVD in the drive, and the system should automatically detect the DVD/CD-ROM. Depending on the Linux distribution, either a DVD/CD-ROM icon appears on the desktop or a file manager automatically opens and displays the contents of the DVD/CD-ROM. If all else fails, you can type a simple `mount` command to associate the DVD/CD-ROM with a directory on your system. This whole process of accessing the files on a CD or a DVD from Linux is called *mounting the CD or the DVD*.

Besides the hard drive and DVD/CD-ROM drive, of course, your PC may have other drives, such as a floppy disk or Zip drive, and using those disks in Linux is also simple: You insert a disk and double-click the icon that represents the disk drive on the GUI desktop. Doing so mounts the disk so that you can begin using it.



## Peripheral devices

Anything connected to your PC is a peripheral device, and so are some components like sound cards that are installed inside the system box. You can configure and manage these peripheral devices in Linux.

One of the common peripherals is a printer, typically hooked up to the parallel port of your PC. (Many distributions come with a graphical Printer Configuration tool that you can use to configure the printer.)

Another peripheral device that needs configuration is the sound card. Most Linux distributions detect and configure sound cards, just as Windows does. However, if Linux cannot detect the sound card correctly, you may have to run a text-mode or graphical tool to configure the sound card.

Linux configures other peripheral devices such as the mouse and keyboard at the time of installation. You can pretty much leave them alone after installation.

Nowadays, PCs come with the USB (Universal Serial Bus) interface; many devices, including printers and scanners, plug into a PC's USB port. One nice feature of USB devices is that you can plug them into the USB port and unplug them at any time — the device does not have to be connected when you power up the system. These devices are called *hot plug* because you can plug in a device when the system is hot, meaning while it's running. Linux supports many hot plug USB devices. When you plug in a device into the USB port, Linux loads the correct driver and makes the device available to applications.

## File systems and sharing

The whole organization of directories and files is called the *file system*. You can, of course, manage the file system using Linux. When you browse the files from the GNOME or KDE graphical desktop, you work with the familiar folder icons.



A key task in caring for a file system is to back up important files. In Linux, you can use the `tar` program to archive one or more directories on a floppy or a Zip drive. You can even back up files on a tape (if you have a tape drive). If you have a CD burner, you can also burn a CD with the files you want to back up or save for posterity.

Linux can also share parts of the file system with other systems on a network. For example, you can use the Network File System (NFS) to share files with other systems on the network. To a user on the system, the remote system's files appear to be in a directory on the local system.

Linux also comes with the Samba package, which supports file sharing with Microsoft Windows systems. Samba makes a Linux system work just like a Windows file or print server. You can also access shared folders on other Windows systems on your network.

### *Network*

Now that most PCs are either in a local area network or connected to the Internet, you need to manage the network as well. Linux comes with a Network Configuration tool to set up the local area network. For connecting to the Internet using a modem, there is usually a GUI Internet Dialup Tool.

If you connect to the Internet using DSL (that's the fast Internet connection from the phone company) or cable modem, you need a PC with an Ethernet card that connects to the cable or DSL modem. It also means that you have to set up a local area network and configure the Ethernet card. But fortunately, these steps are typically a part of Linux installation. If you want to do the configurations later, you can by using a GUI network configuration tool.

Linux also includes tools for configuring a *firewall*, a protective buffer that helps keep your system relatively secure from anyone trying to snoop over your Internet connection. You can configure the firewall by using `iptables` commands or by running a GUI firewall configuration tool.

## *Getting Started*

Based on my personal experience in learning new subjects, I prescribe a four-step process to get started with Linux:

1. **Install** Linux on your PC.
2. **Configure** Linux so that everything works to your liking.
3. **Explore** the GUI desktops and the applications.
4. **Learn** the details of specific subjects such as Internet servers.

In the following sections, I explain this prescription a bit more.

### *Install*

Microsoft Windows comes installed on your new PC, but Linux usually doesn't. So your first hurdle is to get Linux onto your PC.

After you overcome that initial human fear of the unknown, I'll bet you find Linux fairly easy to install — but where do you *get* it in the first place? Well,

the good news is that it's free — available just for the downloading. For example, you can visit the Linux Online Web site at [www.linux.org](http://www.linux.org) and click the Download button.



Because the complete distribution is *huge* — it takes up several CDs or a single DVD — your best bet is to buy a book (such as this one) that comes with Linux on a DVD-ROM. You can then do the installation by following the instructions in the book.

Just to pique your curiosity, installation involves creating space on the hard drive for both Windows and Linux. Then a step creates the Linux partitions and installs Linux from the DVD. Along the way, you configure many items — from the Ethernet card (if any) to the X Window System.

## Configure

When you finish installing Linux, the next step is to configure individual system components (for example, the sound card and the printer) and tweak any needed settings that aren't configured during installation.

If you aren't getting a graphical login screen, for example, Linux comes with tools that help you troubleshoot that problem (typically by configuring the X Window System).

You also want to configure your GUI desktop of choice — GNOME or KDE. Each has configuration tools. You can use these tools to adjust the look and feel of the desktop (background, title fonts, even the entire color scheme).

After you're through with the configuration step, all the hardware on your system and the applications run to your liking.

## Explore

With a properly configured Linux PC at your disposal, you are ready to explore Linux itself. You can begin the exploration from the GUI desktop that you get after logging in.

Explore the GUI desktops — GNOME and KDE — and the folders and files that make up the Linux file system. You can also try out the applications from the desktop. You find office and multimedia applications and databases to explore.

Also try out the *shell* — open up a terminal window and type some Linux commands in that window. You can also explore the text editors that work in text mode. Knowing how to edit text files without the GUI just in case the GUI is not available is a good idea. At least you won't be helpless.

### ***Learn***

After you explore the Linux landscape and know what is what, you can then dig in deeper and find out more about specific subject areas. For example, you may be interested in setting up Internet servers. You can then learn the details of setting up individual servers such as `sendmail` for e-mail, Apache for a Web server, and the INN server for news.

You can choose to find out about many more areas, such as security, programming, and system administration.

Of course, you can expect this step to go on and on, even after you have your system running the way you want it — for now. After all, learning is a lifelong journey.

Bon voyage!